


Second Edition

Introduction to Computational Chemistry



Frank Jensen

 **WILEY**

Introduction to Computational Chemistry

Second Edition

Introduction to Computational Chemistry

Second Edition

Frank Jensen

Department of Chemistry, University of Southern Denmark, Odense, Denmark



John Wiley & Sons, Ltd

Copyright © 2007

John Wiley & Sons Ltd
The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England
Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The Publisher is not associated with any product or vendor mentioned in this book.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 42 McDougall Street, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 6045 Freemont Blvd, Mississauga, ONT, Canada, L5R 4J3

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging-in-Publication Data

Jensen, Frank.

Introduction to computational chemistry / Frank Jensen. – 2nd ed.

p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-470-01186-7 (cloth : alk. paper)

ISBN-10: 0-470-01186-6 (cloth : alk. paper)

ISBN-13: 978-0-470-01187-4 (pbk. : alk. paper)

ISBN-10: 0-470-01187-4 (pbk. : alk. paper)

1. Chemistry, Physical and theoretical – Data processing. 2. Chemistry, Physical and theoretical – Mathematics. I. Title.

QD455.3.E4J46 2006

541.0285 – dc22

2006023998

A catalogue record for this book is available from the British Library

ISBN-13 978-0-470-01186-7 (HB) ISBN-13 978-0-470-01187-4 (PB)

ISBN-10 0-470-01186-6 (PB) ISBN-10 0-470-01187-4 (PB)

Typeset in 10/12 Times by SNP Best-set Typesetter Ltd., Hong Kong
Printed and bound in Great Britain by Antony Rowe

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

Contents

Preface to the First Edition	xv
Preface to the Second Edition	xix
1 Introduction	1
1.1 Fundamental Issues	2
1.2 Describing the System	3
1.3 Fundamental Forces	4
1.4 The Dynamical Equation	5
1.5 Solving the Dynamical Equation	8
1.6 Separation of Variables	8
1.6.1 Separating space and time variables	10
1.6.2 Separating nuclear and electronic variables	10
1.6.3 Separating variables in general	11
1.7 Classical Mechanics	12
1.7.1 The Sun–Earth system	12
1.7.2 The solar system	13
1.8 Quantum Mechanics	14
1.8.1 A hydrogen-like atom	14
1.8.2 The helium atom	17
1.9 Chemistry	19
References	21
2 Force Field Methods	22
2.1 Introduction	22
2.2 The Force Field Energy	24
2.2.1 The stretch energy	25
2.2.2 The bending energy	27
2.2.3 The out-of-plane bending energy	30
2.2.4 The torsional energy	30
2.2.5 The van der Waals energy	34
2.2.6 The electrostatic energy: charges and dipoles	40
2.2.7 The electrostatic energy: multipoles and polarizabilities	43

2.2.8	Cross terms	47
2.2.9	Small rings and conjugated systems	48
2.2.10	Comparing energies of structurally different molecules	50
2.3	Force Field Parameterization	51
2.3.1	Parameter reductions in force fields	57
2.3.2	Force fields for metal coordination compounds	58
2.3.3	Universal force fields	62
2.4	Differences in Force Fields	62
2.5	Computational Considerations	65
2.6	Validation of Force Fields	67
2.7	Practical Considerations	69
2.8	Advantages and Limitations of Force Field Methods	69
2.9	Transition Structure Modelling	70
2.9.1	Modelling the TS as a minimum energy structure	70
2.9.2	Modelling the TS as a minimum energy structure on the reactant/ product energy seam	71
2.9.3	Modelling the reactive energy surface by interacting force field functions or by geometry-dependent parameters	73
2.10	Hybrid Force Field Electronic Structure Methods	74
	References	77

3 Electronic Structure Methods: Independent-Particle Models 80

3.1	The Adiabatic and Born–Oppenheimer Approximations	82
3.2	Self-Consistent Field Theory	86
3.3	The Energy of a Slater Determinant	87
3.4	Koopmans’ Theorem	92
3.5	The Basis Set Approximation	93
3.6	An Alternative Formulation of the Variational Problem	98
3.7	Restricted and Unrestricted Hartree–Fock	99
3.8	SCF Techniques	100
3.8.1	SCF convergence	101
3.8.2	Use of symmetry	104
3.8.3	Ensuring that the HF energy is a minimum, and the correct minimum	105
3.8.4	Initial guess orbitals	107
3.8.5	Direct SCF	108
3.8.6	Reduced scaling techniques	110
3.9	Periodic Systems	113
3.10	Semi-Empirical Methods	115
3.10.1	Neglect of Diatomic Differential Overlap Approximation (NDDO)	116
3.10.2	Intermediate Neglect of Differential Overlap Approximation (INDO)	117
3.10.3	Complete Neglect of Differential Overlap Approximation (CNDO)	117
3.11	Parameterization	118
3.11.1	Modified Intermediate Neglect of Differential Overlap (MINDO)	119
3.11.2	Modified NDDO models	119
3.11.3	Modified Neglect of Diatomic Overlap (MNDO)	121
3.11.4	Austin Model 1 (AM1)	121
3.11.5	Modified Neglect of Diatomic Overlap, Parametric Method Number 3 (PM3)	122
3.11.6	Parametric Method number 5 (PM5) and PDDG/PM3 methods	123

3.11.7	The MNDO/d and AM1/d methods	124
3.11.8	Semi Ab initio Method 1	124
3.12	Performance of Semi-Empirical Methods	125
3.13	Hückel Theory	127
3.13.1	Extended Hückel theory	127
3.13.2	Simple Hückel theory	128
3.14	Limitations and Advantages of Semi-Empirical Methods	129
	References	131

4 Electron Correlation Methods 133

4.1	Excited Slater Determinants	135
4.2	Configuration Interaction	137
4.2.1	CI Matrix elements	138
4.2.2	Size of the CI matrix	141
4.2.3	Truncated CI methods	143
4.2.4	Direct CI methods	144
4.3	Illustrating how CI Accounts for Electron Correlation, and the RHF Dissociation Problem	145
4.4	The UHF Dissociation, and the Spin Contamination Problem	148
4.5	Size Consistency and Size Extensivity	153
4.6	Multi-Configuration Self-Consistent Field	153
4.7	Multi-Reference Configuration Interaction	158
4.8	Many-Body Perturbation Theory	159
4.8.1	Møller-Plesset perturbation theory	162
4.8.2	Unrestricted and projected Møller-Plesset methods	168
4.9	Coupled Cluster	169
4.9.1	Truncated coupled cluster methods	172
4.10	Connections between Coupled Cluster, Configuration Interaction and Perturbation Theory	174
4.10.1	Illustrating correlation methods for the beryllium atom	177
4.11	Methods Involving the Interelectronic Distance	178
4.12	Direct Methods	181
4.13	Localized Orbital Methods	182
4.14	Summary of Electron Correlation Methods	183
4.15	Excited States	186
4.16	Quantum Monte Carlo Methods	187
	References	189

5 Basis Sets 192

5.1	Slater and Gaussian Type Orbitals	192
5.2	Classification of Basis Sets	194
5.3	Even- and Well-Tempered Basis Sets	198
5.4	Contracted Basis Sets	200
5.4.1	Pople style basis sets	202
5.4.2	Dunning-Huzinaga basis sets	204
5.4.3	MINI, MIDI and MAXI basis sets	205
5.4.4	Ahlrichs type basis sets	205
5.4.5	Atomic natural orbital basis sets	205
5.4.6	Correlation consistent basis sets	206

5.4.7	Polarization consistent basis sets	207
5.4.8	Basis set extrapolation	208
5.5	Plane Wave Basis Functions	211
5.6	Recent Developments and Computational Issues	212
5.7	Composite Extrapolation Procedures	213
5.8	Isogyric and Isodesmic Reactions	221
5.9	Effective Core Potentials	222
5.10	Basis Set Superposition Errors	225
5.11	Pseudospectral Methods	227
	References	229
6	Density Functional Methods	232
6.1	Orbital-Free Density Functional Theory	233
6.2	Kohn-Sham Theory	235
6.3	Reduced Density Matrix Methods	236
6.4	Exchange and Correlation Holes	240
6.5	Exchange-Correlation Functionals	243
6.5.1	Local Density Approximation	246
6.5.2	Gradient-corrected methods	248
6.5.3	Higher order gradient or meta-GGA methods	250
6.5.4	Hybrid or hyper-GGA methods	252
6.5.5	Generalized random phase methods	253
6.5.6	Functionals overview	254
6.6	Performance and Properties of Density Functional Methods	255
6.7	DFT Problems	258
6.8	Computational Considerations	260
6.9	Final Considerations	263
	References	264
7	Valence Bond Methods	268
7.1	Classical Valence Bond Theory	269
7.2	Spin-Coupled Valence Bond Theory	270
7.3	Generalized Valence Bond Theory	275
	References	276
8	Relativistic Methods	277
8.1	The Dirac Equation	278
8.2	Connections Between the Dirac and Schrödinger Equations	280
8.2.1	Including electric potentials	280
8.2.2	Including both electric and magnetic potentials	282
8.3	Many-Particle Systems	284
8.4	Four-Component Calculations	287
8.5	Relativistic Effects	289
	References	292
9	Wave Function Analysis	293
9.1	Population Analysis Based on Basis Functions	293
9.2	Population Analysis Based on the Electrostatic Potential	296

9.3	Population Analysis Based on the Electron Density	299
9.3.1	Atoms In Molecules	299
9.3.2	Voronoi, Hirshfeld and Stewart atomic charges	303
9.3.3	Generalized atomic polar tensor charges	304
9.4	Localized Orbitals	304
9.4.1	Computational considerations	306
9.5	Natural Orbitals	308
9.6	Natural Atomic Orbital and Natural Bond Orbital Analysis	309
9.7	Computational Considerations	311
9.8	Examples	312
	References	313

10 Molecular Properties 315

10.1	Examples of Molecular Properties	316
10.1.1	External electric field	316
10.1.2	External magnetic field	318
10.1.3	Internal magnetic moments	318
10.1.4	Geometry change	319
10.1.5	Mixed derivatives	319
10.2	Perturbation Methods	321
10.3	Derivative Techniques	321
10.4	Lagrangian Techniques	324
10.5	Coupled Perturbed Hartree–Fock	325
10.6	Electric Field Perturbation	329
10.6.1	External electric field	329
10.6.2	Internal electric field	329
10.7	Magnetic Field Perturbation	329
10.7.1	External magnetic field	331
10.7.2	Nuclear spin	332
10.7.3	Electron spin	333
10.7.4	Classical terms	333
10.7.5	Relativistic terms	334
10.7.6	Magnetic properties	334
10.7.7	Gauge dependence of magnetic properties	338
10.8	Geometry Perturbations	339
10.9	Response and Propagator Methods	343
10.10	Property Basis Sets	348
	References	349

11 Illustrating the Concepts 350

11.1	Geometry Convergence	350
11.1.1	Ab Initio methods	350
11.1.2	Density functional methods	353
11.2	Total Energy Convergence	354
11.3	Dipole Moment Convergence	356
11.3.1	Ab Initio methods	356
11.3.2	Density functional methods	357
11.4	Vibrational Frequency Convergence	358
11.4.1	Ab Initio methods	358
11.4.2	Density functional methods	360

11.5	Bond Dissociation Curves	361
11.5.1	Basis set effect at the Hartree–Fock level	361
11.5.2	Performance of different types of wave function	363
11.5.3	Density functional methods	369
11.6	Angle Bending Curves	370
11.7	Problematic Systems	370
11.7.1	The geometry of F ₂ O	371
11.7.2	The dipole moment of CO	372
11.7.3	The vibrational frequencies of O ₃	373
11.8	Relative Energies of C ₄ H ₆ Isomers	374
	References	378
12	Optimization Techniques	380
12.1	Optimizing Quadratic Functions	381
12.2	Optimizing General Functions: Finding Minima	383
12.2.1	Steepest descent	383
12.2.2	Conjugate gradient methods	384
12.2.3	Newton–Raphson methods	385
12.2.4	Step control	386
12.2.5	Obtaining the Hessian	387
12.2.6	Storing and diagonalizing the Hessian	388
12.2.7	Extrapolations: the GDIIS method	389
12.3	Choice of Coordinates	390
12.4	Optimizing General Functions: Finding Saddle Points (Transition Structures)	394
12.4.1	One-structure interpolation methods: coordinate driving, linear and quadratic synchronous transit, and sphere optimization	394
12.4.2	Two-structure interpolation methods: saddle, line-then-plane, ridge and step-and-slide optimizations	397
12.4.3	Multi-structure interpolation methods: chain, locally updated planes, self-penalty walk, conjugate peak refinement and nudged elastic band	398
12.4.4	Characteristics of interpolation methods	401
12.4.5	Local methods: gradient norm minimization	402
12.4.6	Local methods: Newton–Raphson	403
12.4.7	Local methods: the dimer method	405
12.4.8	Coordinates for TS searches	405
12.4.9	Characteristics of local methods	406
12.4.10	Dynamic methods	406
12.5	Constrained Optimization Problems	407
12.6	Conformational Sampling and the Global Minimum Problem	409
12.6.1	Stochastic and Monte Carlo methods	411
12.6.2	Molecular dynamics	412
12.6.3	Simulated annealing	413
12.6.4	Genetic algorithms	413
12.6.5	Diffusion methods	414
12.6.6	Distance geometry methods	414
12.7	Molecular Docking	415
12.8	Intrinsic Reaction Coordinate Methods	416
	References	419

13	Statistical Mechanics and Transition State Theory	421
13.1	Transition State Theory	421
13.2	Rice–Ramsperger–Kassel–Marcus Theory	424
13.3	Dynamical Effects	425
13.4	Statistical Mechanics	426
13.5	The Ideal Gas, Rigid-Rotor Harmonic-Oscillator Approximation	429
13.5.1	Translational degrees of freedom	430
13.5.2	Rotational degrees of freedom	430
13.5.3	Vibrational degrees of freedom	431
13.5.4	Electronic degrees of freedom	433
13.5.5	Enthalpy and entropy contributions	433
13.6	Condensed Phases	439
	References	443
14	Simulation Techniques	445
14.1	Monte Carlo Methods	448
14.1.1	Generating non-natural ensembles	450
14.2	Time-Dependent Methods	450
14.2.1	Molecular dynamics methods	451
14.2.2	Generating non-natural ensembles	454
14.2.3	Langevin methods	455
14.2.4	Direct methods	455
14.2.5	Extended Lagrange techniques (Car–Parrinello methods)	457
14.2.6	Quantum methods using potential energy surfaces	459
14.2.7	Reaction path methods	460
14.2.8	Non-Born–Oppenheimer methods	463
14.2.9	Constrained sampling methods	463
14.3	Periodic Boundary Conditions	464
14.4	Extracting Information from Simulations	468
14.5	Free Energy Methods	472
14.5.1	Thermodynamic perturbation methods	472
14.5.2	Thermodynamic integration methods	473
14.6	Solvation Models	475
14.7	Continuum Solvation Models	476
14.7.1	Poisson–Boltzmann methods	478
14.7.2	Born/Onsager/Kirkwood models	480
14.7.3	Self-consistent reaction field models	481
	References	484
15	Qualitative Theories	487
15.1	Frontier Molecular Orbital Theory	487
15.2	Concepts from Density Functional Theory	492
15.3	Qualitative Molecular Orbital Theory	494
15.4	Woodward–Hoffmann Rules	497
15.5	The Bell–Evans–Polanyi Principle/Hammond Postulate/Marcus Theory	506
15.6	More O’Ferrall–Jencks Diagrams	510
	References	512

16	Mathematical Methods	514
16.1	Numbers, Vectors, Matrices and Tensors	514
16.2	Change of Coordinate System	520
16.2.1	Examples of changing the coordinate system	525
16.2.2	Vibrational normal coordinates	526
16.2.3	Energy of a Slater determinant	528
16.2.4	Energy of a CI wave function	529
16.2.5	Computational Consideration	529
16.3	Coordinates, Functions, Functionals, Operators and Superoperators	530
16.3.1	Differential operators	531
16.4	Normalization, Orthogonalization and Projection	532
16.5	Differential Equations	535
16.5.1	Simple first-order differential equations	535
16.5.2	Less simple first-order differential equations	536
16.5.3	Simple second-order differential equations	536
16.5.4	Less simple second-order differential equations	537
16.5.5	Second-order differential equations depending on the function itself	537
16.6	Approximating Functions	538
16.6.1	Taylor expansion	539
16.6.2	Basis set expansion	541
16.7	Fourier and Laplace Transformations	541
16.8	Surfaces	543
	References	546
17	Statistics and QSAR	547
17.1	Introduction	547
17.2	Elementary Statistical Measures	549
17.3	Correlation Between Two Sets of Data	550
17.4	Correlation between Many Sets of Data	553
17.4.1	Multiple-descriptor data sets and quality analysis	553
17.4.2	Multiple linear regression	555
17.4.3	Principal component and partial least squares analysis	556
17.4.4	Illustrative example	558
17.5	Quantitative Structure–Activity Relationships (QSAR)	559
	References	561
18	Concluding Remarks	562
	Appendix A	565
	Notation	565
	Appendix B	570
	B.1 The Variational Principle	570
	B.2 The Hohenberg–Kohn Theorems	571
	B.3 The Adiabatic Connection Formula	572
	Reference	573

Appendix C	574
Atomic Units	574
Appendix D	575
Z-Matrix Construction	575
Index	583

Preface to the First Edition

Computational chemistry is rapidly emerging as a subfield of theoretical chemistry, where the primary focus is on solving chemically related problems by calculations. For the newcomer to the field, there are three main problems:

- (1) Deciphering the code. The language of computational chemistry is littered with acronyms, what do these abbreviations stand for in terms of underlying assumptions and approximations?
- (2) Technical problems. How does one actually run the program and what to look for in the output?
- (3) Quality assessment. How good is the number that has been calculated?

Point (1) is part of every new field: there is not much to do about it. If you want to live in another country, you have to learn the language. If you want to use computational chemistry methods, you need to learn the acronyms. I have tried in the present book to include a good fraction of the most commonly used abbreviations and standard procedures.

Point (2) is both hardware and software specific. It is not well suited for a text book, as the information rapidly becomes out of date. The average lifetime of computer hardware is a few years, the time between new versions of software is even less. Problems of type (2) need to be solved “on location”. I have made one exception, however, and have including a short discussion of how to make *Z*-matrices. A *Z*-matrix is a convenient way of specifying a molecular geometry in terms of internal coordinates, and it is used by many electronic structure programs. Furthermore, geometry optimizations are often performed in *Z*-matrix variables, and since optimizations in a good set of internal coordinates are significantly faster than in Cartesian coordinates, it is important to have a reasonable understanding of *Z*-matrix construction.

As computer programs evolve they become easier to use. Modern programs often communicate with the user in terms of a graphical interface, and many methods have become essential “black box” procedures: if you can draw the molecule, you can also do the calculation. This effectively means that you no longer have to be a highly trained theoretician to run even quite sophisticated calculations.

The ease with which calculations can be performed means that point (3) has become the central theme in computational chemistry. It is quite easy to run a series of calculations which produce results that are absolutely meaningless. The program will not tell you whether the chosen method is valid for the problem you are studying. Quality assessment is thus an absolute requirement. This, however, requires much more experience and insight than just running the program. A basic understanding of the theory behind the method is needed, and a knowledge of the performance of the method for other systems. If you are breaking new ground, where there is no previous experience, you need a way of calibrating the results.

The lack of quality assessment is probably one of the reasons why computational chemistry has (had) a somewhat bleak reputation. “If five different computational methods give five widely different results, what has computational chemistry contributed? You just pick the number closest to experiments and claim that you can reproduce experimental data accurately.” One commonly sees statements of the type “The theoretical results for property X are in disagreement. Calculation at the CCSD(T)/6-31G(d,p) level predicts that . . . , while the MINDO/3 method gives opposing results. There is thus no clear consent from theory.” This is clearly a lack of understanding of the quality of the calculations. If the results disagree, there is a very high probability that the CCSD(T) results are basically correct, and the MINDO/3 results are wrong. If you want to make predictions, and not merely reproduce known results, you need to be able to judge the quality of your results. This is by far the most difficult task in computational chemistry. I hope the present book will give some idea of the limitations of different methods.

Computers don’t solve problems, people do. Computers just generate numbers. Although computational chemistry has evolved to the stage where it often can be competitive with experimental methods for generating a value for a given property of a given molecule, the number of possible molecules (there are an estimated 10^{200} molecules with a molecular weight less than 850) and their associated properties is so huge that only a very tiny fraction will ever be amenable to calculations (or experiments). Furthermore, with the constant increase in computational power, a calculation that barely can be done today will be possible on medium-sized machines in 5–10 years. Prediction of properties with methods that do not provide converged results (with respect to theoretical level) will typically only have a lifetime of a few years before being surpassed by more accurate calculations.

The real strength of computational chemistry is the ability to generate data (for example by analyzing the wave function) from which a human may gain *insight*, and thereby rationalize the behaviour of a large class of molecules. Such insights and rationalizations are much more likely to be useful over a longer period of time than the raw results themselves. A good example is the concept used by organic chemists with molecules composed of functional groups, and representing reactions by “pushing electrons”. This may not be particularly accurate from a quantum mechanical point of view, but it is very effective in rationalizing a large body of experimental results, and has good predictive power.

Just as computers do not solve problems, mathematics by itself does not provide insight. It merely provides formulas, a framework for organizing thoughts. It is in this spirit that I have tried to write this book. Only the necessary (obviously a subjective criterion) mathematical background has been provided, the aim being that the reader

should be able to understand the premises and limitations of different methods, and follow the main steps in running a calculation. This means that I in many cases have omitted to tell the reader of some of the finer details, which may annoy the purists. However, I believe the large overview is necessary before embarking on a more stringent and detailed derivation of the mathematics. The goal of this book is to provide an overview of commonly used methods, giving enough theoretical background to understand why for example the AMBER force field is used for modelling proteins but MM2 is used for small organic molecules. Or why coupled cluster inherently is an iterative method, while perturbation theory and configuration interaction inherently are non-iterative methods, although the CI problem in practice is solved by iterative techniques.

The prime focus of this book is on calculating molecular structures and (relative) energies, and less on molecular properties or dynamical aspects. In my experience, predicting structures and energetics are the main uses of computational chemistry today, although this may well change in the coming years. I have tried to include most methods that are already extensively used, together with some that I expect to become generally available in the near future. How detailed the methods are described depends partly on how practical and commonly used the methods are (both in terms of computational resources and software), and partly reflects my own limitations in terms of understanding. Although simulations (e.g. molecular dynamics) are becoming increasingly powerful tools, only a very rudimentary introduction is provided in Chapter 16. The area is outside my expertise, and several excellent textbooks are already available.

Computational chemistry contains a strong practical element. Theoretical methods must be translated into working computer programs in order to produce results. Different algorithms, however, may have different behaviours in practice, and it becomes necessary to be able to evaluate whether a certain type of calculation can be carried out with the available computers. The book thus contains some guidelines for evaluating what type of resources necessary for carrying out a given calculation.

The present book grew out of a series of lecture notes that I have used for teaching a course in computational chemistry at Odense University, and the style of the book reflects its origin. It is difficult to master all disciplines in the vast field of computational chemistry. A special thanks to H. J. Aa. Jensen, K. V. Mikkelsen, T. Saue, S. P. A. Sauer, M. Schmidt, P. M. W. Gill, P.-O. Norrby, D. L. Cooper, T. U. Helgaker and H. G. Petersen for having read various parts of the book and providing input. Remaining errors are of course my sole responsibility. A good part of the final transformation from a set of lecture notes to the present book was done during a sabbatical leave spent with Prof. L. Radom at the Research School of Chemistry, Australia National University, Canberra, Australia. A special thanks to him for his hospitality during the stay.

A few comments on the layout of the book. Definitions, acronyms or common phrases are marked in *italic*; these can be found in the index. Underline is used for emphasizing important points. Operators, vectors and matrices are denoted in **bold**, scalars in normal text. Although I have tried to keep the notation as consistent as possible, different branches in computational chemistry often use different symbols for the same quantity. In order to comply with common usage, I have elected sometimes to switch notation between chapters. The second derivative of the energy, for example, is called the force constant k in force field theory, the corresponding matrix is denoted **F** when discussing vibrations, and called the Hessian **H** for optimization purposes.

I have assumed that the reader has no prior knowledge of concepts specific to computational chemistry, but has a working understanding of introductory quantum mechanics and elementary mathematics, especially linear algebra, vector, differential and integral calculus. The following features specific to chemistry are used in the present book without further introduction. Adequate descriptions may be found in a number of quantum chemistry textbooks (J. P. Lowe, *Quantum Chemistry*, Academic Press, 1993; I. N. Levine, *Quantum Chemistry*, Prentice Hall, 1992; P. W. Atkins, *Molecular Quantum Mechanics*, Oxford University Press, 1983).

- (1) The Schrödinger equation, with the consequences of quantized solutions and quantum numbers.
- (2) The interpretation of the square of the wave function as a probability distribution, the Heisenberg uncertainty principle and the possibility of tunnelling.
- (3) The solutions for the hydrogen atom, atomic orbitals.
- (4) The solutions for the harmonic oscillator and rigid rotor.
- (5) The molecular orbitals for the H_2 molecule generated as a linear combination of two s -functions, one on each nuclear centre.
- (6) Point group symmetry, notation and representations, and the group theoretical condition for when an integral is zero.

I have elected to include a discussion of the variational principle and perturbational methods, although these are often covered in courses in elementary quantum mechanics. The properties of angular momentum coupling are used at the level of knowing the difference between a singlet and triplet state. I do not believe that it is necessary to understand the details of vector coupling to understand the implications.

Although I have tried to keep each chapter as self-contained as possible, there are unavoidable dependencies. The part in Chapter 3 describing HF methods is a prerequisite for understanding Chapter 4. Both these Chapters use terms and concepts for basis sets which are treated in Chapter 5. Chapter 5, in turn, relies on concepts in Chapters 3 and 4, i.e. these three chapters form the core for understanding modern electronic structure calculations. Many of the concepts in Chapters 3 and 4 are also used in Chapters 6, 7, 9, 11 and 15 without further introduction, although these five chapters probably can be read with some benefits without a detailed understanding of Chapters 3 and 4. Chapter 8, and to a certain extent also Chapter 10, are fairly advanced for an introductory textbook, such as the present, and can be skipped. They do, however, represent areas that are probably going to be more and more important in the coming years. Function optimization, which is described separately in Chapter 14, is part of many areas, but a detailed understanding is not required for following the arguments in the other chapters. Chapters 12 and 13 are fairly self-contained, and form some of the background for the methods in the other chapters. In my own course I normally take Chapters 12, 13 and 14 fairly early in the course, as they provide background for Chapters 3, 4 and 5.

If you would like to make comments, advise me of possible errors, make clarifications, add references, etc., or view the current list of misprints and corrections, please visit the author's website (URL: <http://bogense.chem.ou.dk/~icc>).

Preface to the Second Edition

The changes relative to the first edition are as follows:

- Numerous misprints and inaccuracies in the first edition have been corrected. Most likely some new ones have been introduced in the process, please check the book website for the most recent correction list and feel free to report possible problems. Since web addresses have a tendency to change regularly, please use your favourite search engine to locate the current URL.
- The methodologies and references in each chapter have been updated with new developments published between 1998 and 2005.
- More extensive referencing. Complete referencing is impossible, given the large breadth of subjects. I have tried to include references that preferably are recent, have a broad scope and include key references. From these the reader can get an entry into the field.
- Many figures and illustrations have been redone. The use of colour illustrations has been deferred in favour of keeping the price of the book down.
- Each chapter or section now starts with a short overview of the methods, described without mathematics. This may be useful for getting a feel for the methods, without embarking on all the mathematical details. The overview is followed by a more detailed mathematical description of the method, including some key references which may be consulted for more details. At the end of the chapter or section, some of the pitfalls and the directions of current research are outlined.
- Energy units have been converted from kcal/mol to kJ/mol, based on the general opinion that the scientific world should move towards SI units.
- Furthermore, some chapters have undergone major restructuring:
 - Chapter 16 (Chapter 13 in the first edition) has been greatly expanded to include a summary of the most important mathematical techniques used in the book. The goal is to make the book more self-contained, i.e. relevant mathematical techniques used in the book are at least rudimentarily discussed in Chapter 16.

- All the statistical mechanics formalism has been collected in Chapter 13.
- Chapter 14 has been expanded to cover more of the methodologies used in molecular dynamics.
- Chapter 12 on optimization techniques has been restructured.
- Chapter 6 on density functional methods has been rewritten.
- A new Chapter 1 has been introduced to illustrate the similarities and differences between classical and quantum mechanics, and to provide some fundamental background.
- A rudimentary treatment of periodic systems has been incorporated in Chapters 3 and 14.
- A new Chapter 17 has been introduced to describe statistics and QSAR methods.
- I have tried to make the book more modular, i.e. each chapter is more self-contained. This makes it possible to use only selected chapters, e.g. for a course, but has the drawback of repeating the same things in several chapters, rather than simply cross-referencing.

Although the modularity has been improved, there are unavoidable interdependencies. Chapters 3, 4 and 5 contain the essentials of electronic structure theory, and most would include Chapter 6 describing density functional methods. Chapter 2 contains a description of empirical force field methods, and this is tightly coupled to the simulation methods in Chapter 14, which of course leans on the statistical mechanics in Chapter 13. Chapter 1 on fundamental issues is of a more philosophical nature, and can be skipped. Chapter 16 on mathematical techniques is mainly for those not already familiar with this, and Chapter 17 on statistical methods may be skipped as well.

Definitions, acronyms and common phrases are marked in *italic*. In a change from the first edition, where underlining was used, *italic* text has also been used for emphasizing important points.

A number of people have offered valuable help and criticisms during the updating process. I would especially like to thank S. P. A. Sauer, H. J. Aa. Jensen, E. J. Baerends and P. L. A. Popelier for having read various parts of the book and provided input. Remaining errors are of course my sole responsibility.

Specific comments on the preface to the first edition

Bohacek *et al.*¹ have estimated the number of possible compounds composed of H, C, N, O and S atoms with 30 non-hydrogen atoms or fewer to be 10^{60} . Although this number is so large that only a very tiny fraction will ever be amenable to investigation, the concept of functional groups means that one does not need to evaluate all compounds in a given class to determine their properties. The number of alkanes meeting the above criteria is $\sim 10^{10}$; clearly these will all have very similar and well-understood properties, and there is no need to investigate all 10^{10} compounds.

Reference

1. R. S. Bohacek, C. McMartin, W. C. Guida, *Med. Res. Rev.*, **16** (1996), 3.

1 Introduction

Chemistry is the science dealing with construction, transformation and properties of molecules. Theoretical chemistry is the subfield where mathematical methods are combined with fundamental laws of physics to study processes of chemical relevance.¹

Molecules are traditionally considered as “composed” of atoms or, in a more general sense, as a collection of charged particles, positive nuclei and negative electrons. The only important physical force for chemical phenomena is the Coulomb interaction between these charged particles. Molecules differ because they contain different nuclei and numbers of electrons, or because the nuclear centres are at different geometrical positions. The latter may be “chemically different” molecules such as ethanol and dimethyl ether, or different “conformations” of for example butane.

Given a set of nuclei and electrons, theoretical chemistry can attempt to calculate things such as:

- Which geometrical arrangements of the nuclei correspond to stable molecules?
- What are their relative energies?
- What are their properties (dipole moment, polarizability, NMR coupling constants, etc.)?
- What is the rate at which one stable molecule can transform into another?
- What is the time dependence of molecular structures and properties?
- How do different molecules interact?

The only systems that can be solved exactly are those composed of only one or two particles, where the latter can be separated into two pseudo one-particle problems by introducing a “centre of mass” coordinate system. Numerical solutions to a given accuracy (which may be so high that the solutions are essentially “exact”) can be generated for many-body systems, by performing a very large number of mathematical operations. Prior to the advent of electronic computers (i.e. before 1950), the number of systems that could be treated with a high accuracy was thus very limited. During the sixties and seventies, electronic computers evolved from a few very expensive, difficult to use, machines to become generally available for researchers all over the world. The performance for a given price has been steadily increasing since and the use of computers is now widespread in many branches of science. This has spawned a new

field in chemistry, *computational chemistry*, where the computer is used as an “experimental” tool, much like, for example, an NMR spectrometer.

Computational chemistry is focused on obtaining results relevant to chemical problems, not directly at developing new theoretical methods. There is of course a strong interplay between traditional theoretical chemistry and computational chemistry. Developing new theoretical models may enable new problems to be studied, and results from calculations may reveal limitations and suggest improvements in the underlying theory. Depending on the accuracy wanted, and the nature of the system at hand, one can today obtain useful information for systems containing up to several thousand particles. One of the main problems in computational chemistry is selecting a suitable level of theory for a given problem, and to be able to evaluate the quality of the obtained results. The present book will try to put the variety of modern computational methods into perspective, hopefully giving the reader a chance of estimating which types of problems can benefit from calculations.

1.1 Fundamental Issues

Before embarking on a detailed description of the theoretical methods in computational chemistry, it may be useful to take a wider look at the background for the theoretical models, and how they relate to methods in other parts of science, such as physics and astronomy.

A very large fraction of the computational resources in chemistry and physics is used in solving the so-called *many-body problem*. The essence of the problem is that two-particle systems can in many cases be solved exactly by mathematical methods, producing solutions in terms of analytical functions. Systems composed of more than two particles cannot be solved by analytical methods. Computational methods can, however, produce approximate solutions, which in principle may be refined to any desired degree of accuracy.

Computers are not smart – at the core level they are in fact very primitive. Smart programmers, however, can make sophisticated computer programs, which may make the computer appear smart, or even intelligent. But the basics of any computer program consist of a doing a few simple tasks such as:

- Performing a mathematical operation (adding, multiplying, square root, cosine, . . .) on one or two numbers.
- Determining the relationship (equal to, greater than, less than or equal to, . . .) between two numbers.
- Branching depending on a decision (add two numbers if $N > 10$, else subtract one number from the other).
- Looping (performing the same operation a number of times, perhaps on a set of data).
- Reading and writing data from and to external files.

These tasks are the essence of any programming language, although the syntax, data handling and efficiency depend on the language. The main reason why computers are so useful is the sheer speed with which they can perform these operations. Even a cheap off-the-shelf personal computer can perform billions (10^9) of operations per second.

Within the scientific world, computers are used for two main tasks: performing numerically intensive calculations and analyzing large amounts of data. Such data can, for example, be pictures generated by astronomical telescopes or gene sequences in the bioinformatics area that need to be compared. The numerically intensive tasks are typically related to simulating the behaviour of the real world, by a more or less sophisticated computational model. The main problem in such simulations is the multi-scale nature of real-world problems, spanning from sub-nano to millimetres ($10^{-10} - 10^{-3}$) in spatial dimensions, and from femto- to milliseconds ($10^{-15} - 10^{-3}$) in the time domain.

1.2 Describing the System

In order to describe a system we need four fundamental features:

- System description – What are the fundamental units or “particles”, and how many are there?
- Starting condition – Where are the particles and what are their velocities?
- Interaction – What is the mathematical form for the forces acting between the particles?
- Dynamical equation – What is the mathematical form for evolving the system in time?

The choice of “particles” puts limitations on what we are ultimately able to describe. If we choose atomic nuclei and electrons as our building blocks, we can describe atoms and molecules, but not the internal structure of the atomic nucleus. If we choose atoms as the building blocks, we can describe molecular structures, but not the details of the electron distribution. If we choose amino acids as the building blocks, we may be able to describe the overall structure of a protein, but not the details of atomic movements.




Figure 1.1 Hierarchy of building blocks for describing a chemical system

The choice of starting conditions effectively determines what we are trying to describe. The complete phase space (i.e. all possible values of positions and velocities for all particles) is huge, and we will only be able to describe a small part of it. Our choice of starting conditions determines which part of the phase space we sample, for example which (structural or conformational) isomer or chemical reaction we can describe. There are many structural isomers with the molecular formula C_6H_6 , but if we want to study benzene, we should place the nuclei in a hexagonal pattern, and start them with relatively low velocities.

The interaction between particles in combination with the dynamical equation determines how the system evolves in time. At the fundamental level, the only important force at the atomic level is the electromagnetic interaction. Depending on the choice of system description (particles), however, this may result in different effective forces.

In force field methods, for example, the interactions are parameterized as stretch, bend, torsional, van der Waals, etc., interactions.

The dynamical equation describes the time evolution of the system. It is given as a differential equation involving both time and space derivatives, with the exact form depending on the particle masses and velocities. By solving the dynamical equation the particles' position and velocity can be predicted at later (or earlier) times relative to the starting conditions, i.e. how the system evolves in the phase space.

1.3 Fundamental Forces

The interaction between particles can be described in terms of either a force (\mathbf{F}) or a potential (\mathbf{V}). These are equivalent, as the force is the derivative of the potential with respect to the position \mathbf{r} .

$$\mathbf{F}(\mathbf{r}) = -\frac{\partial \mathbf{V}}{\partial \mathbf{r}} \quad (1.1)$$

Current knowledge indicates that there are four fundamental interactions, at least under normal conditions, as listed in Table 1.1.

Table 1.1 Fundamental interactions

Name	Particles	Range (m)	Relative strength
Strong interaction	Quarks	$<10^{-15}$	100
Weak interaction	Quarks, leptons	$<10^{-15}$	0.001
Electromagnetic	Charged particles	∞	1
Gravitational	Mass particles	∞	10^{-40}

Quarks are the building blocks of protons and neutrons, and lepton is a common name for a group of particles including the electron and the neutrino. The strong interaction is the force holding the atomic nucleus together, despite the repulsion between the positively charged protons. The weak interaction is responsible for radioactive decay of nuclei by conversion of neutrons to protons (β decay). The strong and weak interactions are short-ranged and are only important within the atomic nucleus.

Both the electromagnetic and gravitational interactions depend on the inverse distance between the particles, and are therefore of infinite range. The electromagnetic interaction occurs between all charged particles, while the gravitational interaction occurs between all particles with a mass, and they have the same overall functional form.

$$\begin{aligned} \mathbf{V}_{\text{elec}}(\mathbf{r}_{ij}) &= C_{\text{elec}} \frac{q_i q_j}{r_{ij}} \\ \mathbf{V}_{\text{grav}}(\mathbf{r}_{ij}) &= -C_{\text{grav}} \frac{m_i m_j}{r_{ij}} \end{aligned} \quad (1.2)$$

In SI units $C_{\text{elec}} = 9.0 \times 10^9 \text{ N m}^2 \text{ C}^{-2}$ and $C_{\text{grav}} = 6.7 \times 10^{-11} \text{ N m}^2 \text{ kg}^{-2}$, while in atomic units $C_{\text{elec}} = 1$ and $C_{\text{grav}} = 2.4 \times 10^{-43}$. On an atomic scale, the gravitational interaction is completely negligible compared with the electromagnetic interaction. For the interaction

between a proton and an electron, for example, the ratio between V_{elec} and V_{grav} is 10^{39} . On a large macroscopic scale, such as planets, the situation is reversed. Here the gravitational interaction completely dominates, and the electromagnetic interaction is absent.

On a more fundamental level, it is believed that the four forces are really just different manifestations of a single common interaction, because of the relatively low energy regime we are living in. It has been shown that the weak and electromagnetic forces can be combined into a single unified theory, called *quantum electrodynamics* (QED). Similarly, the strong interaction can be coupled with QED into what is known as the *standard model*. Much effort is being devoted to also include the gravitational interaction into a grand unified theory, and *string theory* is currently believed to hold the greatest promise for such a unification.

Only the electromagnetic interaction is important at the atomic and molecular level, and in the large majority of cases, the simple Coulomb form (in atomic units) is sufficient:

$$V_{\text{Coulomb}}(\mathbf{r}_{ij}) = \frac{q_i q_j}{r_{ij}} \quad (1.3)$$

Within QED, the Coulomb interaction is only the zeroth-order term, and the complete interaction can be written as an expansion in terms of the (inverse) velocity of light, c . For systems where relativistic effects are important (i.e. containing elements from the lower part of the periodic table), or when high accuracy is required, the first-order correction (corresponding to an expansion up to $1/c^2$) for the electron–electron interaction may be included:

$$V_{\text{elec}}(\mathbf{r}_{12}) = \frac{1}{r_{12}} \left[1 - \frac{1}{2} \left(\boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2 + \frac{(\boldsymbol{\alpha}_1 \cdot \mathbf{r}_{12})(\boldsymbol{\alpha}_2 \cdot \mathbf{r}_{12})}{r_{12}^2} \right) \right] \quad (1.4)$$

The first-order correction is known as the *Breit* term, and $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$ represent velocity operators. Physically, the first term in the Breit correction corresponds to magnetic interaction between the two electrons, while the second term describes a “retardation” effect, since the interaction between distant particles is “delayed” relative to interactions between close particles, owing to the finite value of c (in atomic units, $c \sim 137$).

1.4 The Dynamical Equation

The mathematical form for the dynamical equation depends on the mass and velocity of the particles, and can be divided into four regimes.

Newtonian mechanics, exemplified by Newton’s second law ($\mathbf{F} = m\mathbf{a}$), applies for “heavy”, “slow-moving” particles. Relativistic effects become important when the velocity is comparable to the speed of light, causing an increase in the particle mass m relative to the rest mass m_0 . A pragmatic borderline between Newtonian and relativistic (Einstein) mechanics is $\sim 1/3c$, corresponding to a relativistic correction of a few percent.

Light particles display both wave and particle characteristics, and must be described by quantum mechanics, with the borderline being approximately the mass of a proton.

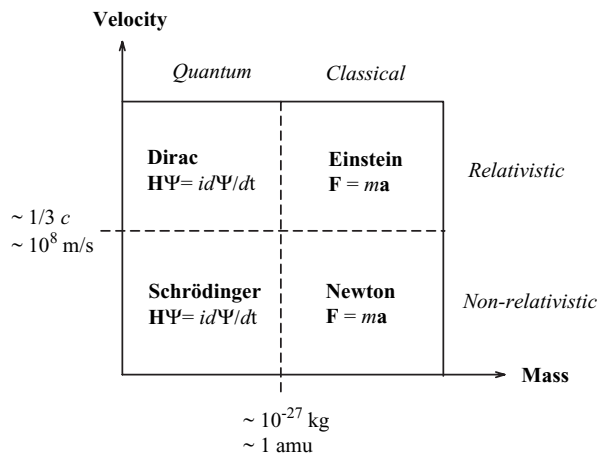


Figure 1.2 Domains of dynamical equations

Electrons are much lighter and can only be described by quantum mechanics, while atoms and molecules, with a few exceptions, behave essentially as classical particles. Hydrogen (protons), being the lightest nucleus, represents a borderline case, which means that quantum corrections in some cases are essential. A prime example is the tunnelling of hydrogen through barriers, allowing reactions involving hydrogen to occur faster than expected from transition state theory.

A major difference between quantum and classical mechanics is that classical mechanics is *deterministic* while quantum mechanics is *probabilistic* (more correctly, quantum mechanics is also deterministic, but the interpretation is probabilistic). Deterministic means that Newton's equation can be integrated over time (forward or backward) and can predict where the particles are at a certain time. This, for example, allows prediction of where and when solar eclipses will occur many thousands of years in advance, with an accuracy of meters and seconds. Quantum mechanics, on the other hand, only allows calculation of the *probability* of a particle being at a certain place at a certain time. The probability function is given as the square of a wave function, $P(\mathbf{r}, t) = \Psi^2(\mathbf{r}, t)$, where the wave function Ψ is obtained by solving either the Schrödinger (non-relativistic) or Dirac (relativistic) equation. Although they appear to be the same in Figure 1.2, they differ considerably in the form of the operator \mathbf{H} .

For classical mechanics at low velocities compared with the speed of light, Newton's second law applies.

$$\mathbf{F} = \frac{d\mathbf{p}}{dt} \quad (1.5)$$

If the particle mass is constant, the derivative of the momentum \mathbf{p} is the mass times the acceleration.

$$\begin{aligned} \mathbf{p} &= m\mathbf{v} \\ \mathbf{F} &= \frac{d\mathbf{p}}{dt} = m \frac{d\mathbf{v}}{dt} = m\mathbf{a} \end{aligned} \quad (1.6)$$

Since the force is the derivative of the potential (eq. (1.1)), and the acceleration is the second derivative of the position \mathbf{r} with respect to time, it may also be written in a differential form.

$$-\frac{\partial \mathbf{V}}{\partial \mathbf{r}} = m \frac{\partial^2 \mathbf{r}}{\partial t^2} \quad (1.7)$$

Solving this equation gives the position of each particle as a function of time, i.e. $\mathbf{r}(t)$.

At velocities comparable with the speed of light, Newton's equation is formally unchanged, but the particle mass becomes a function of the velocity, and the force is therefore not simply a constant (mass) times the acceleration.

$$m = \frac{m_0}{\sqrt{1 - v^2/c^2}} \quad (1.8)$$

For particles with small masses, primarily electrons, quantum mechanics must be employed. At low velocities, the relevant equation is the time-dependent Schrödinger equation.

$$\mathbf{H}\Psi = i \frac{\partial \Psi}{\partial t} \quad (1.9)$$

The Hamiltonian operator is given as a sum of kinetic and potential energy operators.

$$\begin{aligned} \mathbf{H}_{\text{Schrödinger}} &= \mathbf{T} + \mathbf{V} \\ \mathbf{T} &= \frac{\mathbf{p}^2}{2m} = -\frac{1}{2m} \nabla^2 \end{aligned} \quad (1.10)$$

Solving the Schrödinger equation gives the wave function as a function of time, and the probability of observing a particle at a position \mathbf{r} and time t is given as the square of the wave function.

$$P(\mathbf{r}, t) = \Psi^2(\mathbf{r}, t) \quad (1.11)$$

For light particles moving at a significant fraction of the speed of light, the Schrödinger equation is replaced by the Dirac equation.

$$\mathbf{H}\Psi = i \frac{\partial \Psi}{\partial t} \quad (1.12)$$

Although it is formally identical to the Schrödinger equation, the Hamiltonian operator is significantly more complicated.

$$\mathbf{H}_{\text{Dirac}} = (c\boldsymbol{\alpha} \cdot \mathbf{p} + \boldsymbol{\beta}mc^2) + \mathbf{V} \quad (1.13)$$

The $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are 4×4 matrices, and the relativistic wave function consequently has four components. Traditionally, these are labelled the *large* and *small* components, each having an α and β spin function (note the difference between the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ matrices and α and β spin functions). The large component describes the electronic part of the wave function, while the small component describes the positronic (electron antiparticle) part of the wave function, and the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ matrices couple these components. In the limit of $c \rightarrow \infty$, the Dirac equation reduces to the Schrödinger equation, and the two

large components of the wave function reduce to the α and β spin-orbitals in the Schrödinger picture.

1.5 Solving the Dynamical Equation

Both the Newton/Einstein and Schrödinger/Dirac dynamical equations are differential equations involving the derivative of either the position vector or wave function with respect to time. For two-particle systems with simple interaction potentials \mathbf{V} , these can be solved analytically, giving $\mathbf{r}(t)$ or $\Psi(\mathbf{r}, t)$ in terms of mathematical functions. For systems with more than two particles, the differential equation must be solved by numerical techniques involving a sequence of small finite time steps.

Consider a set of particles described by a position vector \mathbf{r}_i at a given time t_i . A small time step Δt later, the positions can be calculated from the velocities, acceleration, hyperaccelerations, etc., corresponding to a Taylor expansion with time as the variable.

$$\mathbf{r}_{i+1} = \mathbf{r}_i + \mathbf{v}_i(\Delta t) + \frac{1}{2}\mathbf{a}_i(\Delta t)^2 + \frac{1}{6}\mathbf{b}_i(\Delta t)^3 + \dots \quad (1.14)$$

The positions a small time step Δt earlier were (replacing Δt with $-\Delta t$)

$$\mathbf{r}_{i-1} = \mathbf{r}_i - \mathbf{v}_i(\Delta t) + \frac{1}{2}\mathbf{a}_i(\Delta t)^2 - \frac{1}{6}\mathbf{b}_i(\Delta t)^3 + \dots \quad (1.15)$$

Addition of these two equations gives a recipe for predicting the positions a time step Δt later from the current and previous positions, and the current acceleration, a method known as the *Verlet* algorithm.

$$\mathbf{r}_{i+1} = (2\mathbf{r}_i - \mathbf{r}_{i-1}) + \mathbf{a}_i(\Delta t)^2 + \dots \quad (1.16)$$

Note that all odd terms in the Verlet algorithm disappear, i.e. the algorithm is correct to third order in the time step. The acceleration can be calculated from the force, or equivalently, the potential.

$$\mathbf{a} = \frac{\mathbf{F}}{m} = -\frac{1}{m} \frac{\partial \mathbf{V}}{\partial \mathbf{r}} \quad (1.17)$$

The time step Δt is an important control parameter for a simulation. The *largest* value of Δt is determined by the *fastest* process occurring in the system, typically being an order of magnitude smaller than the fastest process. For simulating nuclear motions, the fastest process is the motion of hydrogens, being the lightest particles. Hydrogen vibrations occur with a typical frequency of 3000 cm^{-1} , corresponding to $\sim 10^{14}\text{ s}^{-1}$, and therefore necessitating time steps of the order of one femtosecond (10^{-15} s).

1.6 Separation of Variables

As discussed in the previous section, the problem is solving a differential equation with respect to either the position (classical) or wave function (quantum) for the particles in the system. The standard method of solving differential equations is to find a set of coordinates where the differential equation can be separated into less complicated equations. The first step is to introduce a *centre of mass* coordinate system, defined as the mass-weighted sum of the coordinates of all particles, which allows the translation

of the combined system with respect to a fixed coordinate system to be separated from the internal motion. For a two-particle system, the internal motion is described in terms of a reduced mass moving relative to the centre of mass, and this can be further transformed by introducing a coordinate system that reflects the symmetry of the interaction between the two particles. If the interaction only depends on the interparticle distance (e.g. Coulomb or gravitational interaction), the coordinate system of choice is normally a polar (two-dimensional) or spherical polar (three-dimensional) system. In these coordinate systems, the dynamical equation can be transformed into solving one-dimensional differential equations.

For more than two particles, it is still possible to make the transformation to the centre of mass system. However, it is no longer possible to find a set of coordinates that allows a separation of the degrees of freedom for the internal motion, thus preventing an analytical solution. For many-body ($N > 2$) systems, the dynamical equation must therefore be solved by computational (numerical) methods. Nevertheless, it is often possible to achieve an approximate separation of variables based on physical properties, for example particles differing considerably in mass (such as nuclei and electrons). A two-particle system consisting of one nucleus and one electron can be solved exactly by introducing a centre of mass system, thereby transforming the problem into a pseudo-particle with a reduced mass ($\mu = m_1 m_2 / (m_1 + m_2)$) moving relative to the centre of mass. In the limit of the nucleus being infinitely heavier than the electron, the centre of mass system becomes identical to that of the nucleus. In this limit, the reduced mass becomes equal to that of the electron which moves relative to the (stationary) nucleus. For large, but finite, mass ratios, the approximation $\mu \approx m_e$ is unnecessary but may be convenient for interpretative purposes. For many-particle systems, an exact separation is not possible, and the *Born–Oppenheimer* approximation corresponds to assuming that the nuclei are infinitely heavier than the electrons. This allows the electronic problem to be solved for a given set of stationary nuclei. Assuming that the electronic problem can be solved for a large set of nuclear coordinates, the electronic energy forms a parametric hypersurface as a function of the nuclear coordinates, and the motion of the nuclei on this surface can then be solved subsequently.

If an approximate separation is not possible, the many-body problem can often be transformed into a pseudo one-particle system by taking the average interaction into account. For quantum mechanics, this corresponds to the Hartree–Fock approximation, where the average electron–electron repulsion is incorporated. Such pseudo one-particle solutions often form the conceptual understanding of the system, and provide the basis for more refined computational methods.

Molecules are sufficiently heavy that their motions can be described quite accurately by classical mechanics. In condensed phases (solution or solid state), there is a strong interaction between molecules, and a reasonable description can only be attained by having a large number of individual molecules moving under the influence of each other's repulsive and attractive forces. The forces in this case are complex and cannot be written in a simple form such as the Coulomb or gravitational interaction. No analytical solutions can be found in this case, even for a two-particle (molecular) system. Similarly, no approximate solution corresponding to a Hartree–Fock model can be constructed. The only method in this case is direct simulation of the full dynamical equation.

1.6.1 Separating space and time variables

The time-dependent Schrödinger equation involves differentiation with respect to both time and position, the latter contained in the kinetic energy of the Hamiltonian operator.

$$\begin{aligned}\mathbf{H}(\mathbf{r}, t)\Psi(\mathbf{r}, t) &= i \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} \\ \mathbf{H}(\mathbf{r}, t) &= \mathbf{T}(\mathbf{r}) + \mathbf{V}(\mathbf{r}, t)\end{aligned}\tag{1.18}$$

For (bound) systems where the potential energy operator is time-independent ($\mathbf{V}(\mathbf{r}, t) = \mathbf{V}(\mathbf{r})$), the Hamiltonian operator becomes time-independent and yields the total energy when acting on the wave function. The energy is a constant, independent of time, but depends on the space variables.

$$\begin{aligned}\mathbf{H}(\mathbf{r}, t) &= \mathbf{H}(\mathbf{r}) = \mathbf{T}(\mathbf{r}) + \mathbf{V}(\mathbf{r}) \\ \mathbf{H}(\mathbf{r})\Psi(\mathbf{r}, t) &= E(\mathbf{r})\Psi(\mathbf{r}, t)\end{aligned}\tag{1.19}$$

Inserting this in the time-dependent Schrödinger equation shows that the time and space variables of the wave function can be separated.

$$\begin{aligned}\mathbf{H}(\mathbf{r})\Psi(\mathbf{r}, t) &= E(\mathbf{r})\Psi(\mathbf{r}, t) = i \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} \\ \Psi(\mathbf{r}, t) &= \Psi(\mathbf{r})e^{-iEt}\end{aligned}\tag{1.20}$$

The latter follows from solving the first-order differential equation with respect to time, and shows that the time dependence can be written as a simple phase factor multiplied with the spatial wave function. For time-independent problems, this phase factor is normally neglected, and the starting point is taken as the time-independent Schrödinger equation.

$$\mathbf{H}(\mathbf{r})\Psi(\mathbf{r}) = E(\mathbf{r})\Psi(\mathbf{r})\tag{1.21}$$

1.6.2 Separating nuclear and electronic variables

Electrons are very light particles and cannot be described by classical mechanics, while nuclei are sufficiently heavy that they display only small quantum effects. The large mass difference indicates that the nuclear velocities are much smaller than the electron velocities, and the electrons therefore adjust very fast to a change in the nuclear geometry.

For a general N -particle system, the Hamiltonian operator contains kinetic (\mathbf{T}) and potential (\mathbf{V}) energy for all particles.

$$\begin{aligned}\mathbf{H} &= \mathbf{T} + \mathbf{V} \\ \mathbf{T} &= \sum_{i=1}^N \mathbf{T}_i = - \sum_{i=1}^N \frac{1}{2m_i} \nabla_i^2 \\ \nabla_i^2 &= \left(\frac{\partial^2}{\partial x_i^2} + \frac{\partial^2}{\partial y_i^2} + \frac{\partial^2}{\partial z_i^2} \right) \\ \mathbf{V} &= \sum_{i>j}^N V_{ij}\end{aligned}\tag{1.22}$$

The potential energy operator is the Coulomb potential (eq. (1.3)). Denoting nuclear coordinates with \mathbf{R} and subscript n, and electron coordinates with \mathbf{r} and subscript e, this can be expressed as follows.

$$\begin{aligned}
 \mathbf{H}_{\text{tot}} \Psi_{\text{tot}}(\mathbf{R}, \mathbf{r}) &= E_{\text{tot}} \Psi_{\text{tot}}(\mathbf{R}, \mathbf{r}) \\
 \mathbf{H}_{\text{tot}} &= \mathbf{H}_e + \mathbf{T}_n \\
 \mathbf{H}_e &= \mathbf{T}_e + \mathbf{V}_{ne} + \mathbf{V}_{ee} + \mathbf{V}_{nn} \\
 \Psi_{\text{tot}}(\mathbf{R}, \mathbf{r}) &= \Psi_n(\mathbf{R}) \Psi_e(\mathbf{r}) \\
 \mathbf{H}_e \Psi_e(\mathbf{r}) &= E_e(\mathbf{R}) \Psi_e(\mathbf{r}) \\
 (\mathbf{T}_n + E_e(\mathbf{R})) \Psi_n(\mathbf{R}) &= E_{\text{tot}} \Psi_n(\mathbf{R})
 \end{aligned} \tag{1.23}$$

The above approximation corresponds to neglecting the coupling between the nuclear and electronic velocities, i.e. the nuclei are stationary from the electronic point of view. The electronic wave function thus depends *parametrically* on the nuclear coordinates, since it only depends on the *position* of the nuclei, not on their *momentum*. To a good approximation, the electronic wave function thus provides a *potential energy surface* upon which the nuclei move, and this separation is known as the *Born–Oppenheimer* approximation.

The Born–Oppenheimer approximation is usually very good. For the hydrogen molecule (H_2) the error is of the order of 10^{-4} au, and for systems with heavier nuclei the approximation becomes better. As we shall see later, it is possible only in a few cases to solve the electronic part of the Schrödinger equation to an accuracy of 10^{-4} au, i.e. neglect of the nuclear-electron coupling is usually only a minor approximation compared with other errors.

1.6.3 Separating variables in general

Assume that a set of variables can be found where the Hamiltonian operator \mathbf{H} for two particles/variables can be separated into two independent terms, with each only depending on one particle/variable:

$$\mathbf{H} = \mathbf{h}_1 + \mathbf{h}_2 \tag{1.24}$$

Assume furthermore that the Schrödinger equation for one particle/variable can be solved (exactly or approximately):

$$\mathbf{h}_i \phi_i = \varepsilon_i \phi_i \tag{1.25}$$

The solution to the two-particle problem can then be composed of solutions of one-variable Schrödinger equations.

$$\begin{aligned}
 \Psi &= \phi_1 \phi_2 \\
 E &= \varepsilon_1 + \varepsilon_2
 \end{aligned} \tag{1.26}$$

This can be generalized to the case of N particles/variables:

$$\begin{aligned}
 \mathbf{H} &= \sum_i \mathbf{h}_i \\
 \Psi &= \prod_i \phi_i \\
 E &= \sum_i \varepsilon_i
 \end{aligned} \tag{1.27}$$

The properties in eq. (1.27) may be verified by inserting the entities in the Schrödinger equation (1.21).

1.7 Classical Mechanics

1.7.1 The Sun–Earth system

The motion of the Earth around the Sun is an example of a two-body system that can be treated by classical mechanics. The interaction between the two “particles” is the gravitational force.

$$\mathbf{V}(\mathbf{r}_{12}) = -C_{\text{grav}} \frac{m_1 m_2}{r_{12}} \quad (1.28)$$

The dynamical equation is Newton’s second law, which in differential form can be written as in eq. (1.29).

$$-\frac{\partial \mathbf{V}}{\partial \mathbf{r}} = m \frac{\partial^2 \mathbf{r}}{\partial t^2} \quad (1.29)$$

The first step is to introduce a centre of mass system, and the internal motion becomes motion of a “particle” with a reduced mass given by eq. (1.30).

$$\mu = \frac{M_{\text{Sun}} m_{\text{Earth}}}{M_{\text{Sun}} + m_{\text{Earth}}} = \frac{m_{\text{Earth}}}{(1 + m_{\text{Earth}}/M_{\text{Sun}})} \cong m_{\text{Earth}} \quad (1.30)$$

Since the mass of the Sun is 3×10^5 times larger than that of the Earth, the reduced mass is essentially identical to the Earth’s mass ($\mu = 0.999997 m_{\text{Earth}}$). To a very good approximation, the system can therefore be described as the Earth moving around the Sun, which remains stationary.

The motion of the Earth around the Sun occurs in a plane, and a suitable coordinate system is a polar coordinate system (two-dimensional) consisting of r and θ .




Figure 1.3 A polar coordinate system

The interaction depends only on the distance r , and the differential equation (Newton’s equation) can be solved analytically. The bound solutions are elliptical orbits with the Sun (more precisely, the centre of mass) at one of the foci, but for most of the planets, the actual orbits are close to circular. Unbound solutions corresponding to hyperbolas also exist, and could for example describe the path of a (non-returning) comet.

Each bound orbit can be classified in terms of the dimensions (largest and smallest distance to the Sun), with an associated total energy. In classical mechanics, there are no constraints on the energy, and all sizes of orbits are allowed. If the zero point for




Figure 1.4 Bound and unbound solutions to the classical two-body problem

the energy is taken as the two particles at rest infinitely far apart, positive values of the total energy correspond to unbound solutions (hyperbolas, with the kinetic energy being larger than the potential energy) while negative values correspond to bound orbits (ellipsoids, with the kinetic energy being less than the potential energy). Bound solutions are also called *stationary* orbits, as the particle position returns to the same value with well-defined time intervals.

1.7.2 The solar system

Once we introduce additional planets in the Sun–Earth system, an analytical solution for the motions of all the planets can no longer be obtained. Since the mass of the Sun is so much larger than the remaining planets (the Sun is 1000 times heavier than Jupiter, the largest planet), the interactions between the planets can to a good approximation be neglected. For the Earth, for example, the second most important force is from the Moon, with a contribution that is 180 times smaller than that from the Sun. The next largest contribution is from Jupiter, being approximately 30 000 times smaller (on average) than the gravitational force from the Sun. In this *central field model*, the orbit of each planet is determined as if it were the only planet in the solar system, and the resulting computational task is a two-particle problem, i.e. elliptical orbits with the Sun at one of the foci. The complete solar system is the unification of nine such orbits, and the total energy is the sum of all nine individual energies.

A formal refinement can be done by taking the *average* interaction between the planets into account, i.e. a Hartree–Fock type approximation. In this model, the orbit of one planet (e.g. the Earth) is determined by taking the average interaction with all the other planets into account. The average effect corresponds to spreading the mass of the other planets evenly along their orbits.

The Hartree–Fock model represents only a very minute improvement over the independent orbit model for the solar system, since the planetary orbits do not cross. The effect of a planet inside the Earth’s orbit corresponds to adding its mass to the Sun, while the effect of the spread-out mass of a planet outside the Earth’s orbit is zero. The Hartree–Fock model for the Earth thus consists of increasing the Sun’s effective mass with that of Mercury and Venus, i.e. a change of only 0.0003%. For the solar system there is thus very little difference between totally neglecting the planetary interactions and taking the average effect into account.

The real system, of course, includes all interactions, where each pair interaction depends on the actual distance between the planets. The resulting planetary motions




Figure 1.5 A Hartree-Fock model for the solar system




Figure 1.6 Modelling the solar system with actual interactions

cannot be solved analytically, but can be simulated numerically. From a given starting condition, the system is allowed to evolve for many small time steps, and all interactions are considered constant within each time step. By sufficiently small time steps, this yields a very accurate model of the real many-particle dynamics, and will display small wiggles of the planetary motion around the elliptical orbits calculated by either of the two independent-particle models.

Since the perturbations due to the other planets are significantly smaller than the interaction with the Sun, the “wiggles” are small compared with the overall orbital motion, and a description of the solar system as planets orbiting the Sun in elliptical orbits is a very good approximation to the true dynamics of the system.

1.8 Quantum Mechanics

1.8.1 A hydrogen-like atom

A quantum analogue of the Sun–Earth system is a nucleus and one electron, i.e. a hydrogen-like atom. The force holding the nucleus and electron together is the Coulomb interaction.

$$\mathbf{V}(\mathbf{r}_{12}) = \frac{q_1 q_2}{r_{12}} \quad (1.31)$$

The interaction again only depends on the distance, but owing to the small mass of the electron, Newton's equation must be replaced with the Schrödinger equation. For bound states, the time-dependence can be separated out, as shown in Section 1.6.1, giving the time-independent Schrödinger equation.

$$\mathbf{H}\Psi = E\Psi \quad (1.32)$$

The Hamiltonian operator for a hydrogen-like atom (nuclear charge of Z) can in Cartesian coordinates and atomic units be written as eq. (1.33), with M being the nuclear and m the electron mass ($m = 1$ in atomic units).

$$\mathbf{H} = -\frac{1}{2M}\nabla_1^2 - \frac{1}{2m}\nabla_2^2 - \frac{Z}{\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2}} \quad (1.33)$$

The Laplace operator is given by eq. (1.34).

$$\nabla_i^2 = \frac{\partial^2}{\partial x_i^2} + \frac{\partial^2}{\partial y_i^2} + \frac{\partial^2}{\partial z_i^2} \quad (1.34)$$

The two kinetic energy operators are already separated, since each only depends on three coordinates. The potential energy operator, however, involves all six coordinates. The centre of mass system is introduced by the following six coordinates.

$$\begin{aligned} X &= \frac{(Mx_1 + mx_2)}{(M + m)}; & x &= x_1 - x_2 \\ Y &= \frac{(My_1 + my_2)}{(M + m)}; & y &= y_1 - y_2 \\ Z &= \frac{(Mz_1 + mz_2)}{(M + m)}; & z &= z_1 - z_2 \end{aligned} \quad (1.35)$$

Here the X, Y, Z coordinates define the centre of mass system, and the x, y, z coordinates specify the relative position of the two particles. In these coordinates the Hamiltonian operator can be rewritten as eq. (1.36).

$$\mathbf{H} = -\frac{1}{2}\nabla_{XYZ}^2 - \frac{1}{2\mu}\nabla_{xyz}^2 - \frac{Z}{\sqrt{x^2 + y^2 + z^2}} \quad (1.36)$$

The first term only involves the X, Y and Z coordinates, and the ∇_{XYZ}^2 operator is obviously separable in terms of X, Y and Z . Solution of the XYZ part gives translation of the whole system in three dimensions relative to the laboratory-fixed coordinate system. The xyz coordinates describe the relative motion of the two particles in terms of a pseudo-particle with a reduced mass μ relative to the centre of mass.

$$\mu = \frac{M_{\text{nuc}}m_{\text{elec}}}{M_{\text{nuc}} + m_{\text{elec}}} = \frac{m_{\text{elec}}}{\left(1 + \frac{m_{\text{elec}}}{M_{\text{nuc}}}\right)} \equiv m_{\text{elec}} \quad (1.37)$$

For the hydrogen atom, the nucleus is approximately 1800 times heavier than the electron, giving a reduced mass of $0.9995m_{\text{elec}}$. Similar to the Sun–Earth system, the hydrogen atom can therefore to a good approximation be considered as an electron moving around a stationary nucleus, and for heavier elements the approximation becomes better (with a uranium nucleus, for example, the nucleus/electron mass ratio is $\sim 430\,000$). Setting the reduced mass equal to the electron mass corresponds to making the assumption that the nucleus is infinitely heavy and therefore stationary.

The potential energy again only depends on the distance between the two particles, but in contrast to the Sun–Earth system, the motion occurs in three dimensions, and it is therefore advantageous to transform the Schrödinger equation into a spherical polar set of coordinates.




Figure 1.7 A spherical polar coordinate system

The potential energy becomes very simple, but the kinetic energy operator becomes complicated.

$$\mathbf{H} = -\frac{1}{2\mu} \nabla_{r\theta\phi}^2 - \frac{Z}{r} \quad (1.38)$$







$$\nabla_{r\theta\phi}^2 = \frac{1}{r^2} \frac{\partial}{\partial r} r^2 \frac{\partial}{\partial r} + \frac{1}{r^2 \sin \theta} \frac{\partial}{\partial \theta} \sin \theta \frac{\partial}{\partial \theta} + \frac{1}{r^2 \sin^2 \theta} \frac{\partial^2}{\partial \phi^2}$$

The kinetic energy operator, however, is almost separable in spherical polar coordinates, and the actual method of solving the differential equation can be found in a number of textbooks. The bound solutions (negative total energy) are called orbitals and can be classified in terms of three *quantum numbers*, n , l and m , corresponding to the three spatial variables r , θ and ϕ . The quantum numbers arise from the boundary conditions on the wave function, i.e. it must be periodic in the θ and ϕ variables, and must decay to zero as $r \rightarrow \infty$. Since the Schrödinger equation is not completely separable in spherical polar coordinates, there exist the restrictions $n > l \geq |m|$. The n quantum number describes the *size* of the orbital, the l quantum number describes the *shape* of the orbital, while the m quantum number describes the *orientation* of the orbital relative to a fixed coordinate system. The l quantum number translates into names for the orbitals:

- $l = 0$: s-orbital
- $l = 1$: p-orbital
- $l = 2$: d-orbital, etc.

The orbitals can be written as a product of a *radial* function, describing the behaviour in terms of the distance r between the nucleus and electron, and *spherical harmonic* functions Y_{lm} representing the angular part in terms of the angles θ and ϕ . The orbitals can be visualized by plotting three-dimensional objects corresponding to the wave function having a specific value, e.g. $\Psi^2 = 0.10$.

Table 1.2 Hydrogenic orbitals obtained from solving the Schrödinger equation

n	l	m	$\Psi_{n,l,m}(r,\theta,\varphi)$	Shape and size
1	0	0	$Y_{0,0}(\theta,\varphi)e^{-Zr}$	
2	0	0	$Y_{0,0}(\theta,\varphi)(2 - Zr)e^{-Zr/2}$	
	1	$\pm 1, 0$	$Y_{1,m}(\theta,\varphi)Zre^{-Zr/2}$	
3	0	0	$Y_{0,0}(\theta,\varphi)(27 - 18Zr + 2Z^2r^2)e^{-Zr/3}$	
	1	$\pm 1, 0$	$Y_{1,m}(\theta,\varphi)Zr(6 - Zr)e^{-Zr/3}$	
	2	$\pm 2, \pm 1, 0$	$Y_{2,m}(\theta,\varphi)Z^2r^2e^{-Zr/3}$	

The orbitals for different quantum numbers are orthogonal and can be chosen to be normalized.

$$\langle \Psi_{n,l,m} | \Psi_{n',l',m'} \rangle = \delta_{n,n'} \delta_{l,l'} \delta_{m,m'} \quad (1.39)$$

The orthogonality of the orbitals in the angular part (l and m quantum numbers) follows from the shape of the spherical harmonic functions, as these have l nodal planes (points where the wave function is zero). The orthogonality in the radial part (n quantum number) is due to the presence of $(n-l-1)$ radial nodes in the wave function.

In contrast to classical mechanics, where all energies are allowed, wave functions and associated energies are quantized, i.e. only certain values are allowed. The energy only depends on n for a given nuclear charge Z , and is given by eq. (1.40).

$$E = -\frac{Z^2}{2n^2} \quad (1.40)$$

Unbound solutions have a positive total energy and correspond to scattering of an electron by the nucleus.

1.8.2 The helium atom

Like the solar system, it is not possible to find a set of coordinates where the Schrödinger equation can be solved analytically for more than two particles (i.e. for many-electron atoms). Owing to the dominance of the Sun's gravitational field, a central field approximation provides a good description of the actual solar system, but this is not the case for an atomic system. The main differences between the solar system and an atom such as helium are:

- (1) The interaction between the electrons is only a factor of two smaller than between the nucleus and electrons, compared with a ratio of at least 1000 for the solar system.
- (2) The electron–electron interaction is repulsive, compared with the attraction between planets.

- (3) The motion of the electrons must be described by quantum mechanics owing to the small electron mass, and the particle position is determined by an orbital, the square of which gives the probability of finding the electron at a given position.
- (4) Electrons are indistinguishable particles having a spin of $1/2$. This fermion character requires the total wave function to be antisymmetric, i.e. it must change sign when interchanging two electrons. The antisymmetry results in the so-called exchange energy, which is a non-classical correction to the Coulomb interaction.

The simplest atomic model would be to neglect the electron–electron interaction, and only take the nucleus–electron attraction into account. In this model each orbital for the helium atom is determined by solving a hydrogen-like system with a nucleus and one electron, yielding hydrogen-like orbitals, 1s, 2s, 2p, 3s, 3p, 3d, etc., with $Z = 2$. The total wave function is obtained from the resulting orbitals subject to the aufbau and Pauli principles. These principles say that the lowest energy orbitals should be filled first and only two electrons (with different spin) can occupy each orbital, i.e. the electron configuration becomes $1s^2$. The antisymmetry condition is conveniently fulfilled by writing the total wave function as a Slater determinant, since interchanging any two rows or columns changes the sign of the determinant. For a helium atom, this would give the following (unnormalized) wave function, with the orbitals given in Table 1.2 with $Z = 2$.

$$\Phi = \begin{vmatrix} \phi_{1s\alpha}(1) & \phi_{1s\beta}(1) \\ \phi_{1s\alpha}(2) & \phi_{1s\beta}(2) \end{vmatrix} = \phi_{1s\alpha}(1)\phi_{1s\beta}(2) - \phi_{1s\beta}(1)\phi_{1s\alpha}(2) \quad (1.41)$$

The total energy calculated by this wave function is simply twice the orbital energy, -4.000 au, which is in error by 38% compared with the experimental value of -2.904 au. Alternatively, we can use the wave function given by eq. (1.41), but include the electron–electron interaction in the energy calculation, giving a value of -2.750 au.

A better approximation can be obtained by taking the *average* repulsion between the electrons into account when determining the orbitals, a procedure known as the Hartree–Fock approximation. If the orbital for one of the electrons were somehow known, the orbital for the second electron could be calculated in the electric field of the nucleus and the first electron, described by its orbital. This argument could just as well be used for the *second* electron with respect to the *first* electron. The goal is therefore to calculate a set of self-consistent orbitals, and this can be done by iterative methods.

For the solar system, the non-crossing of the planetary orbitals makes the Hartree–Fock approximation only a very minor improvement over a central field model. For a many-electron atom, however, the situation is different since the position of the electrons is described by three-dimensional probability functions (square of the orbitals), i.e. the electron “orbits” “cross”. The average nucleus–electron distance for an electron in a 2s-orbital is larger than for one in a 1s-orbital, but there is a finite probability that a 2s-electron is closer to the nucleus than a 1s-electron. If the 1s-electrons in lithium were completely inside the 2s-orbital, the latter would experience an effective nuclear charge of 1.00, but owing to the 2s-electron penetrating the 1s-orbital, the effective nuclear charge for an electron in a 2s-orbital is 1.26. The 2s-electron in return screens the nuclear charge felt by the 1s-electrons, making the effective nuclear charge felt by the 1s-electrons less than 3.00. The mutual screening of the two 1s-electrons in helium produces an effective nuclear charge of 1.69, yielding a total

energy of -2.848 au, which is a significant improvement relative to the model with orbitals employing a fixed nuclear charge of 2.00.

Although the effective nuclear charge of 1.69 represents the lowest possible energy with the functional form of the orbitals in Table 1.2, it is possible to further refine the model by relaxing the functional form of the orbitals from a strict exponential. Although the exponential form is the exact solution for a hydrogen-like system, this is not the case for a many-electron atom. Allowing the orbitals to adopt best possible form, and simultaneously optimizing the exponents (“effective nuclear charge”), gives an energy of -2.862 au. This represents the best possible independent-particle model for the helium atom, and any further refinement must include the instantaneous correlation between the electrons. By using the electron correlation methods described in Chapter 4, it is possible to reproduce the experimental energy of -2.904 au.

Table 1.3 Helium atomic energies in various approximations

Wave function	Z_{eff}	Energy (au)
He ⁺ exponential orbital, no electron–electron repulsion	2.00	-4.000
He ⁺ exponential orbital, including electron–electron repulsion	2.00	-2.750
Optimum single exponential orbital	1.69	-2.848
Best orbital, Hartree–Fock limit		-2.862
Experimental		-2.904

The equal mass of all the electrons and the strong interaction between them makes the Hartree–Fock model less accurate than desirable, but it is still a big improvement over an independent orbital model. The Hartree–Fock model typically accounts for $\sim 99\%$ of the total energy, but the remaining *correlation energy* is usually very important for chemical purposes. The correlation between the electrons describes the “wiggles” relative to the Hartree–Fock orbitals due to the instantaneous interaction between the electrons, rather than just the average repulsion. The goal of correlated methods for solving the Schrödinger equation is to calculate the remaining correction due to the electron–electron interaction.

1.9 Chemistry

The Born–Oppenheimer separation of the electronic and nuclear motions is a cornerstone in computational chemistry. Once the electronic Schrödinger equation has been solved for a large number of nuclear geometries (and possibly also for several electronic states), the *potential energy surface* (PES) is known. The motion of the nuclei on the PES can then be solved either classically (Newton) or by quantum (Schrödinger) methods. If there are N nuclei, the dimensionality of the PES is $3N$, i.e. there are $3N$ nuclear coordinates that define the geometry. Of these coordinates, three describe the overall translation of the molecule, and three describe the overall rotation of the molecule with respect to three axes. For a linear molecule, only two coordinates are necessary for describing the rotation. This leaves $3N - 6(5)$ coordinates to describe the internal movement of the nuclei, which for small displacements may be chosen as “vibrational normal coordinates”.

It should be stressed that nuclei are heavy enough that quantum effects are almost negligible, i.e. they behave to a good approximation as classical particles. Indeed, if nuclei showed significant quantum aspects, the concept of molecular structure (i.e. different configurations and conformations) would not have any meaning, since the nuclei would simply tunnel through barriers and end up in the global minimum. Dimethyl ether, for example, would spontaneously transform into ethanol. Furthermore, it would not be possible to speak of a molecular geometry, since the Heisenberg uncertainty principle would not permit a measure of nuclear positions with an accuracy smaller than the molecular dimension.

Methods aimed at solving the electronic Schrödinger equation are broadly referred to as “electronic structure calculations”. An accurate determination of the electronic wave function is very demanding. Constructing a complete PES for molecules containing more than three or four atoms is virtually impossible. Consider, for example, mapping the PES by calculating the electronic energy for every 0.1 Å over say a 1 Å range (a very coarse mapping). With three atoms, there are three internal coordinates, giving 10^3 points to be calculated. Four atoms already produce six internal coordinates, giving 10^6 points, which is possible to calculate, but only with a very determined effort. Larger systems are out of reach. Constructing global PES’s for all but the smallest molecules is thus impossible. By restricting the calculations to the “chemically interesting” part of the PES, however, it is possible to obtain useful information. The interesting parts of a PES are usually nuclear arrangements that have low energies. For example, nuclear movements near a minimum on the PES, which corresponds to a stable molecule, are molecular vibrations. Chemical reactions correspond to larger movements, and may in the simplest approximation be described by locating the lowest energy path leading from one minimum on the PES to another.

These considerations lead to the following definition:

Chemistry is knowing the energy as a function of the nuclear coordinates.

The large majority of what are commonly referred to as molecular properties may similarly be defined as:

Properties are knowing how the energy changes upon adding a perturbation.

In the following we will look at some aspects of solving the electronic Schrödinger equation or otherwise construct a PES, how to deal with the movement of nuclei on the PES, and various technical points of commonly used methods. A word of caution here: although it is the nuclei that move, and the electrons follow “instantly” (according to the Born–Oppenheimer approximation), it is common also to speak of “atoms” moving. An isolated atom consists of a nucleus and some electrons, but in a molecule the concept of an atom is not well defined. Analogously to the isolated atom, an atom in a molecule should consist of a nucleus and some electrons. But how does one partition the total electron distribution in a molecule such that a given portion belongs to a given nucleus? Nevertheless, the words nucleus and atom are often used interchangeably.

Much of the following will concentrate on describing individual molecules. Experiments are rarely done on a single molecule; rather they are performed on macroscopic samples with perhaps 10^{20} molecules. The link between the properties of a single molecule, or a small collection of molecules, and the macroscopic observable is statistical

mechanics. Briefly, macroscopic properties, such as temperature, heat capacity, entropy, etc., are the net effect of a very large number of molecules having a certain distribution of energies. If all the possible energy states can be determined for an individual molecule or a small collection of molecules, statistical mechanics can be used for calculating macroscopic properties.

References

1. T. Clark, *A Handbook of Computational Chemistry*, Wiley, 1985; D. M. Hirst, *A Computational Approach to Chemistry*, Blackwell, 1990; A. Hinchcliffe, *Modelling Molecular Structure*, Wiley, 1996; D. Young, *Computational Chemistry*, Wiley-Interscience, 2001; A. R. Leach, *Molecular Modelling. Principles and Applications*, Longman, 2001; C. J. Cramer, *Essentials of Computational Chemistry*, Wiley 2002; T. Schlick, *Molecular Modeling and Simulation*, Springer, 2002.

2 Force Field Methods

2.1 Introduction

As mentioned in Chapter 1, one of the major problems is calculating the electronic energy for a given nuclear configuration to give a potential energy surface. In *force field* (FF) methods, this step is bypassed by writing the electronic energy as a parametric function of the nuclear coordinates, and fitting the parameters to experimental or higher level computational data. The “building blocks” in force field methods are atoms, i.e. electrons are not considered as individual particles. This means that bonding information must be provided *explicitly*, rather than being the *result* of solving the electronic Schrödinger equation.

In addition to bypassing the solution of the electronic Schrödinger equation, the quantum aspects of the nuclear motion are also neglected. This means that the dynamics of the atoms is treated by classical mechanics, i.e. Newton’s second law. For time-independent phenomena, the problem reduces to calculating the energy at a given geometry. Often the interest is in finding geometries of stable molecules and/or different conformations, and possibly also interconversion between conformations. The problem is then reduced to finding energy minima (and possibly also some first-order saddle points) on the potential energy surface.

Molecules are described by a “ball and spring” model in force field methods, with atoms having different sizes and “softness” and bonds having different lengths and “stiffness”.¹ Force field methods are also referred to as *molecular mechanics* (MM) methods. Many different force fields exist, and in this Chapter we will use Allinger’s MM2 and MM3 (Molecular Mechanics versions 2 and 3) to illustrate specific details.²

The foundation of force field methods is the observation that molecules tend to be composed of units that are structurally similar in different molecules. All C—H bond lengths, for example, are roughly constant in all molecules, being between 1.06 and 1.10 Å. The C—H stretch vibrations are also similar, between 2900 and 3300 cm⁻¹, implying that the C—H force constants are also comparable. If the C—H bonds are further divided into groups, for example those attached to single-, double- or triple-bonded carbon, the variation within each of these groups becomes even smaller. The

Table 2.1 MM2(91) atom types

Type	Symbol	Description	Type	Symbol	Description
1	C	sp ³ -carbon	28	H	enol or amide
2	C	sp ² -carbon, alkene	48	H	ammonium
3	C	sp ² -carbon, carbonyl, imine	36	D	deuterium
4	C	sp-carbon	20	lp	lone pair
22	C	cyclopropane	15	S	sulfide (R ₂ S)
29	C·	radical	16	S ⁺	sulfonium (R ₃ S ⁺)
30	C ⁺	carbocation	17	S	sulfoxide (R ₂ SO)
38	C	sp ² -carbon, cyclopropene	18	S	sulfone (R ₂ SO ₂)
50	C	sp ² -carbon, aromatic	42	S	sp ² -sulfur, thiophene
56	C	sp ³ -carbon, cyclobutane	11	F	fluoride
57	C	sp ² -carbon, cyclobutene	12	Cl	chloride
58	C	carbonyl, cyclobutanone	13	Br	bromide
67	C	carbonyl, cyclopropanone	14	I	iodide
68	C	carbonyl, ketene	26	B	boron, trigonal
71	C	ketonium carbon	27	B	boron, tetrahedral
8	N	sp ³ -nitrogen	19	Si	silane
9	N	sp ² -nitrogen, amide	25	P	phosphine
10	N	sp-nitrogen	60	P	phosphor, pentavalent
37	N	azo or pyridine (–N=)	51	He	helium
39	N ⁺	sp ³ -nitrogen, ammonium	52	Ne	neon
40	N	sp ² -nitrogen, pyrrole	53	Ar	argon
43	N	azoxy (–N=N–O)	54	Kr	krypton
45	N	azide, central atom	55	Xe	xenon
46	N	nitro (–NO ₂)	31	Ge	germanium
72	N	imine, oxime (=N–)	32	Sn	tin
6	O	sp ³ -oxygen	33	Pb	lead
7	O	sp ² -oxygen, carbonyl	34	Se	selenium
41	O	sp ² -oxygen, furan	35	Te	tellurium
47	O [–]	carboxylate	59	Mg	magnesium
49	O	epoxy	61	Fe	iron (II)
69	O	amine oxide	62	Fe	iron (III)
70	O	ketonium oxygen	63	Ni	nickel (II)
5	H	hydrogen, except on N or O	64	Ni	nickel (III)
21	H	alcohol (OH)	65	Co	cobalt (II)
23	H	amine (NH)	66	Co	cobalt (III)
24	H	carboxyl (COOH)			

Note that special atom types are defined for carbon atoms involved in small rings, such as cyclopropane and cyclobutane. The reason for this will be discussed in Section 2.2.2.

same grouping holds for other feature as well, e.g. all C=O bonds are approximately 1.22 Å long and have vibrational frequencies of approximately 1700 cm^{–1}, all double-bonded carbons are essentially planar, etc. The transferability also holds for energetic features. A plot of the heat of formation for linear alkanes, i.e. CH₃(CH₂)_nCH₃, against the chain length *n* produces a straight line, showing that each CH₂ group contributes essentially the same amount of energy. (For a general discussion of estimating heat of formation from group additivities, see Benzon.³)

The picture of molecules being composed of structural units (“functional groups”) that behave similarly in different molecules forms the very basis of organic chemistry. Molecular structure drawings, where alphabetic letters represent atoms and lines represent bonds, are used universally. Organic chemists often build ball and stick, or CPK space-filling, models of their molecules to examine their shapes. Force field methods are in a sense a generalization of these models, with the added feature that the atoms and bonds are not fixed at one size and length. Furthermore, force field calculations enable predictions of relative energies and barriers for interconversion of different conformations.

The idea of molecules being composed of atoms, which are structurally similar in different molecules, is implemented in force field models as atom *types*. The atom type depends on the atomic number and the type of chemical bonding it is involved in. The type may be denoted with either a number or a letter code. In MM2, for example, there are 71 different atom types (type 44 is missing). Type 1 is an sp^3 -hybridized carbon, and an sp^2 -hybridized carbon may be type 2, 3 or 50, depending on the neighbour atom(s). Type 2 is used if the bonding is to another sp^2 -carbon (simple double bond), type 3 is used if the carbon is bonded to an oxygen (carbonyl group) and type 50 is used if the carbon is part of an aromatic ring with delocalized bonds. Table 2.1 gives a complete list of the MM2(91) atom types, where (91) indicates the year when the parameter set was released. The atom type numbers roughly reflect the order in which the corresponding functional groups were parameterized.

2.2 The Force Field Energy

The force field energy is written as a sum of terms, each describing the energy required for distorting a molecule in a specific fashion.

$$E_{\text{FF}} = E_{\text{str}} + E_{\text{bend}} + E_{\text{tors}} + E_{\text{vdw}} + E_{\text{el}} + E_{\text{cross}} \quad (2.1)$$

E_{str} is the energy function for *stretching* a bond between two atoms, E_{bend} represents the energy required for *bending* an angle, E_{tors} is the *torsional* energy for rotation around a bond, E_{vdw} and E_{el} describe the *non-bonded* atom–atom interactions, and finally E_{cross} describes coupling between the first three terms.




Figure 2.1 Illustration of the fundamental force field energy terms

Given such an energy function of the nuclear coordinates, geometries and relative energies can be calculated by optimization. Stable molecules correspond to minima on the potential energy surface, and they can be located by minimizing E_{FF} as a function of the nuclear coordinates. Conformational transitions can be described by locating

transition structure on the E_{FF} surface. Exactly how such a multi-dimensional function optimization may be carried out is described in Chapter 12.

2.2.1 The stretch energy

E_{str} is the energy function for stretching a bond between two atom types A and B. In its simplest form, it is written as a Taylor expansion around a “natural”, or “equilibrium”, bond length, R_0 . Terminating the expansion at second order gives eq. (2.2).

$$E_{\text{str}}(R^{\text{AB}} - R_0^{\text{AB}}) = E(0) + \frac{dE}{dR}(R^{\text{AB}} - R_0^{\text{AB}}) + \frac{1}{2} \frac{d^2E}{dR^2}(R^{\text{AB}} - R_0^{\text{AB}})^2 \quad (2.2)$$

The derivatives are evaluated at $R = R_0$ and the $E(0)$ term is normally set to zero, since this is just the zero point for the energy scale. The second term is zero as the expansion is around the equilibrium value. In its simplest form the stretch energy can thus be written as eq. (2.3).

$$E_{\text{str}}(R^{\text{AB}} - R_0^{\text{AB}}) = k^{\text{AB}}(R^{\text{AB}} - R_0^{\text{AB}})^2 = k^{\text{AB}}(\Delta R^{\text{AB}})^2 \quad (2.3)$$

Here k^{AB} is the “force constant” for the A—B bond. This is the form of a harmonic oscillator, with the potential being quadratic in the displacement from the minimum.

The harmonic form is the simplest possible, and sufficient for determining most equilibrium geometries. There are certain strained and crowded systems where the results from a harmonic approximation are significantly different from experimental values, and if the force field should be able to reproduce features such as vibrational frequencies, the functional form for E_{str} must be improved. The straightforward approach is to include more terms in the Taylor expansion.

$$E_{\text{str}}(\Delta R^{\text{AB}}) = k_2^{\text{AB}}(\Delta R^{\text{AB}})^2 + k_3^{\text{AB}}(\Delta R^{\text{AB}})^3 + k_4^{\text{AB}}(\Delta R^{\text{AB}})^4 + \dots \quad (2.4)$$

This of course has a price: more parameters have to be assigned.

Polynomial expansions of the stretch energy do not have the correct limiting behaviour. The cubic anharmonicity constant k_3 is normally negative, and if the Taylor expansion is terminated at third order, the energy will go toward $-\infty$ for long bond lengths. Minimization of the energy with such an expression can cause the molecule to fly apart if a poor starting geometry is chosen. The quartic constant k_4 is normally positive and the energy will go toward $+\infty$ for long bond lengths if the Taylor series is terminated at fourth order. The correct limiting behaviour for a bond stretched to infinity is that the energy should converge towards the dissociation energy. A simple function that satisfies this criterion is the Morse potential.⁴

$$E_{\text{Morse}}(\Delta R) = D(1 - e^{-\alpha\Delta R})^2 \quad (2.5)$$

$$\alpha = \sqrt{\frac{k}{2D}}$$

Here D is the dissociation energy and α is related to the force constant. The Morse function reproduces the actual behaviour quite accurately over a wide range of distances, as seen in Figure 2.2. There are, however, some difficulties with the Morse potential in actual applications. For long bond lengths the restoring force is quite small. Distorted structures, which may either be a poor starting geometry or one that develops during a simulation, will therefore display a slow convergence towards the

equilibrium bond length. For minimization purposes and simulations at ambient temperatures (e.g. 300 K) it is sufficient that the potential is reasonably accurate up to ~ 40 kJ/mol above the minimum (the average kinetic energy is 3.7 kJ/mol at 300 K). In this energy range there is little difference between a Morse potential and a Taylor expansion, and most force fields therefore employ a simple polynomial for the stretch energy. The number of parameters is often reduced by taking the cubic, quartic, etc., constants as a predetermined fraction of the harmonic force constant. A popular method is to require that the n th-order derivative at R_0 matches the corresponding derivative of the Morse potential. For a fourth-order expansion this leads to the following expression.

$$E_{\text{str}}(\Delta R^{\text{AB}}) = k_2^{\text{AB}}(\Delta R^{\text{AB}})^2 \left[1 - \alpha(\Delta R^{\text{AB}}) + \frac{7}{12}\alpha^2(\Delta R^{\text{AB}})^2 \right] \quad (2.6)$$

The α constant is the same as that appearing in the Morse function, but may be taken as a fitting parameter. An alternative method for introducing anharmonicity is to use the harmonic form in eq. (2.3) but allow the force constant to depend on the bond distance.⁵

Figure 2.2 compares the performance of various functional forms for the stretch energy in CH_4 . The “exact” form is taken from electronic structure calculations ([8,8]-CASSCF/aug-cc-pVTZ). The simple harmonic approximation (P2) is seen to be accurate to about ± 0.1 Å from the equilibrium geometry and the quartic approximation (P4) up to ± 0.3 Å. The Morse potential reproduces the real curve quite accurately up to an elongation of 0.8 Å, and becomes exact again in the dissociation limit.

For the large majority of systems, including simulations, the only important chemical region is within ~ 40 kJ/mol of the bottom of the curve. In this region, a fourth-order polynomial is essentially indistinguishable from either a Morse or the exact curve, as shown in Figure 2.3, and even a simple harmonic approximation does a quite good job.




Figure 2.2 The stretch energy for CH_4




Figure 2.3 The stretch energy for CH_4

Until now, we have used two different words for the R_0 parameter, the “natural” or the “equilibrium” bond length. The latter is slightly misleading. The R_0 parameter is *not* the equilibrium bond length for *any* molecule! Instead it is the parameter which, when used to calculate the minimum energy structure of a molecule, will produce a geometry having the experimental equilibrium bond length. If there were only one stretch energy in the whole force field energy expression (i.e. a diatomic molecule), R_0 would be the equilibrium bond length. However, in a polyatomic molecule the other terms in the force field energy will usually produce a minimum energy structure with bond lengths slightly longer than R_0 . R_0 is the hypothetical bond length if no other terms are included, and the word “natural” bond length is a better description of this parameter than “equilibrium” bond length. Essentially all molecules have bond lengths that deviate very little from their “natural” values, typically by less than 0.03 \AA . For this reason a simple harmonic is usually sufficient for reproducing experimental geometries.

For each bond type, i.e. a bond between two atom types A and B, there are at least two parameters to be determined, k^{AB} and R_0^{AB} . The higher order expansions, and the Morse potential, have one additional parameter (α or D) that needs to be determined.

2.2.2 The bending energy

E_{bend} is the energy required for bending an angle formed by three atoms A—B—C, where there is a bond between A and B, and between B and C. Similarly to E_{str} , E_{bend} is usually expanded as a Taylor series around a “natural” bond angle and terminated at second order, giving the harmonic approximation.

$$E_{\text{bend}}(\theta^{\text{ABC}} - \theta_0^{\text{ABC}}) = k^{\text{ABC}}(\theta^{\text{ABC}} - \theta_0^{\text{ABC}})^2 \quad (2.7)$$

While the simple harmonic expansion is adequate for most applications, there may be cases where higher accuracy is required. The next improvement is to include a third-order term, analogous to E_{str} . This can give a very good description over a large range of angles, as illustrated in Figure 2.4 for CH_4 . The “exact” form is again taken from electronic structure calculations (MP2/aug-cc-pVTZ). The simple harmonic approximation (P2) is seen to be accurate to about $\pm 30^\circ$ from the equilibrium geometry and the cubic approximation (P3) up to $\pm 70^\circ$. Higher order terms are often included in order also to reproduce vibrational frequencies. Analogous to E_{str} , the higher order force constants are often taken as a fixed fraction of the harmonic constant. The constants beyond third order can rarely be assigned values with high confidence owing to insufficient experimental information. Fixing the higher order constant in terms of the harmonic constant of course reduces the quality of the fit. While a third-order polynomial is capable of reproducing the actual curve very accurately if the cubic constant is fitted independently, the assumption that it is a fixed fraction (independent of the atom type) of the harmonic constant deteriorates the fit, but it still represent an improvement relative to a simple harmonic approximation.




Figure 2.4 The bending energy for CH_4

In the chemically important region below ~ 40 kJ/mol above the bottom of the energy curve, a second-order expansion is normally sufficient.

Angles where the central atom is di- or trivalent (ethers, alcohols, sulfides, amines and enamines) present a special problem. In these cases, an angle of 180° corresponds to an energy maximum, i.e. the derivative of the energy with respect to the angle should be zero and the second derivative should be negative. This may be enforced by suitable boundary conditions on Taylor expansions of at least order three. A third-order

polynomial fixes the barrier for linearity in terms of the harmonic force constant and the equilibrium angle ($\Delta E^\ddagger = k(\theta - \theta_0)^2/6$). A fourth-order polynomial enables an independent fit of the barrier to linearity, but such constrained polynomial fittings are rarely done. Instead, the bending function is taken to be identical for all atom types, for example a fourth-order polynomial with cubic and quartic constants as a fixed fraction of the harmonic constant.

These features are illustrated for H₂O in Figure 2.5, where the “exact” form is taken from a parametric fit to a large amount of spectroscopic data.⁶ The simple harmonic approximation (P2) is seen to be accurate to about $\pm 20^\circ$ from the equilibrium geometry and the cubic approximation (P3) up to $\pm 40^\circ$. Enforcing the cubic polynomial to have a zero derivative at 180° (P3') gives a qualitatively correct behaviour, but reduces the overall fit, although it is still better than a simple harmonic approximation.




Figure 2.5 The bending energy for H₂O

Although such refinements over a simple harmonic potential clearly improve the overall performance, they have little advantage in the chemically important region up to ~ 40 kJ/mol above the minimum. As for the stretch energy term, the energy cost for bending is so large that most molecules only deviate a few degrees from their natural bond angles. This again indicates that including only the harmonic term is adequate for most applications.

As noted above, special atom types are often defined for small rings, owing to the very different equilibrium angles for such rings. In cyclopropane, for example, the carbons are formally sp^3 -hybridized, but have equilibrium CCC angles of 60° , in contrast to 110° in an acyclic system. With a low-order polynomial for the bend energy, the energy cost for such a deformation is large. For cyclobutane, for example, E_{bend} will dominate the total energy and cause the calculated structure to be planar, in contrast to the puckered geometry found experimentally.

For each combination of three atom types, A, B and C, there are at least two bending parameters to be determined, k^{ABC} and θ_0^{ABC} .

2.2.3 The out-of-plane bending energy

If the central B atom in the angle ABC is sp^2 -hybridized, there is a significant energy penalty associated with making the centre pyramidal, since the four atoms prefer to be located in a plane. If the four atoms are exactly in a plane, the sum of the three angles with B as the central atom should be exactly 360° , however, a quite large pyramidalization may be achieved without seriously distorting any of these three angles. Taking the bond distances to 1.5 \AA , and moving the central atom 0.2 \AA out of the plane, only reduces the angle sum to 354.8° (i.e. only a 1.7° decrease per angle). The corresponding out-of-plane angle, χ , is 7.7° for this case. Very large force constants must be used if the ABC, ABD and CBD angle distortions are to reflect the energy cost associated with the pyramidalization. This would have the consequence that the in-plane angle deformations for a planar structure would become unrealistically stiff. Thus a special *out-of-plane energy bend* term (E_{oop}) is usually added, while the in-plane angles (ABC, ABD and CBD) are treated as in the general case above. E_{oop} may be written as a harmonic term in the angle χ (the equilibrium angle for a planar structure is zero) or as a quadratic function in the distance d , as given in eq. (2.8) and shown in Figure 2.6.

$$E_{\text{oop}}(\chi) = k^{\text{B}} \chi^2 \quad \text{or} \quad E_{\text{oop}}(d) = k^{\text{B}} d^2 \quad (2.8)$$

Such energy terms may also be used for increasing the inversion barrier in sp^3 -hybridized atoms (i.e. an extra energy penalty for being planar), and E_{oop} is also sometimes called E_{inv} . Inversion barriers are in most cases (e.g. in amines, NR_3) adequately modelled without an explicit E_{inv} term, the barrier arising naturally from the increase in bond angles upon inversion. The energy cost for non-planarity of sp^2 -hybridized atoms may also be accounted for by an “improper” torsional energy, as described in Section 2.2.4.

For each sp^2 -hybridized atom there is one additional out-of-plane force constant to be determined, k^{B} .




Figure 2.6 Out-of-plane variable definitions

2.2.4 The torsional energy

E_{tors} describes part of the energy change associated with rotation around a B—C bond in a four-atom sequence A—B—C—D, where A—B, B—C and C—D are bonded.




Figure 2.7 Torsional angle definition

Looking down the B—C bond, the torsional angle is defined as the angle formed by the A—B and C—D bonds as shown in Figure 2.7. The angle ω may be taken to be in the range $[0^\circ, 360^\circ]$ or $[-180^\circ, 180^\circ]$.

The torsional energy is fundamentally different from E_{str} and E_{bend} in three aspects:

- (1) A rotational barrier has contributions from both the non-bonded (van der Waals and electrostatic) terms, as well as the torsional energy, and the torsional parameters are therefore intimately coupled to the non-bonded parameters.
- (2) The torsional energy function must be periodic in the angle ω : if the bond is rotated 360° the energy should return to the same value.
- (3) The cost in energy for distorting a molecule by rotation around a bond is often low, i.e. large deviations from the minimum energy structure may occur, and a Taylor expansion in ω is therefore not a good idea.

To encompass the periodicity, E_{tors} is written as a Fourier series.

$$E_{\text{tors}}(\omega) = \sum_{n=1} V_n \cos(n\omega) \quad (2.9)$$

The $n = 1$ term describes a rotation that is periodic by 360° , the $n = 2$ term is periodic by 180° , the $n = 3$ term is periodic by 120° , and so on. The V_n constants determine the size of the barrier for rotation around the B—C bond. Depending on the situation, some of these V_n constants may be zero. In ethane, for example, the most stable conformation is one where the hydrogens are staggered relative to each other, while the eclipsed conformation represents an energy maximum. As the three hydrogens at each end are identical, it is clear that there are three energetically equivalent staggered, and three equivalent eclipsed, conformations. The rotational energy profile must therefore have three minima and three maxima. In the Fourier series only those terms that have $n = 3, 6, 9$, etc., can therefore have V_n constants different from zero.

For rotation around single bonds in substituted systems, other terms may be necessary. In the butane molecule, for example, there are still three minima, but the two *gauche* (torsional angle $\sim \pm 60^\circ$) and *anti* (torsional angle $\sim 180^\circ$) conformations now have different energies. The barriers separating the two *gauche* and the *gauche* and *anti* conformations are also of different height. This may be introduced by adding a term corresponding to $n = 1$.

For the ethylene molecule, the rotation around the C=C bond must be periodic by 180° , and thus only $n = 2, 4$, etc., terms can enter. The energy cost for rotation around a double bond is of course much higher than that for rotation around a single bond in

ethane, which would be reflected in a larger value of the V_2 constant. For rotation around the C=C bond in a molecule such as 2-butene, there would again be a large V_2 constant, analogous to ethylene, but in addition there are now two different orientations of the two methyl groups relative to each other, *cis* and *trans*. The full rotation is periodic with a period of 360° , with deep energy minima at 0° and 180° , but slightly different energies of these two minima. This energy difference would show up as a V_1 constant, i.e. the V_2 constant essentially determines the barrier and location of the minima for rotation around the C=C bond, and the V_1 constant determines the energy difference between the *cis* and *trans* isomers.

Molecules that are composed of atoms having a maximum valence of four (essentially all organic molecules) are with a few exceptions found to have rotational profiles showing at most three minima. The first three terms in the Fourier series in eq. (2.9) are sufficient for qualitatively reproducing such profiles. Force fields that are aimed at large systems often limit the Fourier series to only one term, depending on the bond type (e.g. single bonds only have $\cos(3\omega)$ and double bonds only $\cos(2\omega)$).


Systems with bulky substituents on sp^3 -hybridized atoms are often found to have four minima, the *anti* conformation being split into two minima with torsional angles of approximately $\pm 170^\circ$. Other systems, notably polyfluoroalkanes, also split the *gauche* minima into two, often called *gauche* (angle of approximately $\pm 50^\circ$) and *ortho* (angle of approximately $\pm 90^\circ$) conformations, creating a rotational profile with six minima.⁷ Rotations around a bond connecting sp^3 - and sp^2 -hybridized atoms (such as CH_3NO_2) also display profiles with six minima.⁸ These exceptions from the regular three minima rotational profile around single bonds are caused by repulsive and attractive van der Waals interactions, and can still be modelled by having only terms up to $\cos(3\omega)$ in the torsional energy expression. Higher order terms may be included to modify the detailed shape of the profile, and a few force fields employ terms with $n = 4$ and 6 . Cases where higher order terms probably are necessary are rotation around bonds to octahedral coordinated metals, such as $\text{Ru}(\text{pyridine})_6$ or a dinuclear complex such as $\text{Cl}_4\text{Mo}-\text{MoCl}_4$. Here the rotation is periodic by 90° and thus requires a $\cos(4\omega)$ term.

It is customary to shift the zero point of the potential by adding a factor of one to each term. Most rotational profiles resemble either the ethane or ethylene examples above, and a popular expression for the torsional energy is given in eq. (2.10).

$$E_{\text{tors}}(\omega^{\text{ABCD}}) = \frac{1}{2} V_1^{\text{ABCD}} [1 + \cos(\omega^{\text{ABCD}})] \\ + \frac{1}{2} V_2^{\text{ABCD}} [1 - \cos(2\omega^{\text{ABCD}})] \\ + \frac{1}{2} V_3^{\text{ABCD}} [1 + \cos(3\omega^{\text{ABCD}})] \quad (2.10)$$

The + and – signs are chosen such that the one-fold rotational term has a minimum for an angle of 180° , the two-fold rotational term has minima for angles of 0° and 180° , and the three-fold rotational term has minima for angles of 60° , 180° and 300° (-60°). The factor $1/2$ is included such that the V_i parameters directly give the height of the barrier if only one term is present. A more general form for eq. (2.10) includes a phase factor, i.e. $\cos(n\omega - \tau)$, but for the most common cases of $\tau = 0^\circ$ or 180° , it is completely equivalent to eq. (2.10). Figure 2.8 illustrates the functional behaviour of the three individual terms in eq. (2.10).

The V_i parameters may also be negative, which corresponds to changing the minima on the rotational energy profile to maxima, and vice versa. Most commonly

**Figure 2.8** Torsional energy functions




Figure 2.9 Rotational profile corresponding to eq. 2.10 with $V_1 = 0.5$, $V_2 = -0.2$, $V_3 = 0.5$

encountered rotational profiles can be obtained by combining the three V_i parameters. Figure 2.9 shows an example with one *anti* and two less stable *gauche* minima and with a significant *cis* barrier, corresponding to the combination $V_1 = 0.5$, $V_2 = -0.2$, $V_3 = 0.5$ in eq. (2.10).

As mentioned in Section 2.2.3, the out-of-plane energy may also be described by an “improper” torsional angle. For the example shown in Figure 2.6, a torsional angle ABCD may be defined, even though there is no bond between C and D. The out-of-plane E_{oop} may then be described by an angle ω^{ABCD} , for example as a harmonic function $(\omega - \omega_0)^2$ or eq. (2.10) with a large V_2 constant. Note that the definition of such improper torsional angles is not unique, the angle ω^{ABDC} (for example) is equally good. In practice there is little difference between describing E_{oop} as in eq. (2.8) or as an improper torsional angle.

For each combination of four atom types, A, B, C and D, there are generally three torsional parameters to be determined, V_1^{ABCD} , V_2^{ABCD} and V_3^{ABCD} .

2.2.5 The van der Waals energy

E_{vdw} is the van der Waals energy describing the repulsion or attraction between atoms that are not directly bonded. Together with the electrostatic term E_{el} (Section 2.2.6), it describes the non-bonded energy. E_{vdw} may be interpreted as the non-polar part of the interaction not related to electrostatic energy due to (atomic) charges. This may for example be the interaction between two methane molecules, or two methyl groups at different ends of the same molecule.

E_{vdw} is zero at large interatomic distances and becomes very repulsive for short distances. In quantum mechanical terms, the latter is due to the overlap of the electron clouds of the two atoms, as the negatively charged electrons repel each other. At intermediate distances, however, there is a slight attraction between two such electron clouds from induced dipole–dipole interactions, physically due to electron correlation (discussed in Chapter 4). Even if the molecule (or part of a molecule) has no

permanent dipole moment, the motion of the electrons will create a slightly uneven distribution at a given time. This dipole moment will induce a charge polarization in the neighbour molecule (or another part of the same molecule), creating an attraction, and it can be derived theoretically that this attraction varies as the inverse sixth power of the distance between the two fragments.

The induced dipole–dipole interaction is only the leading term of such induced multipole interactions: there are also contributions from induced dipole–quadrupole, quadrupole–quadrupole, etc., interactions. These vary as R^{-8} , R^{-10} , etc., and the R^{-6} dependence is only the asymptotic behaviour at long distances. The force associated with this potential is often referred to as a “dispersion” or “London” force.⁹ The van der Waals term is the only interaction between rare gas atoms (and thus the reason why say argon can become a liquid and a solid) and it is the main interaction between non-polar molecules such as alkanes.

E_{vdw} is very positive at small distances, has a minimum that is slightly negative at a distance corresponding to the two atoms just “touching” each other, and approaches zero as the distance becomes large. A general functional form that fits these conditions is given in eq. (2.11).

$$E_{\text{vdw}}(R^{\text{AB}}) = E_{\text{repulsion}}(R^{\text{AB}}) - \frac{C^{\text{AB}}}{(R^{\text{AB}})^6} \quad (2.11)$$

It is not possible to derive theoretically the functional form of the repulsive interaction, it is only required that it goes toward zero as R goes to infinity, and it should approach zero faster than the R^{-6} term, as the energy should go towards zero from below.

A popular function that obeys these general requirements is the *Lennard-Jones* (LJ) potential,¹⁰ where the repulsive part is given by an R^{-12} dependence (C_1 and C_2 are suitable constants).

$$E_{\text{LJ}}(R^{\text{AB}}) = \frac{C_1}{(R^{\text{AB}})^{12}} - \frac{C_2}{(R^{\text{AB}})^6} \quad (2.12)$$

The Lennard-Jones potential can also be written as in eq. (2.13).

$$E_{\text{LJ}}(R) = \epsilon \left[\left(\frac{R_0}{R} \right)^{12} - 2 \left(\frac{R_0}{R} \right)^6 \right] \quad (2.13)$$

Here R_0 is the minimum energy distance and ϵ the depth of the minimum. There are no theoretical arguments for choosing the exponent in the repulsive part to be 12, this is purely a computational convenience, and there is evidence that an exponent of 9 or 10 gives better results.

The *Merck Molecular Force Field* (MMFF) uses a generalized Lennard-Jones potential where the exponents and two empirical constants are derived from experimental data for rare gas atoms.¹¹ The resulting *buffered 14-7* potential is shown in eq. (2.14).

$$E_{\text{buf14-7}}(R) = \epsilon \left(\frac{1.07R_0}{R + 0.07R_0} \right)^7 \left(\frac{1.12R_0^7}{R^7 + 0.12R_0^7} - 2 \right) \quad (2.14)$$

From electronic structure theory it is known that the repulsion is due to overlap of the electronic wave functions, and furthermore that the electron density falls off approximately exponentially with the distance from the nucleus (the exact wave function for the hydrogen atom is an exponential function). There is therefore some justification for choosing the repulsive part as an exponential function. The general form of the “Exponential – R^{-6} ” E_{vdw} function, also known as a “Buckingham” or “Hill” type potential,¹² is given in eq. (2.15).

$$E_{\text{Hill}}(R) = Ae^{-BR} - \frac{C}{R^6} \quad (2.15)$$

A , B and C are here suitable constants. It is sometimes written in a slightly more convoluted form as shown in eq. (2.16).

$$E_{\text{Hill}}(R) = \varepsilon \left[\frac{6}{\alpha - 6} e^{\alpha(1-R/R_0)} - \frac{\alpha}{\alpha - 6} \left(\frac{R_0}{R} \right)^6 \right] \quad (2.16)$$

Here R_0 and ε have been defined in eq. (2.13), and α is a free parameter. Choosing an α value of 12 gives a long-range behaviour identical to the Lennard-Jones potential, while a value of 13.772 reproduces the Lennard-Jones force constant at the equilibrium distance. The α parameter may also be taken as a fitting constant. The Buckingham potential has a problem for short interatomic distances where it “turns over”. As R goes toward zero, the exponential becomes a constant while the R^{-6} term goes toward $-\infty$. Minimizing the energy of a structure that accidentally has a very short distance between two atoms will thus result in nuclear fusion! Special precautions therefore have to be taken for avoiding this when using Buckingham-type potentials.

A third functional form, which has an exponential dependence and the correct general shape, is the Morse potential, eq. (2.5). It does not have the R^{-6} dependence at long range, but as mentioned above, in reality there are also R^{-8} , R^{-10} , etc., terms. The D and α parameters of a Morse function describing E_{vdw} will of course be much smaller than for E_{str} , and R_0 will be longer.

For small systems, where accurate interaction energy profiles are available, it has been shown that the Morse function actually gives a slightly better description than a Buckingham potential, which again performs significantly better than a Lennard-Jones 12-6 potential.¹³ This is illustrated for the H_2 —He interaction in Figure 2.10, where the Buckingham and Morse parameters have been derived from the minimum energy and –distance (ε and R_0) and by matching the force constant at the minimum.

The main difference between the three functions is in the repulsive part at short distances, the Lennard-Jones potential is much too hard, and the Buckingham also tends to overestimate the repulsion. Furthermore, it has the problem of “inverting” at short distances. For chemical purposes, however, these “problems” are irrelevant, since energies in excess of 400 kJ/mol are sufficient to break most bonds and will never be encountered in actual calculations. The behaviour in the attractive part of the potential, which is essential for intermolecular interactions, is very similar for the three functions, as shown in Figure 2.11.

Part of the better description for the Morse and Buckingham potentials is due to the fact that they have three parameters, while the Lennard-Jones only employs two. Since the equilibrium distance and the well depth fix two constants, there is no




Figure 2.10 Comparison of E_{vdw} functionals for the H_2-He potential




Figure 2.11 Comparison of E_{vdw} functionals for the attractive part of the H_2-He potential

additional flexibility in the Lennard-Jones function to fit the form of the repulsive interaction.

Most force fields employ the Lennard-Jones potential, despite the known inferiority to an exponential-type function. Let us examine the reason for this in a little more detail.

Essentially all force field calculations use atomic Cartesian coordinates as the variables in the energy expression. To obtain the distance between two atoms one needs to calculate the quantity shown in eq. (2.17).

$$R_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2} \quad (2.17)$$

In the exponential-type potentials, the distance is multiplied by a constant and used as the argument for the exponential. Computationally, it takes significantly more time (typical factor of ~ 5) to perform mathematical operations such as taking the square root and calculating exponential functions than to do simple multiplication and addition. The Lennard-Jones potential has the advantage that the distance itself is not needed, only R raised to *even* powers. Using square roots and exponential functions is thus avoided. The power of 12 in the repulsive part is chosen as it is simply the square of the power of 6 in the attractive part. Calculating E_{vdw} for an exponential-type potential is computationally more demanding than for the Lennard-Jones potential. For large molecules, the calculation of the non-bonded energy in the force field energy expression is by far the most time-consuming, as will be demonstrated in Section 2.5. The difference between the above functional forms is in the repulsive part of E_{vdw} , which is usually not very important. In actual calculations, the Lennard-Jones potential gives results comparable with the more accurate functions, and it is computationally more efficient.

The van der Waals distance, R_0^{AB} , and softness parameters, ϵ^{AB} , depend on both atom types A and B. These parameters are in all force fields written in terms of parameters for the individual atom types. There are several ways of combining atomic parameters to di-atomic parameters, some of them being quite complicated.¹¹ A commonly used method is to take the van der Waals minimum distance as the sum of two van der Waals radii, and the interaction parameter as the geometrical mean of the atomic “softness” constants.

$$\begin{aligned} R_0^{\text{AB}} &= R_0^{\text{A}} + R_0^{\text{B}} \\ \epsilon^{\text{AB}} &= \sqrt{\epsilon^{\text{A}} \epsilon^{\text{B}}} \end{aligned} \quad (2.18)$$

In some force fields, especially those using the Lennard-Jones form in eq. (2.12), the R_0^{AB} parameter is defined as the geometrical mean of atomic radii, implicitly via the geometrical mean rule used for the C_1 and C_2 constants.

For each atom type there are two parameters to be determined, the van der Waals radius and atomic softness, R_0^{A} and ϵ^{A} . It should be noted that since the van der Waals energy is calculated between pairs of atoms, but parameterized against experimental data, the derived parameters represent an *effective* pair potential, which at least partly includes many-body contributions.

The van der Waals energy is the interaction between the electron clouds surrounding the nuclei. In the above treatment, the atoms are assumed to be spherical, but there are two instances where this may not be a good approximation. The first is when one

(or both) of the atoms is hydrogen. Hydrogen has only one electron, which always is involved in bonding to the neighbour atom. For this reason the electron distribution around the hydrogen nucleus is not spherical; rather the electron distribution is displaced towards the other atom. One way of modelling this anisotropy is to displace the position, which is used in calculating E_{vdw} , inwards along the bond. MM2 and MM3 use this approach with a scale factor of ~ 0.92 , i.e. the distance used in calculating E_{vdw} is between points located 0.92 times the X—H bond distance, as shown in Figure 2.12.




Figure 2.12 Illustration of the distance reduction that can be used for E_{vdw} involving hydrogens

The electron density around the hydrogen will also depend significantly on the nature of the X atom. For example, electronegative atoms such as oxygen or nitrogen will lead to smaller effective van der Waals radius for the hydrogen than when it is bonded to carbon. Many force fields therefore have several different types of hydrogen, depending on whether they are bonded to carbon, nitrogen, oxygen, etc., and this may depend further on the type of the neighbour (e.g. alcohol or acid oxygen) – see Table 2.1.

The other case where the spherical approximation may be less than optimal is for atoms having lone pairs, such as oxygen and nitrogen. The lone pair electrons are more diffuse than the electrons involved in bonding, and the atom is thus “larger” in the lone pair direction. Some force fields choose to model lone pairs by assigning *pseudo-atoms* at the lone pair positions. Pseudo-atoms (type 20 in Table 2.1) behave as any other atom type with bond distances and angles, and have their own van der Waals parameters. They are significantly smaller than normal hydrogen atoms, and thus make the oxygen or nitrogen atom “bulge” in the lone pair direction. In some cases, sulfur is also assigned lone pairs, although it has been argued that the second row atoms are more spherical owing to the increased number of electrons, and therefore should not need lone pairs.

It is unclear whether it is necessary to include these effects to achieve good models. The effects are small, and it may be that the error introduced by assuming spherical atoms can be absorbed in the other parameters. Introducing lone pair pseudo-atoms, and making special treatment for hydrogens, again make the time-consuming part of the calculation, the non-bonded energy, even more demanding. Most force fields thus neglect these effects.

Hydrogen bonds require special attention. Such bonds are formed between hydrogens attached to electronegative atoms such as oxygen and nitrogen, and lone pairs, especially on oxygen and nitrogen. They have bond strengths of typically 10–20 kJ/mol, where normal single bonds are 250–450 kJ/mol and van der Waals interactions are 0.5–1.0 kJ/mol. The main part of the hydrogen bond energy comes from electrostatic attraction between the positively charged hydrogen and negatively charged

heteroatom (see Section 2.2.6). Additional stabilization may be modelled by assigning special deep and short van der Waals interactions (via large ϵ and small R_0 parameters). This does not mean that the van der Waals radius for a hydrogen bonded to an oxygen is especially short, since this would affect all van der Waals interactions involving this atom type. Only those *pairs* of interactions that are capable of forming hydrogen bonds are identified (by their atoms types) and the normal E_{vdw} parameters are replaced by special “hydrogen bonding” parameters. The functional form of E_{vdw} may also be different. One commonly used function is a modified Lennard-Jones potential of the form shown in eq. (2.19).

$$E_{\text{H-bond}}(R) = \epsilon \left[5 \left(\frac{R_0}{R} \right)^{12} - 6 \left(\frac{R_0}{R} \right)^{10} \right] \quad (2.19)$$

In some cases $E_{\text{H-bond}}$ also includes a directional term such as $(1 - \cos \omega^{\text{XHY}})$ or $(\cos \omega^{\text{XHY}})^4$ multiplied with the distance-dependent part in eq. (2.19). The current trend seems to be that force fields are moving away from such specialized parameters and/or functional forms, and instead are accounting for hydrogen bonding purely by electrostatic interactions.

2.2.6 The electrostatic energy: charges and dipoles

The other part of the non-bonded interaction is due to internal (re)distribution of the electrons, creating positive and negative parts of the molecule. A carbonyl group, for example, has a negatively charged oxygen and a positively charged carbon. At the lowest approximation, this can be modelled by assigning (partial) charges to each atom. Alternatively, the bond may be assigned a bond dipole moment. These two descriptions give similar (but not identical) results. Only in the long distance limit of interaction between such molecules do the two descriptions give identical results.

The interaction between point charges is given by the Coulomb potential, with ϵ being a dielectric constant.

$$E_{\text{el}}(R^{\text{AB}}) = \frac{Q^{\text{A}} Q^{\text{B}}}{\epsilon R^{\text{AB}}} \quad (2.20)$$

The atomic charges can be assigned by empirical rules,¹⁴ but are more commonly assigned by fitting to the electrostatic potential calculated by electronic structure methods, as discussed in the next section. Since hydrogen bonding is to a large extent due to attraction between the electron-deficient hydrogen and an electronegative atom such as oxygen or nitrogen, a proper choice of partial charges may adequately model this interaction.

The MM2 and MM3 force fields use a bond dipole description for E_{el} . The interaction between two dipoles is given by eq. (2.21).

$$E_{\text{el}}(R^{\text{AB}}) = \frac{\mu^{\text{A}} \mu^{\text{B}}}{\epsilon (R^{\text{AB}})^3} (\cos \chi - 3 \cos \alpha_{\text{A}} \cos \alpha_{\text{B}}) \quad (2.21)$$

The angles χ , α_{A} and α_{B} are defined as shown in Figure 2.13.




Figure 2.13 Definition of variables for a dipole–dipole interaction

When properly parameterized, there is little difference in the performance of the two ways of representing E_{el} . There are exceptions where two strong bond dipoles are immediate neighbours (for example α -halogen ketones). The dipole model will here lead to a stabilizing electrostatic interaction for a *transoid* configuration (torsional angle of 180°), while the atomic charge model will be purely repulsive for all torsional angles (since all 1,3-interactions are neglected). In either case, however, a proper rotational profile may be obtained by suitable choices of the constants in E_{tors} . The atomic charge model is easier to parameterize by fitting to an electronic wave function, and is preferred by almost all force fields.

The “effective” dielectric constant ϵ can be included to model the effect of surrounding molecules (solvent) and the fact that interactions between distant sites may be “through” part of the same molecule, i.e. a polarization effect. A value of 1 for ϵ corresponds to a vacuum, while a large ϵ reduces the importance of long-range charge–charge interactions. Typically, a value between 1 and 4 is used, although there is little theoretical justification for any specific value. In some applications the dielectric constant is made distance dependent (e.g. $\epsilon = \epsilon_0 R^{\text{AB}}$, changing the Coulomb interaction to $Q_A Q_B / \epsilon_0 (R^{\text{AB}})^2$) to model the “screening” by solvent molecules. There is little theoretical justification for this, but it increases the efficiency of the calculation as a square root operation is avoided (discussed in Section 2.2.5), and it seems to provide reasonable results.

How far apart (in terms of number of bonds between them) should two atoms be before a non-bonded energy term contributes to E_{FF} ? It is clear that two atoms directly bonded should not have an E_{vdw} or E_{el} term – their interaction is described by E_{str} . It is also clear that the interaction between two hydrogens at each end of say $\text{CH}_3(\text{CH}_2)_{50}\text{CH}_3$ is identical to the interaction between two hydrogens belonging to two different molecules, and they therefore should have an E_{vdw} and an E_{el} term. But where should the dividing line be? Most force fields included E_{vdw} and E_{el} for atom pairs that are separated by three bonds or more, although 1,4-interactions are in many cases scaled down by a factor between 1 and 2. This means that the rotational profile for an A–B–C–D sequence is determined both by E_{tors} and E_{vdw} and E_{el} terms for the A–D pair. In a sense, E_{tors} may be considered as a correction necessary for obtaining the correct rotational profile once the non-bonded contribution has been accounted

for. Some force fields have chosen also to include E_{vdw} for atoms that are 1,3 with respect to each other – these are called *Urey–Bradley* force fields. In this case, the energy required to bend a three atom sequence is a mixture of E_{bend} and E_{vdw} . Most modern force fields calculate E_{str} between all atoms pairs that are 1,2 with respect to each other in terms of bonding, E_{bend} for all pairs that are 1,3, E_{tors} between all pairs that are 1,4, and E_{vdw} and E_{el} between all pairs that are 1,4 or higher.

For polar molecules, the electrostatic energy dominates the force field energy function, and an accurate representation is therefore important for obtaining good results. Within the partial charge model, the atomic charges are normally assigned by fitting to the *molecular electrostatic potential* (MEP) calculated by an electronic structure method. The electrostatic potential ϕ_{esp} at a point \mathbf{r} is given by the nuclear charges and electronic wave function as shown in eq. (2.22).

$$\phi_{\text{esp}}(\mathbf{r}) = \sum_a^{N_{\text{nuc}}} \frac{Z_a}{|\mathbf{R}_a - \mathbf{r}|} - \int \frac{\Psi^2(\mathbf{r}')}{|\mathbf{r}' - \mathbf{r}|} d\mathbf{r}' \quad (2.22)$$

The fitting is done by minimizing an error function of the form shown in eq. (2.23), under the constraint that the sum of the partial charges Q_i is equal to the total molecular charge. The electrostatic potential is sampled at a few thousand points in the near vicinity of the molecule.

$$\text{ErrF}(\mathbf{Q}) = \sum_r^{N_{\text{points}}} \left(\phi_{\text{esp}}(\mathbf{r}) - \sum_a^{N_{\text{atoms}}} \frac{Q_a(\mathbf{R}_a)}{|\mathbf{R}_a - \mathbf{r}|} \right)^2 \quad (2.23)$$

The set of linear equations arising from minimizing the error function are often poorly conditioned, i.e. the calculated partial charges are sensitive to small details in the fitting data.¹⁵ The physical reason for this is that the electrostatic potential is primarily determined by the atoms near the surface of the molecule, while the atoms buried within the molecule have very little influence on the external electrostatic potential. A straightforward fitting therefore often results in unrealistically large charges for the non-surface atoms. The problem can to some extent be avoided by adding a hyperbolic penalty term for having non-zero partial charges, since this ensures that only those charges that are important for the electrostatic potential have values significantly different from zero.¹⁶ This *Restrained ElectroStatic Potential* (RESP) fitting scheme has been used in for example the AMBER force field. Other constraints are also often imposed, such as constraining the charges on the three hydrogens in a methyl group to be equal, or the sum of all charges in a subgroups (such as a methyl group or an amino acid) to be zero.

The non-bonded energies are modelled by pair-interactions, but if the parameters are obtained by fitting to experimental data, they will include the average part of many-body effects. For the electrostatic energy, the three-body effect may be considered as the interaction between two atomic charges being modified because a third atom or molecule polarizes the charges. The dipole moment of water, for example, increases from 1.8 debye in the gas phase to an effective value of 2.5 debye in the solid state.¹⁷ Even when the partial charges are obtained by fitting to electronic structure data, the average many-body effect is often approximately accounted for by increasing the fitted charges by 10–15%, or by fitting to data that are known to overestimate the polarity (e.g. Hartree–Fock results).

2.2.7 The electrostatic energy: multipoles and polarizabilities

Obtaining a good description of the electrostatic interaction between molecules (or between different parts of the same molecule) is one of the big problems in force field work. Many commercial applications of force field methods are aimed at designing molecules that interact in a specific fashion. Such interactions are usually pure non-bonded, and for polar molecules such as proteins, the electrostatic interaction is very important.

The modelling of the electrostatic energy by (fixed) atomic charges has four main deficiencies:

- (1) The fitting of atomic charges to electrostatic potentials focuses on reproducing *intermolecular* interactions, but the electrostatic energy also plays a strong role in the *intramolecular* energy, which determines conformational energies. For polar molecules the (relative) conformational energies are therefore often of significantly lower accuracy than for non-polar systems.
- (2) The partial charge model gives a rather crude representation of the electrostatic potential surrounding a molecule, with errors often being in the 10–20 kJ/mol range. For a given (fixed) geometry, the molecular electrostatic potential can be improved either by adding non-nuclear-centred partial charges, or by including higher order (dipole, quadrupole, etc.) electric moments.
- (3) The coupling of electric charges and higher order moments with the geometry is neglected. Analysis has shown that both partial charges and higher order electric moments depend significantly on the geometry, i.e. these quantities do not fulfil the requirement of “transferability”.
- (4) Only two-body interactions are included, but for polar species the three-body contribution is quite significant, perhaps 10–20% of the two-body term.¹⁸ A rigorous modelling of these effects requires inclusion of atomic polarizabilities, but can be partly included in the two-body interaction by empirically increasing the interaction by 10–20%.

The non-bonded terms together with the torsional energy determine the internal (conformational) degrees of freedom, and the torsional parameters are usually obtained as a residual correction to reproduce rotational barriers and energetic preferences after the assignment of the non-bonded parameters. The electrostatic energy is unimportant for non-polar systems such as hydrocarbons and is therefore often completely neglected. The conformational space in such cases is consequently determined by E_{vdw} and E_{tors} . Since the van der Waals interaction is short-ranged, this means that the transferability assumption inherent in force field methods is valid, and torsional parameters determined for small model systems can also be used for predicting conformations for large systems.

For polar systems, however, the long-range electrostatic interaction is often the dominating energy term, and the transferability of torsional parameters determined for small model systems becomes more problematic. Clearly the variation of the electrostatic energy with the geometry must be accurately modelled in order for the torsional parameters to be sufficiently transferable. The use of partial atomic charges determined by fitting to the electrostatic potential is an integral part of most force

fields, but only a few experimental force fields are at present capable of including higher order electric moment and/or atomic polarizabilities.

The representation of the electrostatic potential for a fixed geometry can be systematically improved by including non-atom-centred charges¹⁹ or by including higher order moments. The *Distributed Multipole Analysis* (DMA) developed by A. Stone provides an exact method for expanding the electrostatic potential in terms of multipole moments distributed at a number of positions within the molecule, and these moments can be derived directly from the wave function without a fitting procedure (see Section 9.2).²⁰ By restricting the positions to only atoms and bond midpoints, an accurate representation of the electrostatic potential can be obtained by including up to quadrupole moments at each site. The DMA-derived multipoles are sensitive to the reference data (i.e. wave function quality and basis set), and a better stability can be obtained by fitting a set of multipoles to the electrostatic potential on a molecular surface,²¹ by fitting multipoles to the DMA multipoles,²² or by fitting atomic charges to match the atomic multipole moments.²³ Such fitted multipole methods typically reduce the required moments by one or two, i.e. fitted charges and dipoles can reproduce DMA results including up to quadrupoles or octopoles. Unfortunately, the addition of non-nuclear-centred charges or multipoles significantly increases the computational time for force field calculations as they add to the non-bonded terms, the number of which grows as N_{atom}^2 .

The coupling of the electrostatic energy with the geometry can be modelled by the *fluctuating charge* model, where the charges are allowed to adjust to changes in the geometry based on electronegativity equalization.²⁴ Consider the following expansion of the energy as a function of the number of electrons N .

$$E = E_0 + \frac{\partial E}{\partial N} \Delta N + \frac{1}{2} \frac{\partial^2 E}{\partial N^2} (\Delta N)^2 + \dots \quad (2.24)$$

The first derivative is the *electronegativity* χ , except for a change in sign ($\partial E / \partial N = -\chi$), while the second derivative is the *hardness* η . Although these have well-defined finite-value approximations in terms of ionization potentials and electron affinities (Section 15.2), they are usually treated as empirical parameters within a force field environment. For an atom in a molecule, the change in the number of electrons is equal to minus the change in the atomic charge ($-\Delta N = \Delta Q$). Taking the expansion point as the one with no atomic charges gives eq. (2.25).

$$E = \chi Q + \frac{1}{2} \eta Q^2 + \dots \quad (2.25)$$

Terminating the expansion at second order, adding a term corresponding to the interaction of the charges with an external potential ϕ , and summing over all sites gives an expression for the electrostatic energy.

$$E_{\text{el}} = \sum_a \phi_a Q_a + \sum_a \chi_a Q_a + \frac{1}{2} \sum_{ab} \eta_{ab} Q_a Q_b \quad (2.26)$$

Switching to a vector-matrix notation and requiring that the energy is stationary with and without an external potential gives eq. (2.27).

$$\begin{aligned}\left.\frac{\partial E_{\text{el}}}{\partial \mathbf{Q}}\right|_{\phi \neq 0} &= \phi + \boldsymbol{\chi} + \boldsymbol{\eta} \mathbf{Q} = 0 \\ \left.\frac{\partial E_{\text{el}}}{\partial \mathbf{Q}}\right|_{\phi=0} &= \boldsymbol{\chi} + \boldsymbol{\eta} \mathbf{Q}^0 = 0\end{aligned}\quad (2.27)$$

Subtraction of these two equations leads to a recipe for the charge transfer due to the external potential.

$$\Delta \mathbf{Q} = -\boldsymbol{\eta}^{-1} \phi \quad (2.28)$$

Since the potential depends on the charges at all sites, this must be solved iteratively.

$$\phi(\mathbf{r}) = \sum_a \frac{Q_a}{|\mathbf{r} - \mathbf{R}_a|} \quad (2.29)$$

Once the iterations have converged, the electrostatic energy is given by the simple Coulomb form in eq. (2.20).

The fluctuating charge model is a simple way of introducing a coupling between the electrostatic energy and geometry, but it should only be considered as a first approximation, as it is unable to account for example for the charge polarization of a planar molecule where the external field is perpendicular to the molecule plane. An explicit incorporation of polarization can be done by including an atomic polarization tensor,²⁵ and this also implicitly accounts for some of the geometry dependence of the atomic charges. The polarization contribution to the electrostatic interaction is at the lowest order given by a dipolar term ($\boldsymbol{\mu}_{\text{ind}}$) arising from the electric field ($\mathbf{F} = \partial\phi/\partial\mathbf{r}$) created by the electric moments at other sites multiplied by the polarizability tensor ($\boldsymbol{\alpha}$).²⁶

$$\begin{aligned}\boldsymbol{\mu}_{\text{ind}} &= \boldsymbol{\alpha} \mathbf{F} \\ E_{\text{el}}^{\text{pol}} &= \frac{1}{2} \boldsymbol{\mu}_{\text{ind}} \mathbf{F} = \frac{1}{2} \boldsymbol{\alpha} \mathbf{F}^2\end{aligned}\quad (2.30)$$

Note that since the (atomic) hardness is inversely related to the average polarizability, the charge transfer in eq. (2.28) is essentially the average polarizability times the potential. As each of the atoms contributes to the electric field at a given position (eq. (2.29) with additional contributions from dipole moments), the set of atomic dipoles must be solved self-consistently by iterative methods. For molecular dynamics simulations, the change in the induced dipoles or charges with geometry can be treated by an extended Lagrange method (Section 14.2.5), where fictive masses are assigned to the dipoles or charges, and progressed along with the other variables in a simulation.²⁴

When adding multipole and/or polarizability terms, a decision has to be made on which interactions to include and which to neglect, and either the multipole order or the distance dependence of the interaction can be used as a guiding criterion. The distance dependence on the interactions between multipoles is given in Table 2.2.

Table 2.2 Distance dependence of multipole interactions

	Q	μ	Θ	Ξ
Q	R^{-1}	R^{-2}	R^{-3}	R^{-4}
μ	R^{-2}	R^{-3}	R^{-4}	R^{-5}
Θ	R^{-3}	R^{-4}	R^{-5}	R^{-6}
Ξ	R^{-4}	R^{-5}	R^{-6}	R^{-7}

If all interactions between multipoles up to quadrupoles are included, then only some of the interactions having R^{-4} and R^{-5} distance dependencies are accounted for. Alternatively, if the distance dependence is taken as the deciding factor, then quadrupoles are required for including all interaction of order R^{-3} or lower, but the dipole–quadrupole and quadrupole–quadrupole interactions should not be included.

The polarizability in eq. (2.30) corresponds to an electric field inducing a dipole moment, but higher order polarizabilities giving induced quadrupole and octopole moments are also possible, although these will usually be significantly smaller. It is at present unclear how many multipole moments, which interaction and level of polarizability should be included for a balanced description. The charge-induced dipole interaction has a distance dependence of R^{-4} , while the dipole-induced dipole interaction is R^{-6} , suggesting that the former should be included when quadrupole moments are incorporated. There is also no clear picture of what kind of improvement for the calculated results can be obtained by including these higher order effects.

Incorporation of electric multipole moments, fluctuating charges and atomic polarizabilities significantly increases the number of fitting parameters for each atom type or functional unit, and only electronic structure methods are capable of providing a sufficient number of reference data. Electronic structure calculations, however, automatically include all of the above effects, and also have higher order terms. The data must therefore be “deconvoluted” in order to extract suitable multipole and polarization parameters for use in force fields.²⁷ A calculated set of distributed dipole moments, for example, must be decomposed into permanent and induced contributions, based on an assigned polarizability tensor. Furthermore, only the lowest non-vanishing multipole moment is independent of the origin of the coordinate system, i.e. for a non-centrosymmetric neutral molecule the dipole moment is unique, but the quadrupole moment depends on where the origin is placed.

It should be noted that the transfer of polarization data from gas-phase calculations to a condensed phase may lead to errors. The close spatial arrangement of the molecules in a condensed phase will display quantum mechanical exchange phenomena, which will reduce the effective polarization. The possibility of charge transfer between molecules, however, can lead to an enhancement of the polarization relative to the gas-phase result.²⁸ It is therefore likely that polarizable force fields will need to be re-tuned to reproduce experimental results, unless of course the quantum mechanical effects are incorporated directly.

The addition of multipole moments increases the computational time for the electrostatic energy, since there now are several components for each pair of sites, and for multipoles up to quadrupoles the evaluation time increases by almost an order of magnitude. If bond midpoints are added as multipole sites, the number of non-bonded terms furthermore increases by a factor of ~ 4 over only using atomic sites. Inclusion of polarization further increases the computational complexity by adding an iterative procedure for evaluating the induced dipole moments, although recent advances have reduced the computational overhead to only a factor of 2 over a fixed charge model.²⁹ Advanced force fields with a description beyond fixed partial charges of the electrostatic energy have consequently only seen limited use so far. Nevertheless, the neglect of multipole moments and polarization is probably the main limitation of modern force

fields, at least for polar systems, and future improvements in force field techniques must include such effects.

2.2.8 Cross terms

The first five terms in the general energy expression, eq. (2.1), are common to all force fields. The last term, E_{cross} , covers coupling between these fundamental, or diagonal, terms. Consider for example a molecule such as H_2O . It has an equilibrium angle of 104.5° and an O—H distance of 0.958 \AA . If the angle is compressed to say 90° , and the optimal bond length is determined by electronic structure calculations, the equilibrium distance becomes 0.968 \AA , i.e. slightly longer. Similarly, if the angle is widened, the lowest energy bond length becomes shorter than 0.958 \AA . This may qualitatively be understood by noting that the hydrogens come closer together if the angle is reduced. This leads to an increased repulsion between the hydrogens, which can be partly alleviated by making the bonds longer. If only the first five terms in the force field energy are included, this coupling between bond distance and angle cannot be modelled. It may be taken into account by including a term that depends on both bond length and angle. E_{cross} may in general include a whole series of terms that couple two (or more) of the bonded terms.

The components in E_{cross} are usually written as products of first-order Taylor expansions in the individual coordinates. The most important of these is the stretch/bend term, which for an A—B—C sequence may be written as in eq. (2.31).

$$E_{\text{str/bend}} = k^{\text{ABC}}(\theta^{\text{ABC}} - \theta_0^{\text{ABC}})[(R^{\text{AB}} - R_0^{\text{AB}}) - (R^{\text{BC}} - R_0^{\text{BC}})] \quad (2.31)$$

Other examples of such cross terms are given in eq. (2.32).

$$\begin{aligned} E_{\text{str/str}} &= k^{\text{ABC}}(R^{\text{AB}} - R_0^{\text{AB}})(R^{\text{BC}} - R_0^{\text{BC}}) \\ E_{\text{bend/bend}} &= k^{\text{ABCD}}(\theta^{\text{ABC}} - \theta_0^{\text{ABC}})(\theta^{\text{BCD}} - \theta_0^{\text{BCD}}) \\ E_{\text{str/tors}} &= k^{\text{ABCD}}(R^{\text{AB}} - R_0^{\text{AB}})\cos(n\omega^{\text{ABCD}}) \\ E_{\text{bend/tors}} &= k^{\text{ABCD}}(\theta^{\text{ABC}} - \theta_0^{\text{ABC}})\cos(n\omega^{\text{ABCD}}) \\ E_{\text{bend/tors/bend}} &= k^{\text{ABCD}}(\theta^{\text{ABC}} - \theta_0^{\text{ABC}})(\theta^{\text{BCD}} - \theta_0^{\text{BCD}})\cos(n\omega^{\text{ABCD}}) \end{aligned} \quad (2.32)$$

The constants involved in these cross terms are usually not taken to depend on all the atom types involved in the sequence. The stretch/bend constant, for example, in principle depends on all three atoms, A, B and C. However, it is usually taken to depend only on the central atom, i.e. $k^{\text{ABC}} = k^{\text{B}}$, or chosen as a universal constant independent of atom type. It should be noted that cross terms of the above type are inherently unstable if the geometry is far from equilibrium. Stretching a bond to infinity, for example, will make $E_{\text{str/bend}}$ go toward $-\infty$ if θ is less than θ_0 . If the bond stretch energy itself is harmonic (or quartic) this is not a problem as it approaches $+\infty$ faster. However, if a Morse type potential is used, special precautions will have to be made to avoid long bonds in geometry optimizations and simulations.

Another type of correction, which is related to cross terms, is modification of parameters based on atoms not directly involved in the interaction described by the parameter. Carbon-carbon bond lengths, for example, become shorter if there are

electronegative atoms present at either end. Such electronegativity effects may be modelled by adding a correction to the natural bond length R_0^{AB} based on the atom C attached to the A—B bond.³⁰

$$R_0^{AB-C} = R_0^{AB} + \Delta R_0^C \quad (2.33)$$

Other effects, such as hyperconjugation, can be modelled by allowing the natural bond length to depend on the adjacent torsional angle.³¹ The hyperconjugation effect can be thought of as weakening of a σ -bond by donation of electron density into an adjacent empty π^* -bond, as illustrated in Figure 2.14.




Figure 2.14 Illustrating the elongation of the C—H bond by hyperconjugation

Since the conjugation can only take place when the σ -orbital is aligned with the π -system, the resulting bond elongation will depend on the torsional angle, which can be modelled by an energy term such as in eq. (2.34).

$$\begin{aligned} R_0^{AB} &= R_0^{AB} + \Delta R_0^\omega \\ \Delta R_0^\omega &= k(1 - \cos 2\omega^{ABCD}) \end{aligned} \quad (2.34)$$

2.2.9 Small rings and conjugated systems

It has already been mentioned that small rings present a problem as their equilibrium angles are very different from their acyclic cousins. One way of alleviating this problem is to assign new atom types. If a sufficient number of cross terms is included, however, the necessary number of atom types can actually be reduced. Some force fields have only one sp^3 -carbon atom type, covering bonding situations from cyclopropane to linear alkanes with the same set of parameters. The necessary flexibility in the parameter space is here transferred from the atom types to the parameters in the cross terms, i.e. the cross terms modify the diagonal terms such that a more realistic behaviour is obtained for large deviations from the natural value.

One additional class of bonding that requires special consideration in force fields is conjugated systems. Consider for example 1,3-butadiene. According to the MM2 type convention (Table 2.1), all carbon atoms are of type 2. This means that the same set of parameters is used for the terminal and central C—C bonds. Experimentally, the bond lengths are 1.35 and 1.47 Å, i.e. very different, which is due to the partial delocalization of the π -electrons in the conjugated system.³² The outer C=C bond is slightly reduced in double bond character (and thus has a slightly longer bond length than in ethylene) while the central bond is roughly halfway between a single and a double bond. Similarly, without special precautions, the barriers for rotation around the terminal and central bonds are calculated to be the same, and assume a value characteristic of a localized double bond, ~230 kJ/mol. Experimentally, however, the rotational barrier for the central bond is only ~25 kJ/mol.³³

There are two main approaches for dealing with conjugated systems. One is to identify certain bonding combinations and use special parameters for these cases, analogously to the treatment of hydrogen bonds in E_{vdw} . If four type 2 carbons are located in a linear sequence, for example, they constitute a butadiene unit and special stretch and torsional parameters should be used for the central and terminal bonds. Similarly, if six type 2 carbons are in a ring, they constitute an aromatic ring and a set of special aromatic parameters are used. Or the atom type 2 may be changed to a type 50, identifying from the start that these carbons should be treated with a different parameter set. The main problem with this approach is that there are many such “special” cases requiring separate parameters. Three conjugated double bonds, for example, may either be linearly or cross-conjugated (1,3,5-hexatriene and 2-vinyl-1,3-butadiene), each requiring a set of special parameters different from those used for 1,3-butadiene. The central bond in biphenyl will be different from the central bond in 1,3-butadiene. Modelling the bond alterations in fused aromatics such as naphthalene or phenanthrene requires complicated bookkeeping to keep track of all the different bond lengths, etc.

The other approach, which is somewhat more general, is to perform a simple electronic structure calculation to determine the degree of delocalization within the π -system. This approach is used in the MM2 and MM3 force fields, often denoted MMP2 and MMP3.³⁴ The electronic structure calculation is of the Pariser–Pople–Parr (PPP) type (Section 3.10.3), which is only slightly more advanced than a simple Hückel calculation. From the calculated π -molecular orbitals, the π -bond order ρ for each bond can be calculated. Since the bond length, force constant and rotational energy depend on the π -bond order, these constant can be parameterized based on the calculated ρ . The connections used in MMP2 are given in eq. (2.35), with n_i being the number of electrons in the i th MO and β_{BC} being a resonance parameter.

$$\begin{aligned}\rho_{\text{AB}} &= \sum_i^{N_{\text{occ}}} n_i C_{\text{Ai}} C_{\text{Bi}} \\ R_0^{\text{AB}} &= 1.503 - 0.166\rho_{\text{AB}} \\ k^{\text{AB}} &= 5.0 + 4.6\rho_{\text{AB}} \\ V_2^{\text{ABCD}} &= 62.5\rho_{\text{BC}}\beta_{\text{BC}}\end{aligned}\tag{2.35}$$

The natural bond length varies between 1.503 Å and 1.337 Å for bond orders between 0 and 1 – these are the values for pure single and double bonds between two sp^2 -carbons. Similarly, the force constant varies between the values used for isolated single and double bonds. The rotational barrier for an isolated double bond is 250 kJ/mol, since there are four torsional contributions for a double bond.

This approach allows a general treatment of all conjugated system, but requires the addition of a second level of iterations in a geometry optimization. At the initial geometry, a PPP calculation is performed, the π -bond orders are calculated and suitable bond parameters (R_0^{AB} , k^{AB} and V_2^{ABCD}) are assigned. These parameters are then used for optimizing the geometry. The optimized geometry will usually differ from the initial geometry, thus the parameters used in the optimization are no longer valid. At the “optimized” geometry, a new PPP calculation is performed and a new set of parameters derived. The structure is re-optimized, and a new PPP calculation is carried




Figure 2.15 Illustration of the two-level optimization involved in a MMP2 calculation

out, etc. This is continued until the geometry change between two macro iterations is negligible.

For commonly encountered conjugated systems such as butadiene and benzene, the *ad hoc* assignment of new parameters is usually preferred as it is simpler than the computationally more demanding PPP method. For less common conjugated systems, the PPP approach is more elegant and has the definite advantage that the common user does not need to worry about assigning new parameters. If the system of interest contains conjugation, and a force field that uses the parameter replacement method is chosen, the user should check that proper bond lengths and reasonable rotational barriers are calculated (i.e. that the force field has identified the conjugated moiety and contains suitable substitution parameters). Otherwise, very misleading results may be obtained without any indication from the force field of problems.

2.2.10 Comparing energies of structurally different molecules

The force field energy function has a zero point defined implicitly by the zero points chosen for each of the terms. For the three bonding terms, stretch, bend and torsion, this is at the bottom of the energy curve (natural bond lengths and angles), while for the two non-bonded terms, it is at infinite separation. The zero point for the total force field energy is therefore a hypothetical system, where all the bond distances, angles and torsional angles are at their equilibrium values, and at the same time, all the non-bonded atoms are infinitely removed from each other. Except for small systems such as CH_4 , where there are no non-bonded terms, this is a physically unattainable situation. The force field energy, E_{FF} , is often called the *steric energy*, as it is in some sense the excess energy relative to a hypothetical molecule with non-interacting fragments, but the *numerical value of the force field function has no physical meaning!*

Relative values, however, should ideally reflect conformational energies. If all atom and bond types are the same, as in cyclohexane and methyl-cyclopentane, the energy functions have the same zero point, and relative stabilities can be directly compared. This is a rather special situation, however, and stabilities of different molecules can normally not be calculated by force field techniques. For comparing relative stabilities of chemically different molecules such as dimethyl ether and ethyl alcohol, or for comparing with experimental heat of formations, the zero point of the energy scale must be the same.

In electronic structure calculations, the zero point for the energy function has all particles (electrons and nuclei) infinitely removed from each other, and this common reference state allows energies for systems with different numbers of particles to be directly compared. If the same reference is used in force field methods, the energy function becomes an absolute measure of molecular stability. The difference relative to the normal reference state for force field functions is the sum of all bond dissociation energies, at least for a simple diagonal force field. If correction terms are added to the

normal force field energy function based on average bond dissociation energies for each bond type, the energy scale become absolute, and can be directly compared with e.g. ΔH_f . Such bond dissociation energies again rest on the assumption of transferability, for example that all C—H bonds have dissociation energies close to 400 kJ/mol. In reality, the bond dissociation energy for a C—H bond depends on the environment: the value for the aldehyde C—H bond in CH_3CHO is 366 kJ/mol while it is 410 kJ/mol for C_2H_6 .³⁵ This can be accounted for approximately by assigning an average bond dissociation energy to a C—H bond, and a smaller correction based on larger structural units, such as CH_3 and CHO groups. The MM2 and MM3 force fields use an approach where such bond dissociation energies and structural factors are assigned based on fitting to experimental data, and this approach is quite successful for reproducing experimental ΔH_f values.

$$\Delta H_f = E_{\text{FF}} + \sum_{\text{AB}}^{\text{bonds}} \Delta H_{\text{AB}} + \sum_{\text{G}}^{\text{groups}} \Delta H_{\text{G}} \quad (2.36)$$

These heat of formation parameters may be considered as shifting the zero point of E_{FF} to a common origin. Since corrections from larger moieties are small, it follows that energy differences between systems having the same groups (for example methylcyclohexane and ethylcyclopentane) can be calculated directly from differences in steric energy.

If the heat of formation parameters are derived based on fitting to a large variety of compounds, a specific set of parameters is obtained. A slightly different set of parameters may be obtained if only certain “strainless” molecules are included in the parameterization. Typically molecules such as straight-chain alkanes and cyclohexane are defined to be strainless. By using these strainless heat of formation parameters, a strain energy may be calculated as illustrated in Figure 2.16.




Figure 2.16 Illustrating the difference between steric energy and heat of formation

Deriving such heat of formation parameters requires a large body of experimental ΔH_f values, and for many classes of compounds there are not sufficient data available. Only a few force fields, notably MM2 and MM3, also attempt to parameterize heats of formation. Most force fields are only concerned with reproducing geometries and possibly conformational relative energies, for which the steric energy is sufficient.

2.3 Force Field Parameterization

Having settled on the functional description and a suitable number of cross terms, the problem of assigning numerical values to the parameters arises. This is by no means trivial.³⁶ Consider for example MM2(91) with 71 atom types. Not all of these can form

stable bonds with each other, hydrogens and halogens can only have one bond, etc. For the sake of argument, however, assume that the effective number of atom types capable of forming bonds between each other is 30.

- Each of the 71 atom types has two van der Waals parameters, R_0^A and ϵ^A , giving 142 parameters.
- There are approximately $\frac{1}{2} \times 30 \times 30 = 450$ possible different E_{str} terms, each requiring at least two parameters, k^{AB} and R_0^{AB} , for a total of at least 900 parameters.
- There are approximately $\frac{1}{2} \times 30 \times 30 \times 30 = 13\,500$ possible different E_{bend} terms, each requiring at least two parameters, k^{ABC} and θ_0^{ABC} , for a total of at least 27\,000 parameters.
- There are approximately $\frac{1}{2} \times 30 \times 30 \times 30 \times 30 = 405\,000$ possible different E_{tors} terms, each requiring at least three parameters, V_1^{ABCD} , V_2^{ABCD} and V_3^{ABCD} , for a total of at least 1\,215\,000 parameters.
- Cross terms may add another million possible parameters.

To achieve just a rudimentary assignment of the value of one parameter, at least 3–4 independent data should be available. To parameterize MM2 for all molecules described by the 71 atom types would thus require of the order of 10^7 independent experimental data, not counting cross terms, which clearly is impossible. Furthermore, the parameters that are the most numerous, the torsional constants, are also the ones that are the hardest to obtain experimental data for. Experimental techniques normally probe a molecule near its equilibrium geometry. Getting energetical information about the whole rotational profile is very demanding and has only been done for a handful of small molecules. In recent years, it has therefore become common to rely on data from electronic structure calculations to derive force field parameters. Calculating for example rotational energy profiles is computationally fairly easy. The so-called “Class II” and “Class III” force fields rely heavily on data from electronic structure calculations to derive force field parameters, especially the bonded parameters (stretch, bend and torsional).

While the non-bonded terms are relatively unimportant for the “local” structure, they are the only contributors to intermolecular interactions, and the major factor in determining the global structure of a large molecule, such as protein folding. The electrostatic part of the interaction may be assigned based on fitting parameters to the electrostatic potential derived from an electronic wave function, as discussed in Section 2.2.6. The van der Waals interaction, however, is difficult to calculate reliably by electronic structure methods, requiring a combination of electron correlation and very large basis sets, and these parameters are therefore usually assigned based on fitting to experimental data for either the solid or liquid state.³⁷

For a system containing only a single atom type (e.g. liquid argon), the R_0 (atomic size) and ϵ (interaction strength) parameters can be determined by requiring that the experimental density and heat of evaporation are reproduced, respectively. Since the parameterization implicitly takes many-body effects into account, a (slightly) different set of van der Waals parameters will be obtained if the parameterization instead focuses on reproducing the properties of the crystal phase. For systems where several atom types are involved (e.g. water), there are two van der Waals parameters for each atom type, and the experimental density and heat of evaporation alone therefore give insufficient data for a unique assignment of all parameters. Although one may include

additional experimental data, for example the variation of the density with temperature, this still provides insufficient data for a general system containing many atom types. Furthermore, it is possible that several combinations of van der Waals parameters for different atoms may be able to reproduce properties of a liquid, i.e. even if there are sufficient experimental data, the derived parameter set may not be unique. One approach for solving this problem is to use electronic structure methods to determine *relative* values for van der Waals parameters, for example using a neon atom as the probe, and determine the absolute values by fitting to experimental values.³⁸

An alternative procedure is to derive the van der Waals parameters from other physical (atomic) properties. The interaction strength ϵ_{ij} between two atoms is related to the polarizabilities α_i and α_j , i.e. the ease with which the electron densities can be distorted by an electric field. The Slater–Kirkwood equation³⁹ (2.37) provides an explicit relationship between these quantities, which has been found to give good results for the interaction of rare gas atoms.

$$\epsilon_{ij} = C \frac{\alpha_i \alpha_j}{\sqrt{\frac{\alpha_i}{N_i^{\text{eff}}}} + \sqrt{\frac{\alpha_j}{N_j^{\text{eff}}}}} \quad (2.37)$$

Here C is a constant for converting between the units of ϵ and α , and N_i^{eff} is the effective number of electrons, which may be taken either as the number of valence electrons or treated as a fitting parameter. The R_0 parameter may similarly be taken from atomic quantities. One problem with this procedure is that the atomic polarizability will of course be modified by the bonding situation (i.e. the atom type), which is not taken into account by the Slater–Kirkwood equation.

The above considerations illustrate the inherent contradiction in designing highly accurate force fields. To get a high accuracy for a wide variety of molecules, and a range of properties, many functional complex terms must be included in the force field expression. For each additional parameter introduced in an energy term, the potential number of new parameters to be derived grows with the number of atom types to a power between 1 and 4. The higher accuracy that is needed, the more finely the fundamental units must be separated, i.e. the more atom types must be used. In the extreme limit, each atom that is not symmetry related, in each new molecule is a new atom type. In this limit, each molecule will have its own set of parameters to be used just for this one molecule. To derive these parameters, the molecule must be subjected to many different experiments, or a large number of electronic structure calculations. This is the approach used in “inverting” spectroscopic data to produce a potential energy surface. From a force field point of view, the resulting function is essentially worthless, it just reproduces known results. In order to be useful, a force field should be able to predict unknown properties of molecules from known data on other molecules, i.e. a sophisticated form for inter- or extrapolation. If the force field becomes very complicated, the amount of work required to derive the parameters may be larger than the work required for measuring the property of interest for a given molecule.

The fundamental assumption of force fields is that structural units are transferable between different molecules. A compromise between accuracy and generality must thus be made. In MM2(91) the actual number of parameters compared with the

Table 2.3 Comparison of possible and actual number of MM2(91) parameters

Term	Estimated number of parameters	Actual number of parameters
E_{vdw}	142	142
E_{str}	900	290
E_{bend}	27 000	824
E_{tors}	1 215 000	2466

theoretical estimated possible (based on the 30 effective atom types above) is shown in Table 2.3.

As seen from Table 2.3, there are a large number of possible compounds for which there are no parameters, and on which it is then impossible to perform force field calculations (a good listing of available force field parameters is Osawa and Lipkowitz⁴⁰). Actually, the situation is not as bad as it would appear from Table 2.3. Although only ~0.2% of the possible combinations for the torsional constants has been parameterized, these encompass the majority of the chemically interesting compounds. It has been estimated that ~20% of the ~15 million known compounds can be modelled by the parameters in MM2, the majority with a good accuracy. However, the problem of lacking parameters is very real, and anyone who has used a force field for all but the most rudimentary problems has encountered the problem. How does one progress if there are insufficient parameters for the molecule of interest?

There are two possible routes. The first is to estimate the missing parameters by comparison with force field parameters for similar systems. If, for example, there are missing torsional parameters for rotation around a H—X—Y—O bond in your molecule, but parameters exist for H—X—Y—C, then it is probably a good approximation to use the same values. In other cases, it may be less obvious what to choose. What if your system has an O—X—Y—O torsion, and parameters exist for O—X—Y—C and C—X—Y—O, but they are very different? What do you choose then, one or the other, or the average? After a choice has been made, the results should ideally be evaluated to determine how sensitive they are to the exact value of the guessed parameters. If the guessed parameters can be varied by $\pm 50\%$ without seriously affecting the final results, the property of interest is insensitive to the guessed parameters, and can be trusted to the usual degree of the force field. If, on the other hand, the final results vary by a factor of two when the guessed parameters are changed by 10%, a better estimate of the critical parameters should be sought from external sources. If many parameters are missing from the force field, such an evaluation of the sensitivity to parameter changes becomes impractical, and one should consider either the second route described below, or abandon force field methods altogether.

The second route to missing parameters is to use external information, experimental data or electronic structure calculations. If the missing parameters are bond length and force constant for a specific bond type, it is possible that an experimental bond distance may be obtained from an X-ray structure and the force constant estimated from measured vibrational frequencies, or missing torsional parameters may be obtained from a rotational energy profile calculated by electronic structure calculations. If many parameters are missing, this approach rapidly becomes very time-

consuming, and may not give as good final results as you may have expected from the “rigorous” way of deriving the parameters. The reason for this is discussed below.

Assume now that the functional form of the force field has been settled. The next task is to select a set of reference data – for the sake of argument let us assume that they are derived from experiments, but they could also be taken from electronic structure calculations. The problem is then to assign numerical values to all the parameters such that the results from force field calculations match the reference data set as close as possible. The reference data may be of very different types and accuracy, containing bond distances, bond angles, relative energies, vibrational frequencies, dipole moments, etc. These data of course have different units, and a decision must be made how they should be weighted. How much weight should be put on reproducing a bond length of 1.532 Å relative to an energy difference of 10 kJ/mol? Should the same weight be used for all bond distances, if for example one distance is determined to 1.532 ± 0.001 Å while another is known only to 1.73 ± 0.07 Å? The selection is further complicated by the fact that different experimental methods may give slightly different answers for say the bond distance, even in the limit of no experimental uncertainty. The reason for this is that different experimental methods do not measure the same property. X-ray diffraction, for example, determines the electron distribution, while microwave spectroscopy primarily depends on the nuclear position. The maximum in the electronic distribution may not be exactly identical to the nuclear position, and these two techniques will therefore give slightly different bond lengths.

Once the question of assigning weights for each reference data has been decided, the fitting process can begin. It may be formulated in terms of an error function.⁴¹

$$\text{ErrF}(\text{parameters}) = \sum_i^{\text{data}} \text{weight}_i \cdot (\text{reference value} - \text{calculated value})_i^2 \quad (2.38)$$

The problem is to find the minimum of ErrF with the parameters as variables. From an initial set of guess parameters, force field calculations are performed for the *whole* set of reference molecules and the results compared with the reference data. The deviation is calculated and a new improved set of parameters can be derived. This is continued until a minimum has been found for the ErrF function. To find *the* best set of force field parameters corresponds to finding the *global* minimum for the multi-dimensional ErrF function. The simplest optimization procedure performs a cyclic minimization, reducing the ErrF value by varying one parameter at a time. More advanced methods rely on the ability to calculate the gradient (and possibly also the second derivative) of the ErrF with respect to the parameters. Such information may be used in connection with optimization procedure as described in Chapter 12.

The parameterization process may be done sequentially or in a combined fashion. In the sequential method, a certain class of compounds, such as hydrocarbons, is parameterized first. These parameters are held fixed, and a new class of compounds, for example alcohols and ethers, are then parameterized. This method is in line with the basic assumption of force field, i.e. that parameters are transferable. The advantage is that only a fairly small number of parameters is fitted at a time. The ErrF is therefore a relatively low dimensional function, and one can be reasonably certain that a “good” minimum has been found (although it may not be the global minimum). The

disadvantage is that the final set of parameters necessarily provides a poorer fit (as defined from the value of the ErrF) than if all the parameters are fitted simultaneously.

The combined approach tries to fit all the constants in a single parameterization step. Considering that the number of force field parameters may be many thousands, it is clear that the ErrF function will have a very large number of local minima. To find the global minimum of such a multivariable function is very difficult. It is thus likely that the final set of force field parameters derived by this procedure will in some sense be less than optimal, although it may still be “better” than that derived by the sequential procedure. Furthermore, many of the parameter sets that give low ErrF values (including the global minimum) may be “non-physical”, e.g. force constants for similar bonds being very different. Due to the large dimensionality of the problem, such combined optimizations require the ability to calculate the gradient of the ErrF with respect to the parameters, and writing such programs is not trivial. There is also a more fundamental problem when new classes of compounds are introduced at a later time than the original parameterization. To be consistent, the whole set of parameters should be re-optimized. This has the consequence that (all) parameters change when a new class of compounds is introduced, or whenever more data are included in the reference set. Such “time-dependent” force fields are clearly not desirable. Most parameterization procedures therefore employ a sequential technique, although the number of compound types parameterized in each step varies.

There is one additional point to be mentioned in the parameterization process that is also important for understanding why the addition of missing parameters by comparison with existing data or from external sources is somewhat problematic. This is the question of redundant variables, as can be exemplified by considering acetaldehyde.




Figure 2.17 The structure of acetaldehyde

In the energy bend expression there will be four angle terms describing the geometry around the carbonyl carbon, an HCC, an HCO, a CCO, and an out-of-plane bend. Assuming the latter to be zero for the moment, it is clear that the other three angles are not independent. If the θ_{HCO} and θ_{CCO} angle are given, the θ_{HCC} angle must be $360^\circ - \theta_{\text{HCO}} - \theta_{\text{CCO}}$. Nevertheless, there will be three natural angle parameters, and three force constants associated with these angles. For the whole molecule there are six stretch terms, nine bending terms and six torsional terms (count them!) in addition to at least one out-of-plane term. This means that the force field energy expression has 22 degrees of freedom, in contrast to the 15 ($3N_{\text{atom}} - 6$) independent coordinates necessary to completely specify the system. The force field parameters, as defined by the E_{FF} expression, are therefore not independent.

The implicit assumption in force field parameterization is that, given sufficient amounts of data, this redundancy will cancel out. In the above case, additional data for other aldehydes and ketones may be used (at least partly) for removing this

ambiguity in assigning angle bend parameters, but in general there are more force field parameters than required for describing the system. This clearly illustrates that force field parameters are just that, parameters. They do not necessarily have any direct connection with experimental force constants. Experimental vibrational frequencies can be related to a unique set of force constants, but only in the context of a non-redundant set of coordinates.

It is also clear that errors in the force field due to inadequacies in the functional forms used for each of the energy terms will to some extent be absorbed by the parameter redundancy. Adding new parameters from external sources, or estimating missing parameters by comparison with those for “similar” fragments, may partly destroy this cancellation of errors. This is also the reason why parameters are not transferable between different force fields, the parameter values are dependent on the functional form of the energy terms, and are mutually correlated. The energy profile for rotating around a bond, for example, contains contributions from the electrostatic, the van der Waals and the torsional energy terms. The torsional parameters are therefore intimately related to the atomic partial charges, and cannot be transferred to another force field.

The parameter redundancy is also the reason that care should be exercised when trying to decompose energy differences into individual terms. Although it may be possible to rationalize the preference of one conformation over another by for example increased steric repulsion between certain atom pairs, this is intimately related to the chosen functional form for the non-bonded energy, and the balance between this and the angle bend/torsional terms. The rotational barrier in ethane, for example, may be reproduced solely by an HCCH torsional energy term, solely by an H—H van der Waals repulsion or solely by H—H electrostatic repulsion. Different force fields will have (slightly) different balances of these terms, and while one force field may contribute a conformational difference primarily to steric interactions, another may have the major determining factor to be the torsional energy, and a third may “reveal” that it is all due to electrostatic interactions.

2.3.1 *Parameter reductions in force fields*

The overwhelming problem in developing force fields is the lack of enough high quality reference data. As illustrated above, there are literally millions of possible parameters in even quite simple force fields. The most numerous of these are the torsional parameters, followed by the bending constants. As force fields are designed for predicting properties of unknown molecules, it is inevitable that the problem of lacking parameters will be encountered frequently. Furthermore, many of the existing parameters may be based on very few reference data, and therefore be associated with substantial uncertainty.

Many modern force field programs are commercial. Having the program tell the user that his or her favourite molecule cannot be calculated owing to lack of parameters is not good for business. Making the user derive new parameters, and getting the program to accept them, may require more knowledge than the average user, who is just interested in the answer, has. Many force fields thus have “generic” parameters. This is just a fancy word for the program making more or less educated guesses for the missing parameters.

One way of reducing the number of parameters is to reduce the dependence on atom types. Torsional parameters, for example, can be taken to depend only on the types of the two central atoms. All C—C single bonds would then have the same set of torsional parameters. This does not mean that the rotational barrier for all C—C bonds is identical, since van der Waals and/or electrostatic terms also contribute. Such a reduction replaces all tetra-atomic parameters with diatomic constants, i.e. $V^{ABCD} \rightarrow V^{BC}$. Similarly, the triatomic bending parameters may be reduced to atomic constants by assuming that the bending parameters only depend on the central atom type ($k^{ABC} \rightarrow k^B$, $\theta_0^{ABC} \rightarrow \theta_0^B$). Generic constants are often taken from such reduced parameter sets. In the case of missing torsional parameters, they may also simply be omitted, i.e. setting the constants to zero. A good force field program informs the user of the quality of the parameters used in the calculation, especially if such generic parameters are used, and this is useful for evaluating the quality of the results. Some programs unfortunately use the necessary number of generic parameters to carry out the calculations without notifying the user. In extreme cases, one may perform calculations on molecules for which essentially no “good” parameters exist, and get totally useless results. *The ability to perform a calculation is no guarantee that the results can be trusted!*

The quality of force field parameters is essential for judging how much faith can be put in the results. If the molecule at hand only uses parameters that are based on many good quality experimental results, then the computational results can be trusted to be almost of experimental quality. If, however, the employed parameters are based only on a few experimental data, and/or many generic parameters have been used, the results should be treated with care. Using low quality parameters for describing an “uninteresting” part of the molecule, such as a substituted aromatic ring in a distant side chain, is not problematic. In some cases, such uninteresting parts may simply be substituted by other simpler groups (for example a methyl group). However, if the low quality parameters directly influence the property of interest, the results may potentially be misleading.

2.3.2 Force fields for metal coordination compounds

Coordination chemistry is an area that is especially plagued with the problems of assigning suitable functions for describing the individual energy terms and deriving good parameters.⁴² The bonding around metals is much more varied than for organic molecules, where there are just two, three or four bonds. Furthermore, for a given number of ligands, more than one geometrical arrangement is usually possible. A four-coordinated metal, for example, may either be tetrahedral or square planar, and a five-coordinated metal may either have a square pyramidal or trigonal bipyramidal structure. This is in contrast to four-coordinated atoms such as carbon or sulfur that are always very close to tetrahedral. The increased number of ligands combined with the multitude of possible geometries significantly increases the problem of assigning suitable functional forms for each of the energy terms. Consider for example a “simple” compound such as $\text{Fe}(\text{CO})_5$, which has a trigonal bipyramid structure.

It is immediately clear that a C—Fe—C angle bend must have three energy minima corresponding to 90° , 120° and 180° , indicating that a simple Taylor expansion around a (single) natural value is not suitable. Furthermore, the energy cost for a geometrical distortion (bond stretching and bending) is usually much smaller around a metal atom

than for a carbon atom. This has the consequence that coordination compounds are much more dynamic, displaying phenomena such as pseudo-rotations, ligand exchange and large geometrical variations for changes in the ligands. In iron pentacarbonyl there exists a whole series of equivalent trigonal bipyramid structures that readily interconvert, i.e. the energy cost for changing the C—Fe—C angle from 90° to 120° and to 180° is small. Deviations up to 30° from the “natural” angle by introducing bulky substituents on the ligands are not uncommon. Furthermore, the bond distance for a given metal–ligand is often sensitive to the nature of the other ligands. An example there is the *trans* effect, where a metal–ligand bond distance can vary by perhaps 0.2 Å depending on the nature of the ligand on the opposite side.

Another problem encountered in metal systems is the lack of well-defined bonds. Consider for example an olefin coordinated to a metal – should this be considered as a single bond between the metal and the centre of the C—C bond, or as a metallocyclopropane with two M—C bonds? A cyclopentadiene ligand may similarly be modelled either with a single bond to the centre of the ring, or with five M—C bonds. In reality, these represent limiting behaviours, and the structures on the left in Figure 2.19 correspond to a weak interaction while those on the right involve strong electron donation from the ligand to the metal (and vice versa). A whole range of intermediate cases is found in coordination chemistry. The description with bonds between the metal and all the ligand atoms suffers from the lack of (free) rotation of the ligand. The coordination to the “centre” of the ligand may be modelled by placing a *pseudo-atom* at that position, and relating the ligand atoms to the pseudo-atom (a pseudo-atom is just a point in space, also sometimes called a dummy atom, see Appendix D). Alternatively the coordination may be described entirely by non-bonded interactions (van der Waals and electrostatic).

One possible, although not very elegant, solution to these problems is to assign different atom types for each bonding situation. In the Fe(CO)₅ example, this would mean distinguishing between equatorial and axial CO units. There would then be three different C—Fe—C bending terms, C_{eq}—Fe—C_{eq}, C_{eq}—Fe—C_{ax} and C_{ax}—Fe—C_{ax}, with natural angles of 120°, 90° and 180°, respectively. This approach sacrifices the dynamics of the problem: interchanging an equatorial and axial CO no longer produces energetically equivalent structures. Similarly, the same metal atom in two different geometries (such as tetrahedral and square planar) would be assigned two different types, or in general a new type for each metal in a specific oxidation and spin state, and with a specific number of ligands. This approach encounters the parameter “explosion”, as discussed above. It also biases the results in the direction of the user’s expectations – if a metal atom is assigned a square planar atom type, the structure will end up close to square planar, even though the real geometry may be tetrahedral. The object of a computational study, however, is often a series of compounds that have similar bonding around the metal atom. In such cases, the specific parameterization may be quite useful, but the limitations should of course be kept in mind. Most force field modelling of coordination compounds to date have employed this approach, tailoring an existing method to also reproduce properties (most notably geometries) of a small set of reference systems.

Part of the problems may be solved by using more flexible functional forms for the individual energy terms, most notably the stretching and bending energies. The stretch energy may be chosen as a Morse potential (eq. (2.5)), allowing for quite large distor-




Figure 2.18 The structure of iron pentacarbonyl




Figure 2.19 The ambiguity of modelling metal coordination

tions away from the natural distance, and also being able to account for dissociation. However, phenomena such as the *trans* effect are inherently electronic in nature (similar to the delocalization in conjugated systems) and are not easily accounted for in a force field description.

The multiple minima nature of the bending energy, combined with the low barriers for interconversion, resembles the torsional energy for organic molecules. An expansion of E_{bend} in terms of cosine or sine functions to the angle is therefore more natural than a simple Taylor expansion in the angle. Furthermore, bending around a metal atom often has an energy maximum for an angle of 180° , with a low barrier. The following examples have a zero derivative for a linear angle, and are reasonable for describing bond bending such as that encountered in the H_2O example (Figure 2.5).

$$E_{\text{bend}}(\theta) = k \left(\sin^2 \frac{\theta}{2} - \sin^2 \frac{\theta_0}{2} \right)^2 \quad (2.39)$$

$$E_{\text{bend}}(\theta) = k(1 + \cos(n\theta + \tau)) \quad (2.40)$$

The latter functional form contains a constant n that determines the periodicity of the potential (τ is a phase factor), and allows bending energies with multiple minima, analogously to the torsional energy. It does, however, have problems of unwanted oscillations if an energy minimum with a natural angle close to 180° is desired (this requires




Figure 2.20 Comparing different E_{bend} functionals for the H_2O potential

n to be large, creating many additional minima). It is also unable to describe situations where the minima are not regularly spaced, such as the $\text{Fe}(\text{CO})_5$ system (minima for angles of 90° , 120° and 180°). The performance of eqs (2.39) and (2.40) for the H_2O case is given in Figure 2.20, which can be compared with Figure 2.5.

The barrier towards linearity is given implicitly by the force constant in both the potentials in eqs (2.39) and (2.40). A more general expression, which allows even quite complicated energy functionals to be fitted, is a Fourier expansion.

$$E_{\text{bend}}(\theta) = \sum_n k_n \cos(n\theta) \quad (2.41)$$

An alternative approach consists of neglecting the L—M—L bending terms, and instead includes non-bonded 1,3-interactions. The geometry around the metal is then defined exclusively by the van der Waals and electrostatic contributions (i.e. placing the ligands as far apart as possible), and this model is known as *Points-On-a-Sphere* (POS).⁴³ It is basically equivalent to the *Valence Shell Electron-Pair Repulsion* (VSEPR) model, with VSEPR focusing on the electron pairs that make up a bond, and POS focusing on the atoms and their size.⁴⁴ For alkali, alkaline earth and rare earth metals, where the bonding is mainly electrostatic, POS gives quite reasonable results, but it is unable to model systems where the d-orbitals have a preferred bonding arrangement. Tetracoordinated metal atoms, for example, will in such models always end up being tetrahedral, although d^8 -metals are normally square planar. An explicit coupling between the geometry and the occupancy of the d-orbitals can be achieved by adding a ligand field energy term to the force field energy function,⁴⁵ but this is not (yet) part of the mainstream force field programs.

The final problem encountered in designing force fields for metal complexes is the lack of sufficient numbers of experimental data. Geometrical data for metal compounds are much scarcer than for organic structures, and the soft deformation potentials mean that vibrational frequencies are often difficult to assign to specific modes. Deriving parameters from electron structure calculations is troublesome because the presence of multiple ligands means that the number of atoms is quite large, and the metal atom itself contains many electrons. Furthermore, there are often many different low-lying electronic states owing to partly occupied d-orbitals, indicating that single reference methods (i.e. Hartree–Fock type calculations) are insufficient for even a qualitative correct wave function. Finally, relativistic effects become important for some of the metals in the lower part of the periodic system. These effects have the consequence that electronic structure calculations of a sufficient quality are computationally expensive to carry out.

2.3.3 Universal force fields

The combination of many atom types and the lack of a sufficient number of reference data have prompted the development of force fields with reduced parameters sets, such as the *Universal Force Field* (UFF).⁴⁶ The idea is to derive di-, tri- and tetra-atomic parameters (E_{str} , E_{bend} , E_{tors}) from atomic constants (such as atom radii, ionization potentials, electronegativities, polarizabilities, etc.). Such force fields are in principle capable of describing molecules composed of elements from the whole periodic table, and these have been labelled as “all elements” in Table 2.4 below. They give less accurate results compared with conventional force fields, but geometries are often calculated qualitatively correctly. Relative energies, however, are much more difficult to obtain accurately, and conformational energies for organic molecules are generally quite poor. Another approach is to use simple valence bonding arguments (e.g. hybridization) to derive the functional form for the force field, as employed in the VALBOND approach.⁴⁷

2.4 Differences in Force Fields

There are many different force fields in use. They differ in three main aspects:

- (1) What is the functional form of each energy term?
- (2) How many cross terms are included?
- (3) What types of information are used for fitting the parameters?

There are two general trends. If the force field is designed primarily to treat large systems, such as proteins or DNA, the functional forms are kept as simple as possible. This means that only harmonic functions are used for E_{str} and E_{bend} (or these term are omitted, forcing all bond lengths and angles to be constant), no cross terms are included, and the Lennard-Jones potential is used for E_{vdw} . Such force fields are often called “harmonic”, “diagonal” or “Class I”. The other branch concentrates on reproducing small- to medium-size molecules to a high degree of accuracy. These force fields will include a number of cross terms, use at least cubic or quartic expansions of E_{str} and E_{bend} , and possibly an exponential-type potential for E_{vdw} . The current efforts in developing small-molecule force fields go in the direction of not only striving to

Table 2.4 Comparison of functional forms used in common force fields;⁴⁹ the torsional energy, E_{tors} , is in all cases given as a Fourier series in the torsional angle

Force field	Types	E_{str}	E_{bend}	E_{oop}	E_{vdw}	E_{el}	E_{cross}	Molecules
AMBER	41	P2	P2	imp.	12–6 12–10	charge	none	proteins, nucleic acids, carbohydrates
CFF91/93/95	48	P4	P4	P2	9–6	charge	ss,bb,st, sb,bt,btb	general
CHARMM	29	P2	P2	imp.	12–6	charge	none	proteins
COSMIC	25	P2	P2		Morse	charge	none	general
CVFF	53	P2 or Morse	P2	P2	12–6	charge	ss,bb,sb, btb	general
DREIDING	37	P2 or Morse	P2(cos)	P2(cos)	12–6 or Exp–6	charge	none	general
EAS	2	P2	P3	none	Exp–6	none	none	alkanes
ECEPP		fixed	fixed	fixed	12–6 and 12–10	charge	none	proteins
EFF	2	P4	P3	none	Exp–6	none	ss,bb,sb, st,btb	alkanes
ENCAD	35	P2	P2	imp.	12–6	charge	none	proteins, nucleic acids
ESFF	97	Morse	P2(cos)	P2	9–6	charge	none	all elements
GROMOS		P2	P2	P2(imp.)	12–6	charge	none	proteins, nucleic acids, carbohydrates
MM2	71	P3	P2+P6	P2	Exp–6	dipole	sb	general
MM3	153	P4	P6	P2	Exp–6	dipole or charge	sb,bb,st	general (all elements)
MM4		P6	P6	imp.	Exp–6	charge	ss,bb,sb, tt,st,tb, btb	general
MMFF	99	P4	P3	P2	14–7	charge	sb	general
MOMECC		P2	P2	P2	Exp–6	none	none	metal coordination
NEMO		fixed	fixed	none	Exp–6	quad, polar	none	special
OPLS	41	P2	P2	imp.	12–6	charge	none	proteins, nucleic acids, carbohydrates
PFF		P2	P2	imp.	12–6	polar	none	proteins
PROSA	41	P2	P2	imp.	12–6	polar	none	proteins
QMFF	32	P4	P4	P2	9–6	charge	ss,sb,st,bb, bt,btb	general
SDDF		P4	P4		9–6	charge, dipole, polar	ss,st,tt	hydrocarbons
TraPPE		fixed	P2	fixed	12–6	charge	none	C, N, O compounds
TRIPOS	31	P2	P2	P2	12–6	charge	none	general
UFF	126	P2 or Morse	$\cos(n\theta)$	imp.	12–6	charge	none	all elements
YETI	17	P2	P2	imp.	12–6 and 12–10	charge	none	proteins

Notation: Pn : Polynomial of order n ; $Pn(\cos)$: polynomial of order n in cosine to the angle; $\cos(n\theta)$: Fourier term(s) in cosine to the angle; Exp–6: exponential + R^{-6} ; $n - m$: $R^{-n} + R^{-m}$ Lennard-Jones type potential; quad: electric moments up to quadrupoles; polar: polarizable; fixed: not a variable; imp.: improper torsional angle; ss: stretch–stretch; bb: bend–bend; sb: stretch–bend; st: stretch–torsional; bt: bend–torsional; tt: torsional–torsional; btb: bend–torsional–bend.

reproduce geometries and relative energies, but also vibrational frequencies, and these are often called “Class II” force fields. Further refinements by allowing parameters to depend on neighbouring atom types, e.g. for modelling hyperconjugation, and including electronic polarization effects have been denoted “Class III” force fields.

Force fields designed for treating macromolecules can be simplified by not considering hydrogens explicitly – the so-called *united atom* approach (an option present in for example the AMBER, CHARMM, GROMOS and DREIDING force fields). Instead of modelling a CH₂ group as a carbon and two hydrogens, a single “CH₂ atom” may be assigned, and such a united atom will have a larger van der Waals radius to account also for the hydrogens. The advantage of united atoms is that they effectively reduce the number of variables by a factor of ~2–3, thereby allowing correspondingly larger systems to be treated. Of course the coarser the atomic description is, the less detailed the final results will be. Which description, and thus which type of force field to use, depends on what type of information is sought. If the interest is in geometries and relative energies of different conformations of say hexose, then an elaborate force field is necessary. However, if the interest is in studying the dynamics of a protein consisting of hundreds of amino acids, a crude model where whole amino acids are used as the fundamental unit may be all that is possible, considering the sheer size of the problem.⁴⁸

Table 2.4 gives a description of the functional forms used in some of the common force fields. The torsional energy is in all cases written as a Fourier series, typically of order 3. Many of the force fields undergo developments, and the number of atom types increases as more and more systems become parameterized, and Table 2.4 may thus be considered as a “snapshot” of the situation when the data were collected. The “universal” type force fields, described in Section 2.3.3, are in principle capable of covering molecules composed of elements from the whole periodic table, and these have been labelled as “all elements”.

Even for force fields employing the same mathematical form for an energy term there may be significant differences in the parameters Table 2.5 below shows the variability of the parameters for the stretch energy between different force fields. It should be noted that the stretching parameters are among those that vary the *least* between force fields.

It is perhaps surprising that force constants may differ by almost a factor of two, but this is of course related to the stiffness of the stretch and bending energies. Very few molecules have bond lengths deviating more than a few hundredths of an angstrom

Table 2.5 Comparison of stretch energy parameters for different force fields

Force field	R_0 (Å)				k (mdyn/Å)			
	C—C	C—O	C—F	C=O	C—C	C—O	C—F	C=O
MM2	1.523	1.402	1.392	1.208	4.40	5.36	5.10	10.80
MM3	1.525	1.413	1.380	1.208	4.49	5.70	5.10	10.10
MMFF	1.508	1.418	1.360	1.222	4.26	5.05	6.01	12.95
AMBER	1.526	1.410	1.380	1.220	4.31	4.45	3.48	8.76
OPLS	1.529	1.410	1.332	1.229	3.73	4.45	5.10	7.92

from the reference value, and the associated energy contribution will be small regardless of the force constant value. Stated another way, the minimum energy geometry is insensitive to the exact value of the force constant.

2.5 Computational Considerations

Evaluation of the non-bonded energy is by far the most time-consuming step, and this can be exemplified by the number of individual energy terms for the linear alkanes $\text{CH}_3(\text{CH}_2)_{n-2}\text{CH}_3$ shown in Table 2.6.

Table 2.6 Number of terms for each energy contribution in $\text{CH}_3(\text{CH}_2)_{n-2}\text{CH}_3$

n	N_{atoms}	E_{str}	E_{bend}	E_{tors}	E_{vdw}
10	32	31 (5%)	30 (10%)	81 (14%)	405 (70%)
20	62	61 (3%)	60 (6%)	171 (8%)	1710 (83%)
50	152	151 (1%)	300 (3%)	441 (4%)	11 025 (93%)
100	302	301 (1%)	600 (1%)	891 (2%)	44 550 (96%)
	N	$(N - 1)$	$2(N - 2)$	$3(N - 5)$	$\frac{1}{2}N(N - 1) - 3N + 5$

The number of bonded contributions, E_{str} , E_{bend} and E_{tors} , grow *linearly* with the system size, while the non-bonded contributions, E_{vdw} (and E_{el}), grow as the *square* of the system size. This is fairly obvious as, for a large molecule, most of the atom pairs are not bonded, or not bonded to a common atom, and thus contribute with an E_{vdw} term. For $\text{CH}_3(\text{CH}_2)_{98}\text{CH}_3$, which contains a mere 302 atoms, the non-bonded terms already account for ~96% of the computational effort. For a 1000 atom system, the percentage is 98.8%, and for 10000 atoms it is 99.88%. *In the limit of large molecules, the computational time for calculating the force field energy grows approximately as the square of the number of atoms.* The majority of these non-bonded energy contributions are numerically very small, as the distance between the atom pairs is large. A considerable saving in computational time can be achieved by truncating the van der Waals potential at some distance, say 10 Å. If the distance is larger than this cutoff, the contribution is neglected. This is not quite as clean as it may sound at first. Although it is true that the contribution from a pair of atoms is very small if they are separated by 10 Å, there may be a large number of such atom pairs. The *individual* contribution falls off quickly, but the *number* of contributions also increases. Many force fields use cutoff distances around 10 Å, but it has been shown that the total van der Waals energy only converges if the cutoff distance is of the order of 20 Å. However, using a cutoff of 20 Å may significantly increase the computational time (by a factor of perhaps 5–10) relative to a cutoff of 10 Å.

The introduction of a cutoff distance does not by itself lead to a significant computational saving, since all the distances must be computed prior to the decision on whether to include the contribution or not. A substantial increase in computational efficiency can be obtained by keeping a *non-bonded* or *neighbour list* over atom pairs. From a given starting geometry, a list is prepared over the atom pairs that are within the cutoff distance plus a smaller buffer zone. During a minimization or simulation, only the contributions from the atom pairs on the list are evaluated, which avoids the

calculation of distances between all pairs of atoms. Since the geometry changes during the minimization or simulation, the non-bonded list must be updated once an atom has moved more than the buffer zone or simply at (fixed) suitable intervals, for example every 10 or 20 steps.

The use of a cutoff distance reduces the scaling in the large system limit from N_{atom}^2 to N_{atom} since the non-bonded contributions are then only evaluated within the local “sphere” determined by the cutoff radius. A cutoff distance of $\sim 10 \text{ \AA}$, however, is so large that the large system limit is not achieved in practical calculations. Furthermore, the updating of the neighbour list still involves calculating all distances between atom pairs. The actual scaling is thus N_{atom}^n , where n is between 1 and 2, depending on the details of the system. In most applications, however, it is not the energy of a single geometry that is of interest, but that of an optimized geometry. The larger the molecule, the more degrees of freedom, and the more complicated the geometry optimization is. The gain by introducing a non-bonded cutoff is partly offset by the increase in computational effort in the geometry optimization. Thus as a rough guideline the increase in computational time upon changing the size of the molecule can be taken as N_{atom}^2 .

The introduction of a cutoff distance, beyond which E_{vdw} is set to zero, is quite reasonable as the neglected contributions rapidly become small for any reasonable cutoff distance. This is not true for the other part of the non-bonded energy, the Coulomb interaction. Contrary to the van der Waals energy, which falls off as R^{-6} , the charge–charge interaction varies as R^{-1} . This is actually only true for the interaction between molecules (or fragments) carrying a net charge. The charge distribution in neutral molecules or fragments makes the long range interaction behave as a dipole–dipole interaction. Consider for example the interaction between two carbonyl groups. The carbons carry a positive and the oxygens a negative charge. Seen from a distance it looks like a bond dipole moment, not two net charges. The interaction between two dipoles behaves like R^{-3} , not R^{-1} , but an R^{-3} interaction still requires a significantly larger cutoff than the van der Waals R^{-6} interaction.

Table 2.7 shows the interaction energy between two carbonyl groups in terms of the MM3 E_{vdw} and E_{el} , the latter described either by an atomic point charge or a bond dipole model. The bond dipole moment is 1.86 debye, corresponding to atomic charges of ± 0.32 separated by a bond length of 1.208 \AA . For comparison, the interaction between two net charges of 0.32 is also given.

From Table 2.7 it is clearly seen that E_{vdw} becomes small (less than $\sim 0.01 \text{ kJ/mol}$) beyond a distance of $\sim 10 \text{ \AA}$. The electrostatic interaction reaches the same level of

Table 2.7 Comparing the distance behaviour of non-bonded energy contributions (kJ/mol)

Distance (\AA)	E_{vdw}	$E_{\text{dipole-dipole}}$	$E_{\text{point charges}}$	$E_{\text{net charges}}$
5	-0.92	1.665	1.598	28.5
10	-0.0060	0.208	0.206	14.2
15	-0.00054	0.0617	0.0614	9.5
20	-9.5×10^{-5}	0.0260	0.0259	7.1
30	-8.4×10^{-6}	0.00770	0.00770	4.7
50	-3.9×10^{-7}	0.00167	0.00167	2.8
100	-6.1×10^{-9}	0.000208	0.000208	1.4

importance at a distance of ~ 30 Å. The table also shows that the interaction between point charges behaves as a dipole–dipole interaction, i.e. an R^{-3} dependence. The interaction between net charges is very long-range – even at 100 Å separation there is a 1.4 kJ/mol energy contribution. The “cutoff” distance corresponding to a contribution of ~ 0.01 kJ/mol is of the order of 14 000 Å!

There are different ways of implementing a non-bonded cutoff. The simplest is to neglect all contributions if the distance is larger than the cutoff. This is in general not a very good method as the energy function becomes discontinuous. Derivatives of the energy function also become discontinuous, which causes problems in optimization procedures and when performing simulations. A better method is to use two cutoff distances between which a switching function connects the correct E_{vdw} or E_{el} smoothly with zero. Such interpolations solve the mathematical problems associated with optimization and simulation, but the chemical significance of the cutoff of course still remains. This is especially troublesome in simulation studies where the distribution of solvent molecules can be very dependent on the use of cutoffs. The modern approaches for evaluating the electrostatic contribution is the use of *fast multipole* or *Ewald* sum methods (see Section 14.3), both of which are able to calculate the electrostatic energy exactly (to within a specified numerical precision) with an effort that scales less than quadratic with the number of particles (linear for fast multipole, $N^{3/2}$ for Ewald and $N \ln N$ for particle mesh Ewald methods). These methods require only slightly more computer time than using a cutoff-based method, and give much better results.

2.6 Validation of Force Fields

The quality of a force field calculation depends on two quantities: the appropriateness of the mathematical form of the energy expression, and the accuracy of the parameters. If elaborate forms for the individual interaction terms have been chosen, and a large number of experimental data is available for assigning the parameters, the results of a calculation may be as good as those obtained from experiments, but at a fraction of the cost. This is the case for simple systems such as hydrocarbons. Even a force field with complicated functional forms for each of the energy contributions contains only relatively few parameters when carbon and hydrogen are the only atom types, and experimental data exist for hundreds of such compounds. The parameters can therefore be assigned with a high degree of confidence. Other well-known compound types, such as ethers and alcohol, can achieve almost as good results. For less common species, such as sulfones, or polyfunctional molecules, much less experimental information is available, and the parameters are less well defined.

Force field methods are primarily geared to predicting two properties: geometries and relative energies. Structural features are in general much easier to predict than relative energies. Each geometrical feature depends only on a few parameters. For example bond distances are essentially determined by R_0 and the corresponding force constant, bond angles by θ_0 , and conformational minima by V_1 , V_2 and V_3 . It is therefore relatively easy to assign parameters that reproduce a given geometry. Relative energies of different conformations, however, are much more troublesome since they are a consequence of many small contributions, i.e. the exact functional form of the individual energy terms and the balance between them. The largest contributions to conformational energy differences are the non-bonded and torsional terms, and it is

therefore important to have good representations of the whole torsional energy profile. Even though a given force field may be parameterized to reproduce rotational energy profiles for ethane and ethanol, and contains a good description of hydrogen bonding between two ethanol molecules, there is no guarantee that it will be successful in reproducing the relative energies of different conformations of say 1,2-dihydroxyethane.⁵⁰ For large systems, it is inevitable that small inaccuracies in the functional forms for the energy terms and parameters will influence the shape of the whole energy surface to the point where minima may disappear or become saddle points. Essentially all force fields, no matter how elaborate the functional forms and parameterization, will have artificial minima, and fail to predict real minima, even for quite small systems. For cyclododecane (which is one of the largest molecules to have been subjected to an exhaustive search), the MM2 force field predicts 122 different conformations, but the MM3 surface contains only 98 minima.⁵¹ Given that cyclododecane belongs to a class of well-parameterized molecules, the saturated hydrocarbons, and that MM2 and MM3 are among the most accurate force fields, this clearly illustrates the point.

Validation of a force field is typically done by showing how accurately it reproduces reference data, which may or may not have been used in the actual parameterization. Since different force fields employ different sets of reference data, it is difficult to compare their accuracies directly. Indeed there is no single “best” force field, each has its advantages and disadvantages. They perform best for the type of compounds that have been used in the parameterization, but may give questionable results for other systems. Table 2.8 gives typical accuracies for ΔH_f that can be obtained with the MM2 force field.

Table 2.8 Average errors in heat of formation (kJ/mol) by MM2⁵²

Compound type	Average error in ΔH_f
Hydrocarbons	1.8
Ethers and alcohols	2.1
Carbonyl compounds	3.4
Aliphatic amines	1.9
Aromatic amines	12.1
Silanes	4.5

The average error is the difference between the calculated and experimental ΔH_f . In this connection it should be noted that the average error in the *experimental* data for the hydrocarbons is 1.7 kJ/mol, i.e. MM2 essentially reproduce the experiments to within the experimental uncertainty.

There is one final thing that needs to be mentioned in connection with the validation of a force field, namely the reproducibility. The results of a calculation are determined completely by the mathematical expressions for the energy terms and the parameter set (assuming that the computer program is working correctly). A new force field is usually parameterized for a fairly small set of functional groups initially, and may then evolve by addition of parameters for a larger diversity later. This sometimes has the consequence that some of the initial parameters must be modified to give an acceptable fit. Furthermore, new experimental data may warrant changes in existing parameters. In some cases, different sets of parameters are derived by different

research groups for the same types of functional group. The result is that the parameter set for a given force field is not constant in time, and sometimes not in geographical location either. There may also be differences in the implementation details of the energy terms. The E_{oop} in MM2, for example, is defined as a harmonic term in the bending angle (Figure 2.6), but may be substituted by an improper torsional angle in some computer programs. The consequence is that there often are several different “flavours” of a given force field, depending on the exact implementation, the original parameter set (which may not be the most recent), and any local additions to the parameters. A vivid example is the MM2 force field, which exists in several different implementations that do not give the exact same results but are nevertheless denoted as “MM2” results.

2.7 Practical Considerations

It should be clear that force field methods are models of the real quantum mechanical systems. The neglect of electrons as individual particles forces the user to define explicitly the bonding in the molecule prior to any calculations. The *user* must decide how to describe a given molecule in terms of the selected force field. The input to a calculation consists of three sets of information:

- (1) Which atom types are present?
- (2) How are they connected, i.e. which atoms are bonded to each other?
- (3) A start guess of the geometry.

The first two sets of information determine the functional form of E_{FF} , i.e. enable the calculation of the potential energy surface for the molecule. Normally the molecule will then be optimized by minimizing E_{FF} , which requires a starting guess of the geometry. The information necessary for the program to perform the calculation is read via a file on the computer, and in older programs, the input file had to be prepared manually by the user. All the above three sets of information, however, can be uniquely defined from a (three-dimensional) drawing of a molecule. Modern programs therefore have a graphical interface that allows the molecule simply to be drawn on the screen or constructed from pre-optimized fragments. The interface then automatically assigns suitable atom types based on the selected atomic symbols and the connectivity, and converts the drawing to Cartesian coordinates.

2.8 Advantages and Limitations of Force Field Methods

The main advantage of force field methods is the speed with which calculations can be performed, enabling large systems to be treated. Even with a desktop personal computer, molecules with several thousand atoms can be optimized. This puts the applications in the region of modelling biomolecular macromolecules, such as proteins and DNA, and molecular modelling is now used by most pharmaceutical companies. The ability to treat a large number of particles also makes force field models the only realistic method for performing simulations where solvent effects or crystal packing can be studied (Chapter 14).

For systems where good parameters are available, it is possible to make very good predictions of geometries and relative energies of a large number of molecules in a

short time. It is also possible to determine barriers for interconversion between different conformations, although this is much less automated. One of the main problems is of course the lack of good parameters. If the molecule is slightly out of the ordinary, it is very likely that only poor quality parameters exist, or none at all. Obtaining suitable values for these missing parameters can be a frustrating experience. Force field methods are good for predicting properties for classes of molecules where a lot of information already exists. For unusual molecules, their use is very limited.

Finally, force field methods are “zero-dimensional”. It is not possible to assess the probable error of a given result *within* the method. The quality of the result can only be judged by comparing to other calculations on similar types of molecules for which relevant experimental data exist.

2.9 Transition Structure Modelling

Structural changes can be divided into two general types: those of a conformational nature and those involving bond breaking/forming. There are intermediate cases, such as bonds involving metal coordination, but since metal coordination is difficult to model anyway, we will neglect such systems at present. The bottleneck for structural changes is the highest energy point along the reaction path, called the *Transition State* or *Transition Structure* (TS) (Chapter 13). Conformational TS's have the same atom types and bonding for both the reactant and product, and can be located on the force field energy surface by standard optimization algorithms. Since conformational changes are often localized to rotation around a single bond, simply locating the maximum energy structure for rotation (so-called “torsional angle driving”, see Section 12.4.1) around this bond often represents a good approximation of the real TS.

Modelling TS's for reactions involving bond breaking/forming within a force field methodology is much more difficult.⁵³ In this case, the reactant and product are *not* described by the same set of atom types and/or bonding. There may even be a different number of atoms at each end of the reaction (for example lone pairs disappearing). This means that there are two *different* force field energy functions for the reactant and product, i.e. the energy as a function of the reactant coordinate is not continuous. Nevertheless, methods have been developed for modelling *differences* in activation energies between similar reactions by means of force field techniques, and three approaches are described below.

2.9.1 Modelling the TS as a minimum energy structure

One of the early applications of TS modelling was the work on steric effects in S_N2 reactions by DeTar and coworkers, and it has more recently been advocated by Houk and coworkers.⁵⁴ The approach consists of first locating the TS for a typical example of the reaction with electronic structure methods, often at the Hartree–Fock or density functional theory level. The force field function is then modified such that an energy minimum is created with a geometry that matches the TS geometry found by the electronic structure method. The modification defines new parameters for all the energy terms involving the partly formed/broken bonds. The stretch energy terms have natural bond lengths taken from the electronic structure calculation, and force constants that are typically half the strength of normal bonds. Bond angle terms are similarly

modified with respect to equilibrium values and force constants, the former taken from the electronic structure data and the latter usually estimated. These modifications often necessitate the definition of new “transition state” atom types. Once the force field parameters have been defined, the structure is minimized as usual. Sometimes a few cycles of parameter adjustments and re-optimizations are necessary for obtaining a set of parameters capable of reproducing the desired TS geometry. P.-O. Norrby has described a partly automated method for simultaneously optimizing all the parameters to reproduce the reference structure.⁵⁵ When the modified force field is capable of reproducing the reference TS geometry, it can be used for predicting TS geometries and relative energies of reactions related to the model system. As long as the differences between the systems are purely “steric”, it can be hoped that relative energy differences (between the reactant and the TS model) will correlate with relative activation energies. Purely electronic effects, such as Hammett-type effects due to *para*-substitution in aromatic systems, can of course not be modelled by force field techniques.

2.9.2 Modelling the TS as a minimum energy structure on the reactant/product energy seam

There are two principal problems with the above modelling technique. First, the TS is modelled as a minimum on the energy surface, while it should be a first-order saddle point. This has the consequence that changes in the TS position along the reaction coordinate due to differences in the reaction energy will be in the wrong direction (Section 15.6). In many cases, this is probably not important. For reactions having a reasonable barrier, the TS geometry appears to be relatively constant, which may be rationalized in terms of the Marcus equation (Section 15.5). Comparing reactions that differ in terms of the steric hindrance at the TS, however, may be problematic, as the TS changes along the reaction coordinate will be in the wrong direction. The second problem is the more or less *ad hoc* assignment of parameters. Even for quite simple reactions, many new parameters must be added. Inventing perhaps 40 new parameters for reproducing maybe five relative activation energies raises the nagging question as to whether TS modelling is just a fancy way of describing five data points by 40 variables.⁵⁶

Both of these problems are eliminated in the intersecting potential energy surface modelling technique called *SEAM*.⁵⁷ The force field TS is here modelled as the lowest point on the seam of the reactant and product energy functions, as shown in Figure 2.21. Locating the minimum energy structure on the seam is an example of a constrained optimization; the energy should be minimized subject to the constraint that the reactant and product energies are identical. Although this is computationally somewhat more complicated than the simple minimization required in the Houk approach, it can be handled in a quite efficient manner.

In the SEAM approach only the force field parameters for describing the reactant and products are necessary, alleviating the problem of assigning parameters specific for the TS. Furthermore, differences in reactivity due to differences in reaction energy or steric hindrance at the TS are automatically included. The question is how accurately the lowest energy point on the seam resembles the actual TS. This is difficult to evaluate rigorously as it is intimately connected with the accuracy of the force field used for describing the reactant and product structures. It is clear that the TS will have




Figure 2.21 Modelling a transition structure as a minimum on the intersection of two potential energy surfaces

bond distances and angles significantly different from equilibrium structures. This method of TS modelling therefore requires a force field that is accurate over a much wider range of geometries than normal. Especially important is the stretch energy, which must be able to describe bond breaking. A polynomial expansion is therefore not suitable, and for example a Morse function is necessary. Similarly, the repulsive part of the van der Waals energy must be fairly accurate, which means that Lennard-Jones potentials are not suitable and should be replaced by, for example, Buckingham-type potentials. Furthermore, many of the commonly employed cross terms (Section 2.2.8) become unstable at long bonds lengths and must be modified. When such modifications are incorporated, however, the intersecting energy surface model appears to give surprisingly good results.

There are of course also disadvantages in this approach: these are essentially the same as the advantages! The SEAM method automatically includes the effect of different reaction energies, since a more exothermic reaction will move the TS toward the reactant and lower the activation energy (Section 15.5). This, however, requires that the force field be able to calculate relative energies of the reactant and product, i.e. the ability to convert steric energies to heat of formation. As mentioned in Section 2.2.10, there are only a few force fields that have been parameterized for this. In practice, this is not a major problem since the reaction energy for a prototypical example of the reaction of interest can be obtained from experimental data or estimated. Using the normal force field assumption of transferability of heat of formation parameters, the difference in reaction energy is thus equal to the difference in steric energy. Only the reaction energy for a single reaction of the given type therefore needs to be estimated and relative activation energies are not sensitive to the exact value used.

If the minimum energy seam structure does not accurately represent the actual TS (compared for example with that obtained from an electronic structure calculation) the lack of specific TS parameters becomes a disadvantage. In the Houk approach, it is fairly easy to adjust the relevant TS parameters to reproduce the desired TS geometry. In the intersecting energy surface method, the TS geometry is a complicated result

of the force field parameters for the reactant and product, and the force field energy functions. Modifying the force field parameters, or the functional form of some of the energy terms, in order to achieve the desired TS geometry without destroying the description of the reactant/product, is far from trivial. A final disadvantage, which is inherent to the SEAM method, is the implicit assumption that all the geometrical changes between the reactant and product occurs in a “synchronous” fashion, albeit weighted by the energy costs for each type of distortion. “Asynchronous” or “two-stage” reactions (as opposed to two-step reactions that involve an intermediate), where some geometrical changes occur mainly before the TS, and others mainly after the TS, are difficult to model by this method.

Since the TS is given in terms of the diabatic energy surfaces for the reactant and product, it is also clear that activation energies will be too high. For evaluating relative activation energies of similar reactions this is not a major problem since the important aspect is the relative energies. The overestimation of the activation energy can be improved by adding a “resonance” term to the force field, as discussed in the next section.

2.9.3 Modelling the reactive energy surface by interacting force field functions or by geometry-dependent parameters

Within a valence bond approach (Chapter 7), the reaction energy surface can be considered as arising from the interaction of two *diabatic* surfaces. The *adiabatic* surface can be generated by solving a 2×2 secular equation involving the reactant and product energy surfaces, E_r and E_p .

$$\begin{vmatrix} E_r - E & V \\ V & E_p - E \end{vmatrix} = 0 \quad (2.42)$$

$$E = \frac{1}{2} \left[(E_r + E_p) - \sqrt{(E_r + E_p)^2 + 4V^2} \right]$$

A. Warshel has pioneered the *Extended Valence Bond* (EVB) method,⁵⁸ where the reactant and product surfaces are described by force field energy functions, and Truhlar has more recently generalized the approach by the *Multi-Configurations Molecular Mechanics* (MCMM) method.⁵⁹ In either case, the introduction of the interaction term V generates a continuous energy surface for transforming the reactant into the product configuration, and the TS can be located analogously to energy surfaces generated by electronic structure methods. The main drawback of this method is the somewhat arbitrary interaction element, and the fact that the TS must be located as a first-order saddle point, which is significantly more difficult than locating minima or minima on seams. It can be noted that the SEAM method corresponds to the limiting case where $V \rightarrow 0$ in the EVB method.

Another way of creating a continuous surface connecting the reactant and product energy functions is to make the force field parameters dependent on the geometry, which is an approach used in the *ReaxFF* method.⁶⁰ The force constant for stretching a bond, for example, should decrease and approach zero as the bond length increases towards infinity. The energy function in this case depends directly on the atomic coordinates via the energy term in eq. (2.3), but also indirectly via the geometry

dependence of the parameters. Achieving a smooth and realistic variation of the energy with geometry requires quite elaborate interpolation functions, which makes the parameterization non-trivial.

2.10 Hybrid Force Field Electronic Structure Methods

Force field methods are inherently unable to describe the details of bond breaking/forming or electron transfer reactions, since there is an extensive rearrangement of the electrons. If the system of interest is too large to treat entirely by electronic structure methods, there are two possible approximate methods that can be used. In some cases, the system can be “pruned” to a size that can be treated by replacing “unimportant” parts of the molecule with smaller model groups, e.g. substitution of a hydrogen or methyl group for a phenyl ring. For studying enzymes, however, it is usually assumed that the whole system is important for holding the active site in the proper arrangement, and the “backbone” conformation may change during the reaction. Similarly, for studying solvation, it is not possible to “prune” the number of solvent molecules without severely affecting the accuracy of the model. Hybrid methods have been designed for modelling such cases, where the active site is calculated by electronic structure methods (usually semi-empirical, low-level *ab initio* or density functional methods), while the backbone is calculated by a force field method.⁶¹ Such methods are often denoted *Quantum Mechanics – Molecular Mechanics (QM/MM)*. Formally, the partition can be done by dividing the Hamiltonian and resulting energy into three parts.

$$\begin{aligned}\mathbf{H}_{\text{total}} &= \mathbf{H}_{\text{QM}} + \mathbf{H}_{\text{MM}} + \mathbf{H}_{\text{QM/MM}} \\ E_{\text{total}} &= E_{\text{QM}} + E_{\text{MM}} + E_{\text{QM/MM}}\end{aligned}\quad (2.43)$$

The QM and MM regions are described completely analogously to the corresponding isolated system, using the techniques discussed in Chapters 2–6. The main problem with QM/MM schemes is deciding how the two parts should interact (i.e. $\mathbf{H}_{\text{QM/MM}}$). The easiest situation is when the two regions are not connected by covalent bonds, as for example when using an MM description for modelling the effect of solvation on a QM system. If the two regions are connected by covalent bonding, as for example when using a QM model for the active site in an enzyme and describing the backbone by an MM model, the partitioning is somewhat more difficult.

The lowest level of interaction is called *mechanical embedding*. In this case, only the bonded and steric energies of the two regions are included in the interaction term, i.e. QM atoms have additional forces generated by the MM framework, and vice versa, but there is no interaction between the electronic parts of the two regions. The QM atoms are assigned van der Waals parameters and included in an MM non-bonded energy expression, as illustrated by a Lennard-Jones potential in eq. (2.44).

$$\mathbf{H}_{\text{QM/MM}} = \sum_a^{N_{\text{MM-Atoms}}} \sum_b^{N_{\text{QM-Atoms}}} \epsilon_{ab} \left[\left(\frac{R_0}{R_{ab}} \right)^{12} - 2 \left(\frac{R_0}{R_{ab}} \right)^6 \right] \quad (2.44)$$

The QM atoms may also be assigned partial charges, for example from a population analysis, and charge–charge interactions between the QM and MM atoms included by a classical expression such as eq. (2.20). If the two regions are bonded there are

additional terms corresponding to stretching and bending interactions. The mechanical embedding model is rarely a useful level of approximation, as the wave function of the QM region does not respond to changes in the MM region.

The next level of improvement is called *electronic embedding*, where the atoms in the MM regions are allowed to polarize the QM region. Partial charges on the MM atoms can be incorporated into the QM Hamiltonian analogously to nuclear charges (i.e. adding \mathbf{V}_{nc} -like terms to the one-electron matrix elements in eq. (3.56)), and the QM atoms thus feel the electric potential due to all the MM atoms.

$$\mathbf{V}_{\text{QM/MM}} = \sum_a^{N_{\text{MM-Atoms}}} \frac{Q_a}{|\mathbf{R}_a - \mathbf{r}_i|} \quad (2.45)$$

The non-bonded mechanical term in eq. (2.44) is still needed in order to prevent the MM atoms from drifting into the QM region. The electronic embedding allows the geometry of MM atoms to influence the QM region, i.e. the wave function in the QM region becomes coupled to the MM geometry. An interesting computational issue arises when the number of MM atoms is large and the QM region is small, since the calculation of the one-electron integrals associated with $\mathbf{V}_{\text{QM/MM}}$ may become a dominating factor, rather than the two-electron integrals associated with the QM region itself, but in most cases the inclusion of the $\mathbf{V}_{\text{QM/MM}}$ term only marginally increases the computational effort over a mechanical embedding.

A further refinement, often called *polarizable embedding*, can be made by allowing the QM atoms also to polarize the MM region, i.e. the electric field generated by the QM region influences the MM electric moments (atomic charges and dipoles). This of course requires that a polarizable force field is employed (Section 2.2.7), and necessitates a double iterative procedure for allowing the electric fields in both the QM and MM regions to be determined in a self-consistent fashion. This substantially increases the computational cost, and since polarizable force fields are not yet commonly used anyway, most QM/MM methods employ the electronic embedding approximation. An exception is the *effective fragment method*, often used for modelling solvation, where both quadrupoles and polarizabilities are included for the MM atoms.⁶²

In many cases, the QM and MM regions belong to the same molecule, and the division between the two parts must be done by cutting one or more covalent bonds. This leaves one or more unpaired electrons in the QM part, which must be properly terminated. In most cases, the dangling bonds are terminated by adding “link” atoms, typically a hydrogen. For semi-empirical methods, it can also be a pseudo-halogen atom with parameters adjusted to provide a special link atom.⁶³ Alternatively, the termination can be in the form of a localized molecular or generalized hybrid orbital.⁶⁴ At present, there does not seem to be a clear consensus on whether one or the other approach provides the best results, but the link atom method is somewhat simpler to implement. When the link atom procedure is used, the link atom(s) is only present in the QM calculation, and is not seen by the MM framework. A number of choices must also be made for which and how many of the MM bend and torsional terms that involve one or more QM atoms are included. Bending terms involving two MM and one QM atoms are usually included, but those involving one MM and two QM atoms may be neglected. Similarly, the torsional terms involving only one QM atom are usually included, but those involving two or three QM atoms may or may not be neglected.

The concept of mixing methods of different accuracy has been generalized in the *ONIOM* (*our own n-layered integrated molecular orbital molecular mechanics*) method to include several (usually two or three) layers, for example using relatively high-level theory in the central part, a lower level electronic structure theory in an intermediate layer and force field to treat the outer layer.⁶⁵ The original ONIOM method only employed mechanical embedding for the QM/MM interface, but more recent extensions have also included electronic embedding.⁶⁶ The ONIOM method employs an extrapolation scheme based on assumed additivity, in analogy to the CBS, *Gn* and *Wn* methods discussed in Section 5.7. For a two-layer scheme, the small (model) system is calculated at both the low and high levels of theory, while the large (real) system is calculated at the low level of theory. The result for the real system at the high theoretical level is estimated by adding the change between the high and low levels of theory for the model system to the low level results for the real system, as illustrated in Figure 2.22 and eq. (2.46).

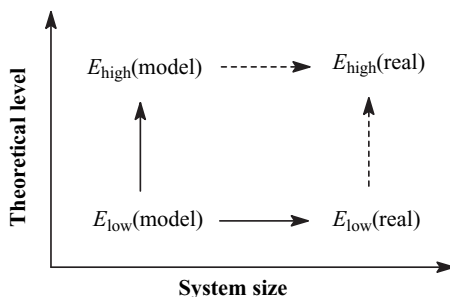


Figure 2.22 Illustration of the ONIOM extrapolation method

$$E_{\text{ONIOM}}(\text{real system, high level}) = E_{\text{high level}}(\text{model system}) - E_{\text{low level}}(\text{model system}) + E_{\text{low level}}(\text{real system}) \quad (2.46)$$

A similar extrapolation can be done for multi-level ONIOM models, although it requires several intermediate calculations. It should be noted that derivatives of the ONIOM model can be constructed straightforwardly from the corresponding derivative of the underlying methods, and it is thus possible to perform geometry optimizations and vibrational analysis using the ONIOM energy function.

QM/MM methods are often used for modelling solvent effects, with the solvent treated by MM methods, but in some cases the first solvation shell is included in the QM region. If such methods are used in connection with dynamical sampling of the configurational space, it is possible that MM solvent molecules can enter the QM regions, or QM solvent molecules can drift into the MM region. In order to handle such situations, there must be a procedure for allowing solvent molecules to switch between a QM and MM description. In order to ensure a smooth transition, a transition region can be defined between the two parts, where a switching function is employed to make a continuous transition between the two descriptions.⁶⁷

The main problem with QM/MM methods is that there is no unique way of deciding which part should be treated by force field and which by quantum mechanics, and QM/MM methods are therefore not “black box” methods. The “stitching” together of

the two regions is certainly not unique, and the many possible combinations of force field and QM methods make QM/MM methods still somewhat experimental. Furthermore, the inability to perform calibration studies of large systems by pure QM methods makes it difficult to evaluate the severity of the approximations included in QM/MM methods.

References

1. U. Dinur, A. T. Hagler, *Rev. Comp. Chem.*, **2** (1991), 99; U. Burkert, N. L. Allinger, *Molecular Mechanics*, ACS, 1982; A. K. Rappe, C. J. Casewit, *Molecular Mechanics Across Chemistry*, University Science Books, 1997; A. D. Mackrell Jr, *J. Comp. Chem.*, **25** (2004), 1584; J. W. Ponder, D. A. Case, *Adv. Protein Chem.*, **66** (2003), 27.
2. N. L. Allinger, *J. Am. Chem. Soc.*, **99** (1977), 8127; N. L. Allinger, Y. H. Yuh, J. H. Lii, *J. Am. Chem. Soc.*, **111** (1989), 8551.
3. S. W. Benson, *Thermochemical Kinetics*, John Wiley and Sons, 1976.
4. P. M. Morse, *Phys. Rev.*, **34** (1929), 57.
5. M. Möllhoff, U. Sternberg, *J. Mol. Mod.*, **7** (2001), 90.
6. P. Jensen, *J. Mol. Spect.*, **133** (1989), 438.
7. B. Albinsson, H. Teramae, J. W. Downing, J. Michl, *Chem. Eur. J.*, **2** (1996), 529.
8. P. C. Chen, *Int. J. Quant. Chem.*, **62** (1997), 213.
9. F. London, *Z. Physik*, **63** (1930), 245.
10. J. E. Lennard-Jones, *Proc. R. Soc. London, Ser. A*, **106** (1924), 463.
11. T. A. Halgren, *J. Am. Chem. Soc.*, **114** (1992), 7827.
12. T. L. Hill, *J. Chem. Phys.*, **16** (1948), 399.
13. J. R. Hart, A. K. Rappé, *J. Chem. Phys.*, **97** (1992), 1109; J. M. Hayes, J. C. Greer, D. A. Morton-Blake, *J. Comp. Chem.*, **25** (2004), 1953.
14. J. Gasteiger, H. Saller, *Ang. Chem. Int. Ed.*, **24** (1985), 687.
15. M. M. Francl, L. E. Chirlian, *Rev. Comp. Chem.*, **14** (2000), 1.
16. C. I. Bayly, P. Cieplak, W. D. Cornell, P. A. Kollman, *J. Phys. Chem.*, **97** (1993), 10269.
17. E. Whalley, *Chem. Phys. Lett.*, **53** (1978), 449.
18. M. P. Hodges, A. J. Stone, S. S. Xantheas, *J. Phys. Chem. A*, **101** (1997), 9163.
19. R. W. Dixon, P. A. Kollman, *J. Comp. Chem.*, **18** (1997), 1632.
20. A. J. Stone, *J. Chem. Theo. Comp.*, **1** (2005), 1128.
21. E. V. Tsiper, K. Burke, *J. Chem. Phys.*, **120** (2004), 1153.
22. G. G. Ferency, P. J. Winn, C. A. Reynolds, *J. Phys. Chem. A*, **101** (1997), 5446.
23. M. Swart, P. Th. Van Duijnen, J. G. Snijders, *J. Comp. Chem.*, **22** (2001), 79.
24. S. Patal, C. L. Brooks III, *J. Comp. Chem.*, **25** (2003), 1.
25. K. Ramnarayan, B. G. Rao, U. C. Singh, *J. Chem. Phys.*, **92** (1990), 7057; S. W. Rick, S. J. Stuart, *Rev. Comp. Chem.*, **18** (2002), 89.
26. A. K. Rappé, W. A. Goddard III, *J. Phys. Chem.*, **95** (1991), 3358; U. Dinur, A. T. Hagler, *J. Comp. Chem.*, **16** (1995), 154.
27. P. Ren, J. W. Ponder, *J. Comp. Chem.*, **23** (2002), 1497.
28. R. Chelli, V. Schettino, P. Procacci, *J. Chem. Phys.*, **122** (2005), 234107.
29. W. Wang, R. D. Skeel, *J. Chem. Phys.*, **123** (2005), 164107.
30. H. D. Thomas, K. Chen, N. L. Allinger, *J. Am. Chem. Soc.*, **116** (1994), 5887.
31. N. L. Allinger, K. Chen, J. A. Katzenellenbogen, S. R. Wilson, G. M. Anstead, *J. Comp. Chem.*, **17** (1996), 747.
32. K. Kveseth, R. Seip, D. A. Kohl, *Acta Chem. Scand.*, **A34** (1980), 31.
33. R. Engeln, D. Consalvo, J. Reuss, *Chem. Phys.*, **160** (1992), 427.
34. J. T. Sprague, J. C. Tai, Y. Yuh, N. L. Allinger, *J. Comp. Chem.*, **8** (1987), 581.

35. A. J. Gordon, R. A. Ford, *The Chemist's Companion*, Wiley (1972).
36. G. Liang, P. C. Fox, J. P. Bowen, *J. Comp. Chem.*, **17** (1996), 940.
37. C. D. Berweger, W. F. van Gunsteren, F. Muller-Plathe, *Chem. Phys. Lett.*, **232** (1995), 429.
38. I. J. Chen, D. Yin, A. D. MacKerell Jr, *J. Comp. Chem.*, **23** (2002), 199.
39. J. C. Slater, J. G. Kirkwood, *Phys. Rev.*, **37** (1931), 682.
40. E. Osawa, K. B. Lipkowitz, *Rev. Comp. Chem.*, **6** (1995), 355.
41. P.-O. Norrby, T. Liljefors, *J. Comp. Chem.*, **19** (1998), 1146.
42. C. R. Landis, D. M. Root, T. Cleveland, *Rev. Comp. Chem.*, **6** (1995), 73; P. Comba, *Coord. Chem. Rev.*, **123** (1993), 1; B. P. Hay, *Coord. Chem. Rev.*, **126** (1993), 177.
43. B. P. Hay, *Coord. Chem. Rev.*, **126** (1993), 177.
44. R. J. Gillespie, I. Hargittai, *The VSEPR Model of Molecular Geometry*, Allyn and Bacon, Boston (1991).
45. V. J. Burton, R. J. Deeth, C. M. Kemp, P. J. Gilbert, *J. Am. Chem. Soc.*, **117** (1995), 8407.
46. A. K. Rappé, C. J. Casewit, K. S. Colwell, W. A. Goddard III, W. M. Skiff, *J. Am. Chem. Soc.*, **114** (1992), 10024.
47. D. M. Root, C. R. Landis, T. Cleveland, *J. Am. Chem. Soc.*, **115** (1993), 4201.
48. T. Hassinen, M. Perdkyld, *J. Comp. Chem.*, **22** (2001), 1229.
49. AMBER: W. D. Cornell, P. Cieplak, C. I. Bayly, I. R. Gould, K. M. Merz Jr, D. M. Ferguson, D. C. Spellmeyer, T. Fox, J. W. Caldwell, P. A. Kollman, *J. Am. Chem. Soc.*, **117** (1995), 5179; CFF91/93/95: M. J. Hwang, J. P. Stockfisch, A. T. Hagler, *J. Am. Chem. Soc.*, **116** (1994), 2515; CHARMM: R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan, M. Karplus, *J. Comp. Chem.*, **4** (1983), 187; COSMIC: S. D. Morley, R. J. Abraham, I. S. Haworth, D. E. Jackson, M. R. Saunders, J. G. Vinter, *J. Comput.-Aided Mol. Des.*, **5** (1991), 475; CVFF: S. Lifson, A. T. Hagler, P. Dauber, *J. Am. Chem. Soc.*, **101** (1979), 5111, 5122, 5131; DREIDING: S. L. Mayo, B. D. Olafson, W. A. Goddard III, *J. Phys. Chem.*, **94** (1990), 8897; EAS: E. M. Engler, J. D. Andose, P. v. R. Schleyer, *J. Am. Chem. Soc.*, **95** (1973), 8005; ECEPP: G. Nemethy, K. D. Gibbsen, K. A. Palmer, C. N. Yoon, G. Paterlini, A. Zagari, S. Rumsey, H. A. Sheraga, *J. Phys. Chem.*, **96** (1992), 6472; EFF: J. L. M. Dillen, *J. Comp. Chem.*, **16** (1995), 595, 610; ENCAD: M. Levitt, M. Hirshberg, R. Sharon, V. Daggett, *Comp. Phys. Commun.*, **91** (1995), 215; ESFF: S. Barlow, A. L. Rohl, S. Shi, C. M. Freeman, D. O'Hare, *J. Am. Chem. Soc.*, **118** (1996), 7578; S. Shi, L. Yan, Y. Yang, J. Fisher-Shaulsky, T. Tacher, *J. Comp. Chem.*, **24** (2003), 1059; GROMOS: W. R. P. Scott, P. H. Hünenberger, I. G. Tironi, A. E. Mark, S. R. Billeter, J. Fennen, A. E. Torda, T. Huber, P. Krüger, W. F. van Gunsteren, *J. Phys. Chem. A*, **103** (1999), 3596; MM2: N. L. Allinger, *J. Am. Chem. Soc.*, **99** (1977), 8127; MM3: N. L. Allinger, Y. H. Yuh, J. H. Lii, *J. Am. Chem. Soc.*, **111** (1989), 8551; J. H. Lii, N. L. Allinger, *J. Am. Chem. Soc.*, **111** (1989), 8566, 8576; N. L. Allinger, X. Zhou, J. Bergsma, *J. Mol. Struct. Theochem.*, **312** (1994), 69; MM4: N. L. Allinger, K. Chen, J.-H. Lii, *J. Comp. Chem.*, **17** (1996), 642; N. Nevins, K. Chen, N. L. Allinger, *J. Comp. Chem.*, **17** (1996), 669, N. Nevins, J.-H. Lii, N. L. Allinger, *J. Comp. Chem.*, **17** (1996), 695, N. L. Allinger, K. Chen, J. A. Katzenellenbogen, S. R. Wilson, G. M. Anstead, *J. Comp. Chem.*, **17** (1996), 747; N. L. Allinger, K.-H. Chen, J.-H. Lii, K. A. Durkin, *J. Comp. Chem.*, **24** (2003), 1447. MMFF: T. A. Halgren, *J. Comp. Chem.*, **17** (1996), 490; MOME: P. Comba, T. W. Hambley, *Molecular Modeling of Inorganic Compounds*, VCH, Weinheim (1995); SHAPES: V. S. Allured, C. M. Kelly, C. R. Landis, *J. Am. Chem. Soc.*, **113** (1991), 1; NEMO: J. M. Hermida-Ramon, S. Brdarski, G. Karlstrom, U. Berg, *J. Comp. Chem.*, **24** (2002), 161;

- OPLS: W. Damm, A. Frontera, J. Tirado-Rives, W. L. Jorgensen; *J. Comput. Chem.*, **18** (1997), 1995;
- PROSA: H. A. Stern, G. A. Kaminski, J. L. Banks, R. Zhou, B. J. Berne, R. A. Friesner, *J. Phys. Chem. B*, **103** (1999), 4730;
- QMFF: C. S. Ewig, R. Berry, U. Dinur, J.-R. Hill, M.-J. Hwang, H. Li, C. Liang, J. Maple, Z. Peng, T. P. Stockfisch, T. S. Thacher, L. Yan, X. Ni, A. T. Hagler, *J. Comp. Chem.*, **22** (2001), 1782;
- PFF: G. A. Kaminski, H. A. Stern, B. J. Berne, R. A. Friesner, Y. X. Cao, R. B. Murphy, R. Zhou, T. A. Halgren, *J. Comp. Chem.*, **23** (2002), 1515;
- SDFF: K. Palmo, B. Mannfors, N. G. Mirkin, S. Krimm, *Biopolymers*, **68** (2003), 383;
- TraPPE: C. D. Wick, J. M. Stubbs, N. Rai, J. I. Siepmann, *J. Phys. Chem. B*, **109** (2005), 18974.
- TRIPOS: M. Clark, R. D. Cramer III, N. van Opdenbosch, *J. Comp. Chem.*, **10** (1989), 982;
- J. R. Maple, M.-J. Hwang, T. P. Stockfisch, U. Dinur, M. Waldman, C. S. Ewig, A. T. Hagler, *J. Comp. Chem.*, **15** (1994), 162;
- UFF: A. K. Rappé, C. J. Casewit, K. S. Colwell, W. A. Goddard III, W. M. Skiff, *J. Am. Chem. Soc.*, **114** (1992), 10024; C. J. Casewit, K. S. Colwell, A. K. Rappé, *J. Am. Chem. Soc.*, **114** (1992), 10035, 10046;
- YETI: A. Vedani, *J. Comp. Chem.*, **9** (1988), 269.
50. S. Reiling, J. Brickmann, M. Schlenkrich, P. A. Bopp, *J. Comp. Chem.*, **17** (1996), 133.
51. I. Kolossvary, W. G. Guida, *J. Am. Chem. Soc.*, **118** (1996), 5011.
52. N. L. Allinger, S. H.-M. Chang, D. H. Glaser, H. Hönl, *Isr. J. Chem.*, **20** (1980), 51; J. P. Bowen, A. Pathiaseril, S. Profeta Jr, N. L. Allinger, *J. Org. Chem.*, **52** (1987), 5162; S. Profeta Jr, N. L. Allinger, *J. Am. Chem. Soc.*, **107** (1985), 1907; J. C. Tai, N. L. Allinger, *J. Am. Chem. Soc.*, **110** (1988), 2050; M. R. Frierson, M. R. Imam, V. B. Zalkow, N. L. Allinger, *J. Org. Chem.*, **53** (1988), 5248.
53. F. Jensen, P.-O. Norrby, *Theor. Chem. Acc.*, **109** (2003), 1.
54. J. E. Eksterowicz, K. N. Houk, *Chem. Rev.*, **93** (1993), 2439.
55. P.-O. Norrby, *J. Mol. Struct. (THEOCHEM)*, **506** (2000), 9.
56. F. M. Menger, M. J. Sherrod, *J. Am. Chem. Soc.*, **112** (1990), 8071.
57. F. Jensen, *J. Comp. Chem.*, **15** (1994), 1199, P. T. Olsen, F. Jensen, *J. Chem. Phys.*, **118** (2003), 3523.
58. A. Warshel, *J. Am. Chem. Soc.*, **102** (1980), 6218, J. Åqvist, A. Warshel, *Chem. Rev.*, **93** (1993), 2523.
59. Y. Kim, J. C. Corchado, J. Villa, J. Xing, D. G. Truhlar, *J. Chem. Phys.*, **112** (2000), 2718.
60. A. D. T. van Duin, S. Dasgupta, F. Lorant, W. A. Goddard III, *J. Phys. Chem. A*, **105** (2001), 9396.
61. M. J. Field, P. A. Bash, M. J. Karplus, *J. Comp. Chem.*, **11** (1990), 700.
62. I. Adamovic, M. S. Gordon, *Mol. Phys.*, **103** (2005), 369.
63. I. Antes, W. Thiel, *J. Phys. Chem. A*, **103** (1999), 9290.
64. V. Thery, D. Rinaldi, J.-L. Rivail, B. Maigret, G. G. Frenczy, *J. Comp. Chem.*, **15** (1994), 269; J. Gao, P. Amara, C. Alhambra, M. J. Field, *J. Phys. Chem. A*, **102** (1998), 4714.
65. M. Svensson, S. Humbel, R. D. J. Froese, T. Matsubara, S. Sieber, K. Morokuma, *J. Phys. Chem.*, **100** (1996), 19357; T. Vreven, K. Morokuma, *J. Comp. Chem.*, **21** (2000), 1419.
66. R. Prabhakar, D. G. Musaev, I. V. Khavrutskii, K. Morokuma, *J. Phys. Chem. B*, **108** (2004), 12643.
67. T. Kerdcharoen, K. Morokuma, *Chem. Phys. Lett.*, **355** (2001), 257.

3 Electronic Structure Methods: Independent-Particle Models

If we are interested in describing the electron distribution in detail, there is no substitute for quantum mechanics. Electrons are very light particles and they cannot be described correctly even qualitatively by classical mechanics. We will in this chapter and in Chapter 4 concentrate on solving the time-independent Schrödinger equation, which in shorthand operator form is given in eq. (3.1).

$$\mathbf{H}\Psi = E\Psi \quad (3.1)$$

If solutions are generated without reference to experimental data, the methods are usually called *ab initio* (latin: “from the beginning”), in contrast to semi-empirical models, which are described in Section 3.10.

An essential part of solving the Schrödinger equation is the Born–Oppenheimer approximation, where the coupling between the nuclei and electronic *motion* is neglected. This allows the electronic part to be solved with the nuclear positions as parameters, and the resulting *potential energy surface* (PES) forms the basis for solving the nuclear motion. The major computational effort is in solving the electronic Schrödinger equation for a given set of nuclear coordinates.

The dynamics of a many-electron system is very complex, and consequently requires elaborate computational methods. A significant simplification, both conceptually and computationally, can be obtained by introducing *independent-particle* models, where the motion of one electron is considered to be independent of the dynamics of all other electrons. An independent-particle model means that the interactions between the particles is approximated, either by neglecting all but the most important one, or by taking all interactions into account in an average fashion. Within electronic structure theory, only the latter has an acceptable accuracy, and is called *Hartree–Fock* (HF) theory. In the HF model, each electron is described by an orbital, and the total wave function is given as a product of orbitals. Since electrons are indistinguishable fermions (particles with a spin of $1/2$), however, the overall wave function must be antisymmetric (change

sign upon interchanging any two electrons), which is conveniently achieved by arranging the orbitals in a Slater determinant. The best set of orbitals is determined by the variational principle, i.e. the HF orbitals give the lowest energy within the restriction of the wave function being a single Slater determinant. The shape of a given molecular orbital describes the probability of finding an electron, where the attraction to all the nuclei and the average repulsion to all the other electrons are included. Since the other electrons are described by their respective orbitals, the HF equations depend on their own solutions, and must therefore be solved iteratively. When the molecular orbitals are expanded in a basis set, the resulting equations can be written as a matrix eigenvalue problem. The elements in the Fock matrix correspond to integrals of one- and two-electron operators over basis functions, multiplied by density matrix elements. The HF equations in a basis set can thus be obtained by repeated diagonalizations of a Fock matrix.

The HF model is a kind of branching point, where either additional approximations can be invoked, leading to semi-empirical methods, or it can be improved by adding additional determinants, thereby generating models that can be made to converge towards the exact solution of the electronic Schrödinger equation.¹

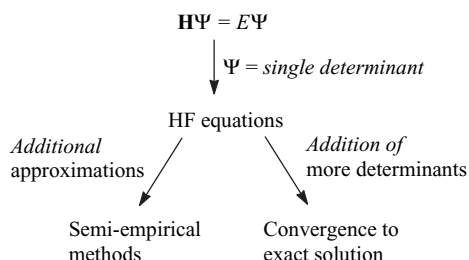


Figure 3.1 The HF model as a starting point for more approximate or more accurate treatments

Semi-empirical methods are derived from the HF model by neglecting all integrals involving more than two nuclei in the construction of the Fock matrix. Since the HF model by itself is only capable of limited accuracy, such approximations will by themselves lead to a poor model. The success of semi-empirical methods relies on turning the remaining integrals into parameters, and fitting these to experimental data, especially molecular energies and geometries. Such methods are computationally much more efficient than the *ab initio* HF method, but are limited to systems for which parameters exist.

HF theory only accounts for the average electron–electron interactions, and consequently neglects the correlation between electrons. Methods that include electron correlation require a multi-determinant wave function, since HF is *the* best single-determinant wave function. Multi-determinant methods are computationally much more involved than the HF model, but can generate results that systematically approach the exact solution of the Schrödinger equation. These methods are described in Chapter 4.

Density Functional Theory (DFT) in the Kohn–Sham version can be considered as an improvement on HF theory, where the many-body effect of electron correlation is modelled by a function of the electron density. DFT is, analogously to HF, an

independent-particle model, and is comparable to HF computationally, but provides significantly better results. The main disadvantage of DFT is that there is no systematic approach to improving the results towards the exact solution. These methods are described in Chapter 6.

We will also neglect relativistic effects in this chapter, which is justifiable for the first three rows in the periodic table (i.e. $Z < 36$) unless high accuracy is required, but the effects become important for the fourth and fifth row elements, and for transition metals. A more detailed discussion can be found in Chapter 8. Spin-dependent effects are relativistic in origin (e.g. spin-orbit interaction), but can be introduced in an *ad hoc* fashion in non-relativistic theory, and calculated as corrections (for example by means of perturbation theory) after the electronic Schrödinger equation has been solved. This will be discussed in more detail in Chapter 10.

A word of caution before we start. A rigorous approach to many of the derivations requires keeping track of several different indices and validating why certain transformations are possible. The derivations will be performed less rigorously here, with the emphasis on illustrating the flow of the argument, rather than focusing on the mathematical details.

It is conventional to use *bra-ket notation* for wave functions and multi-dimensional integrals in electronic structure theory in order to simplify the notation. The equivalences are defined in eq. (3.2).

$$\begin{aligned} |\Psi\rangle &\equiv \Psi; \quad \langle\P| \equiv \Psi^* \\ \int \Psi^* \Psi d\mathbf{r} &= \langle\P|\Psi\rangle \\ \int \Psi^* \mathbf{H} \Psi d\mathbf{r} &= \langle\P|\mathbf{H}|\Psi\rangle \end{aligned} \quad (3.2)$$

The *bra* $\langle n|$ denotes a *complex conjugate* wave function with quantum number n standing to the *left* of the operator, while the *ket* $|m\rangle$ denotes a wave function with quantum number m standing to the *right* of the operator, and the combined *bracket* denotes that the whole expression should be integrated over all coordinates. Such a bracket is often referred to as a *matrix element*, or as an *overlap element* when there is no operator involved.

3.1 The Adiabatic and Born–Oppenheimer Approximations

We will start by reviewing the Born–Oppenheimer approximation in more detail.² The total (non-relativistic) Hamiltonian operator can be written as kinetic and potential energies of the nuclei and electrons.

$$\mathbf{H}_{\text{tot}} = \mathbf{T}_n + \mathbf{T}_e + \mathbf{V}_{ne} + \mathbf{V}_{ee} + \mathbf{V}_{nn} \quad (3.3)$$

The Hamiltonian operator is first transformed to the centre of mass system, where it may be written as (using atomic units, see Appendix D):

$$\begin{aligned} \mathbf{H}_{\text{tot}} &= \mathbf{T}_n + \mathbf{H}_e + \mathbf{H}_{\text{mp}} \\ \mathbf{H}_e &= \mathbf{T}_n + \mathbf{V}_{ne} + \mathbf{V}_{ee} + \mathbf{V}_{nn} \\ \mathbf{H}_{\text{mp}} &= -\frac{1}{2M_{\text{tot}}} \left(\sum_i^{N_{\text{elec}}} \nabla_i \right)^2 \end{aligned} \quad (3.4)$$

Here \mathbf{H}_e is the *electronic Hamiltonian operator* and \mathbf{H}_{mp} is called the *mass-polarization* (M_{tot} is the total mass of all the nuclei). The mass-polarization term arises because it is not possible to rigorously separate the centre of mass motion from the internal motion for a system with more than two particles. We note that \mathbf{H}_e only depends on the nuclear *positions* (via \mathbf{V}_{ne} and \mathbf{V}_{nn} , see eq. (3.23)), but not on their *momenta*.

Assume for the moment that the full set of solutions to the electronic Schrödinger equation is available, where \mathbf{R} denotes nuclear positions and \mathbf{r} electronic coordinates.

$$\mathbf{H}_e(\mathbf{R})\Psi_i(\mathbf{R}, \mathbf{r}) = E_i(\mathbf{R})\Psi_i(\mathbf{R}, \mathbf{r}); \quad i = 1, 2, \dots, \infty \quad (3.5)$$

The Hamiltonian operator is *Hermitian*, eq. (3.6).

$$\int \Psi_i^* \mathbf{H} \Psi_j d\mathbf{r} = \int \Psi_j \mathbf{H}^* \Psi_i^* d\mathbf{r} \quad \leftrightarrow \quad \langle \Psi_i | \mathbf{H} | \Psi_j \rangle = \langle \Psi_j | \mathbf{H} | \Psi_i \rangle^* \quad (3.6)$$

The Hermitian property means that the solutions can be chosen to be orthogonal and normalized (*orthonormal*).

$$\begin{aligned} \int \Psi_i^*(\mathbf{R}, \mathbf{r}) \Psi_j(\mathbf{R}, \mathbf{r}) d\mathbf{r} &= \delta_{ij} \quad \leftrightarrow \quad \langle \Psi_i | \Psi_j \rangle = \delta_{ij} \\ \delta_{ij} &= 1, \quad i = j \\ \delta_{ij} &= 0, \quad i \neq j \end{aligned} \quad (3.7)$$

Without introducing any approximations, the total (exact) wave function can be written as an expansion in the complete set of electronic functions, with the expansion coefficients being functions of the nuclear coordinates.

$$\Psi_{\text{tot}}(\mathbf{R}, \mathbf{r}) = \sum_{i=1}^{\infty} \Psi_{ni}(\mathbf{R}) \Psi_i(\mathbf{R}, \mathbf{r}) \quad (3.8)$$

Inserting eq. (3.8) into the Schrödinger equation (3.1) gives eq. (3.9).

$$\sum_{i=1}^{\infty} (\mathbf{T}_n + \mathbf{H}_e + \mathbf{H}_{mp}) \Psi_{ni}(\mathbf{R}) \Psi_i(\mathbf{R}, \mathbf{r}) = E_{\text{tot}} \sum_{i=1}^{\infty} \Psi_{ni}(\mathbf{R}) \Psi_i(\mathbf{R}, \mathbf{r}) \quad (3.9)$$

The nuclear kinetic energy is a sum of differential operators.

$$\begin{aligned} \mathbf{T}_n &= \sum_a -\frac{1}{2M_a} \nabla_a^2 = \nabla_n^2 \\ \nabla_a &= \left(\frac{\partial}{\partial X_a}, \frac{\partial}{\partial Y_a}, \frac{\partial}{\partial Z_a} \right) \\ \nabla_a^2 &= \left(\frac{\partial^2}{\partial X_a^2} + \frac{\partial^2}{\partial Y_a^2} + \frac{\partial^2}{\partial Z_a^2} \right) \end{aligned} \quad (3.10)$$

We have here introduced the ∇_n^2 symbol, which implicitly includes the mass dependence, sign and summation. Expanding out (3.8) gives eq. (3.11).

$$\begin{aligned}
\sum_{i=1}^{\infty} (\nabla_n^2 + \mathbf{H}_e + \mathbf{H}_{\text{mp}}) \Psi_{ni} \Psi_i &= E_{\text{tot}} \sum_{i=1}^{\infty} \Psi_{ni} \Psi_i \\
\sum_{i=1}^{\infty} \{ \nabla_n^2 \Psi_{ni} \Psi_i + \mathbf{H}_e \Psi_{ni} \Psi_i + \mathbf{H}_{\text{mp}} \Psi_{ni} \Psi_i \} &= E_{\text{tot}} \sum_{i=1}^{\infty} \Psi_{ni} \Psi_i \\
\sum_{i=1}^{\infty} \{ \nabla_n (\Psi_i \nabla_n \Psi_{ni} + \Psi_{ni} \nabla_n \Psi_i) + \Psi_{ni} \mathbf{H}_e \Psi_i + \Psi_{ni} \mathbf{H}_{\text{mp}} \Psi_i \} &= E_{\text{tot}} \sum_{i=1}^{\infty} \Psi_{ni} \Psi_i \\
\sum_{i=1}^{\infty} \left\{ \Psi_i (\nabla_n^2 \Psi_{ni}) + 2(\nabla_n \Psi_i)(\nabla_n \Psi_{ni}) + \right. & \\
\left. \Psi_{ni} (\nabla_n^2 \Psi_i) + \Psi_{ni} E_i \Psi_i + \Psi_{ni} \mathbf{H}_{\text{mp}} \Psi_i \right\} &= E_{\text{tot}} \sum_{i=1}^{\infty} \Psi_{ni} \Psi_i
\end{aligned} \tag{3.11}$$

Here we have used the fact that \mathbf{H}_e and \mathbf{H}_{mp} only act on the electronic wave function, and the fact that Ψ_i is an exact solution to the electronic Schrödinger equation (eq. (3.5)). We will now use the orthonormality of the Ψ_i by multiplying from the left by a *specific* electronic wave function Ψ_j^* and integrate over the electron coordinates.

$$\nabla_n^2 \Psi_{nj} + E_j \Psi_{nj} + \sum_{i=1}^{\infty} \left\{ 2\langle \Psi_j | \nabla_n | \Psi_i \rangle (\nabla_n \Psi_{ni}) + \langle \Psi_j | \nabla_n^2 | \Psi_i \rangle \Psi_{ni} + \right. \\
\left. \langle \Psi_j | \mathbf{H}_{\text{mp}} | \Psi_i \rangle \Psi_{ni} \right\} = E_{\text{tot}} \Psi_{nj} \tag{3.12}$$

The electronic wave function has now been removed from the first two terms while the curly bracket contains terms that couple different electronic states. The first two of these are the first- and second-order *non-adiabatic coupling elements*, respectively, while the last is the mass polarization. The non-adiabatic coupling elements are important for systems involving more than one electronic surface, such as photochemical reactions.

In the *adiabatic* approximation the form of the total wave function is restricted to one electronic surface, i.e. all coupling elements in eq. (3.12) are neglected (only the terms with $i = j$ survive). Except for spatially degenerate wave functions, the diagonal first-order non-adiabatic coupling elements are zero.

$$(\nabla_n^2 + E_j + \langle \Psi_j | \nabla_n^2 | \Psi_j \rangle + \langle \Psi_j | \mathbf{H}_{\text{mp}} | \Psi_j \rangle) \Psi_{nj} = E_{\text{tot}} \Psi_{nj} \tag{3.13}$$

Neglecting the mass-polarization and reintroducing the kinetic energy operator gives eq. (3.14).

$$(\mathbf{T}_n + E_j + \langle \Psi_j | \nabla_n^2 | \Psi_j \rangle) \Psi_{nj} = E_{\text{tot}} \Psi_{nj} \tag{3.14}$$

This can also be written as in eq. (3.15).

$$(\mathbf{T}_n + E_j(\mathbf{R}) + U(\mathbf{R})) \Psi_{nj}(\mathbf{R}) = E_{\text{tot}} \Psi_{nj}(\mathbf{R}) \tag{3.15}$$

The $U(\mathbf{R})$ term is known as the *diagonal correction*, and is smaller than $E_j(\mathbf{R})$ by a factor roughly equal to the ratio of the electronic and nuclear masses. It is usually a slowly varying function of \mathbf{R} , and the shape of the energy surface is therefore determined almost exclusively by $E_j(\mathbf{R})$.³ In the *Born–Oppenheimer* approximation, the diagonal correction is neglected, and the resulting equation takes on the usual Schrödinger form, where the electronic energy plays the role of a potential energy.

$$(\mathbf{T}_n + E_j(\mathbf{R})) \Psi_{nj}(\mathbf{R}) = E_{\text{tot}} \Psi_{nj}(\mathbf{R}) \tag{3.16}$$

In the Born–Oppenheimer picture, the nuclei move on a *potential energy surface* (PES) which is a solution to the *electronic* Schrödinger equation. The PES is independent of

the nuclear masses (i.e. it is the same for isotopic molecules), but this is not the case when working in the adiabatic approximation since the diagonal correction (and mass-polarization) depends on the nuclear masses. Solving eq. (3.16) for the nuclear wave function leads to energy levels for molecular vibrations and rotations (Section 13.5), which in turn are the fundamentals for many forms of spectroscopy, such as infrared (IR), Raman, microwave, etc.

The Born–Oppenheimer (and adiabatic) approximation is usually a good approximation but breaks down when two (or more) solutions to the electronic Schrödinger equation come close together energetically.⁴ Consider for example stretching the bond in the LiF molecule. Near the equilibrium distance the molecule is very polarized, i.e. described essentially by an ionic wave function, Li^+F^- . The molecule, however, dissociates into neutral atoms (all bonds break homolytically in the gas phase), i.e. the wave function at long distance is of a covalent type, $\text{Li}\cdot\text{F}\cdot$. At the equilibrium distance, the covalent wave function is higher in energy than the ionic, but the situation reverses as the bond distance increases. At some point they must “cross”. However, as they have the same symmetry, they do not actually cross, but make an *avoided crossing*. In the region of the avoided crossing, the wave function changes from being mainly ionic to covalent over a short distance, and the adiabatic, and therefore also the Born–Oppenheimer, approximation, breaks down. This is illustrated in Figure 3.2, where the two states have been calculated by a state average MCSCF procedure using the aug-cc-pVTZ basis set. The energy of the ionic state is given by the solid line, while the energy of the covalent state is shown by the dashed line. For bond distances near 6 Å, the lowest energy wave function suddenly switches from being almost ionic to being covalent, and the two states come within ~15 kJ/mol of each other. In this region the Born–Oppenheimer approximation becomes poor.




Figure 3.2 Avoided crossing of potential energy curves for LiF

For the majority of systems the Born–Oppenheimer approximation introduces only very small errors. The *diagonal Born–Oppenheimer correction* (DBOC) can be evaluated relatively easily, as it is just the second derivative of the electronic wave function with respect to the nuclear coordinates, and is therefore closely related to the nuclear gradient and second derivative of the energy (Section 10.8).

$$\Delta E_{\text{DBOC}} = \sum_{a=1}^{N_{\text{nuc}}} -\frac{1}{2M_a} \langle \Psi_e | \nabla_a^2 | \Psi_e \rangle \quad (3.17)$$

The largest effect is expected for hydrogen-containing molecules, since hydrogen has the lightest nucleus. The absolute magnitude of the DBOC for H_2O is $\sim 7 \text{ kJ/mol}$, but the effect for the barrier towards linearity is only $\sim 0.17 \text{ kJ/mol}$.⁵ For the BH molecule, the equilibrium bond length elongates by $\sim 0.0007 \text{ \AA}$ when the DBOC is included, and the harmonic vibrational frequency changes by $\sim 2 \text{ cm}^{-1}$. For systems with heavier nuclei, the effects are expected to be substantially smaller.

When the Born–Oppenheimer approximation is expected to be poor, the non-adiabatic corrections will be large, and a better strategy in such cases may be to take the quantum nature of the nuclei into account directly. Starting from eq. (3.8), both the nuclear and electronic parts may be described by determinantal-based wave functions expanded in Gaussian basis sets. Each of the two wave functions (electronic and nuclear) can be described at different levels of approximations, with mean-field methods (i.e. Hartree–Fock) being the first step. The energy spectrum arising from such methods directly gives both nuclear (e.g. vibrations) and electronic states, but there are still some open questions as to how to formulate a consistent theory for actually carrying out such calculations.⁶

It should be noted that once methods beyond the Born–Oppenheimer approximation are employed, concepts such as molecular geometries become blurred and energy surfaces no longer exist. Nuclei are delocalized in a quantum description, and a “bond length” is no longer a unique quantity, but must be defined according to the experiment that it is compared with. An X-ray structure, for example, measures the scattering of electromagnetic radiation by the electron density, neutron diffraction measures the scattering by the nuclei, while a microwave experiment measures the moments of inertia. With nuclei as delocalized wave packages, these quantities must be obtained as averages over the electronic and nuclear wave function components.

3.2 Self-Consistent Field Theory

Having stated the limitations (non-relativistic Hamiltonian operator and the Born–Oppenheimer approximation), we are ready to consider the electronic Schrödinger equation. It can only be solved exactly for the H_2^+ molecule and similar one-electron systems. In the general case, we have to rely on approximate (numerical) methods. By neglecting relativistic effects, we also have to introduce electron spin as an *ad hoc* quantum effect. Each electron has a spin quantum number of $1/2$, and in the presence of a magnetic field there are two possible states, corresponding to alignment *along* or *opposite* to the field. The corresponding spin functions are denoted α and β , and obey the orthonormality conditions in eq. (3.18).

$$\begin{aligned}\langle\alpha|\alpha\rangle &= \langle\beta|\beta\rangle = 1 \\ \langle\alpha|\beta\rangle &= \langle\beta|\alpha\rangle = 0\end{aligned}\tag{3.18}$$

To generate approximate solutions we will employ the variational principle, which states that any approximate wave function has an energy above or equal to the exact energy (see Appendix B for a proof). The equality holds only if the wave function is the exact function. By constructing a trial wave function containing a number of parameters, we can generate the “best” trial function of the given form by minimizing the energy as a function of these parameters.

The energy of an approximate wave function can be calculated as the expectation value of the Hamiltonian operator, divided by the norm of the wave function.

$$E_e = \frac{\langle\P|\mathbf{H}_e|\Psi\rangle}{\langle\P|\Psi\rangle}\tag{3.19}$$

For a normalized wave function the denominator is 1, and therefore $E_e = \langle\P|\mathbf{H}_e|\Psi\rangle$. The total electronic wave function must be antisymmetric (change sign) with respect to interchange of any two electron coordinates (since electrons are fermions, having a spin of $1/2$). The *Pauli principle*, which states that two electrons cannot have all quantum numbers equal, is a direct consequence of this antisymmetry requirement. The antisymmetry of the wave function can be achieved by building it from *Slater determinants* (SDs). The columns in a Slater determinant are single-electron wave functions, *orbitals*, while the electron coordinates are along the rows. Let us in the following assume that we are interested in solving the electronic Schrödinger equation for a molecule. The one-electron functions are thus *molecular orbitals* (MO), which are given as the product of a spatial orbital and a spin function (α or β), also known as *spin-orbitals*, which may be taken as orthonormal. For the general case of N electrons and N spin-orbitals, the Slater determinant is given in eq. (3.20).

$$\Phi_{\text{SD}} = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(1) & \phi_2(1) & \cdots & \phi_N(1) \\ \phi_1(2) & \phi_2(2) & \cdots & \phi_N(2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(N) & \phi_2(N) & \cdots & \phi_N(N) \end{vmatrix} ; \quad \langle\phi_i|\phi_j\rangle = \delta_{ij}\tag{3.20}$$

We now make one further approximation, by taking the trial wave function to consist of a *single* Slater determinant. As will be seen later, this implies that electron correlation is neglected, or equivalently, the electron–electron repulsion is only included as an average effect. Having selected a single-determinant trial wave function the variational principle can be used to derive the *Hartree–Fock* (HF) equations, by minimizing the energy.

3.3 The Energy of a Slater Determinant

In order to derive the HF equations, we need an expression for the energy of a single Slater determinant. For this purpose, it is convenient to write it as an *antisymmetrizing operator* \mathbf{A} working on the “diagonal” of the determinant, where \mathbf{A} can be expanded as a sum of permutations. We will denote the diagonal product by Π , and use the symbol Φ to represent the determinant wave function.

$$\begin{aligned}\Phi &= \mathbf{A}[\phi_1(1)\phi_2(2)\dots\phi_N(N)] = \mathbf{A}\Pi \\ \mathbf{A} &= \frac{1}{\sqrt{N!}} \sum_{p=0}^{N-1} (-1)^p \mathbf{P} = \frac{1}{\sqrt{N!}} \left[\mathbf{1} - \sum_{ij} \mathbf{P}_{ij} + \sum_{ijk} \mathbf{P}_{ijk} - \dots \right]\end{aligned}\quad (3.21)$$

The $\mathbf{1}$ operator is the identity, while the sum over \mathbf{P}_{ij} generates all possible permutations of two electron coordinates, the sum over \mathbf{P}_{ijk} all possible permutations of three electron coordinates, etc. It may be shown that \mathbf{A} commutes with \mathbf{H} , and that \mathbf{A} acting twice gives the same as \mathbf{A} acting once, multiplied with the square root of N factorial.

$$\begin{aligned}\mathbf{A}\mathbf{H} &= \mathbf{H}\mathbf{A} \\ \mathbf{A}\mathbf{A} &= \sqrt{N!}\mathbf{A}\end{aligned}\quad (3.22)$$

Consider now the Hamiltonian operator. The nuclear–nuclear repulsion does not depend on electron coordinates and is a constant for a given nuclear geometry. The nuclear–electron attraction is a sum of terms, each depending only on one electron coordinate. The same holds for the electron kinetic energy. The electron–electron repulsion, however, depends on two electron coordinates.

$$\begin{aligned}\mathbf{H}_e &= \mathbf{T}_e + \mathbf{V}_{ne} + \mathbf{V}_{ee} + \mathbf{V}_{nn} \\ \mathbf{T}_e &= - \sum_i^{N_{elec}} \frac{1}{2} \nabla_i^2 \\ \mathbf{V}_{ne} &= - \sum_a^{N_{nuclei}} \sum_i^{N_{elec}} \frac{Z_a}{|\mathbf{R}_a - \mathbf{r}_i|} \\ \mathbf{V}_{ee} &= \sum_i^{N_{elec}} \sum_{j>i}^{N_{elec}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} \\ \mathbf{V}_{nn} &= \sum_a^{N_{nuclei}} \sum_{b>a}^{N_{nuclei}} \frac{Z_a Z_b}{|\mathbf{R}_a - \mathbf{R}_b|}\end{aligned}\quad (3.23)$$

We note that the zero point of the energy corresponds to the particles being at rest ($\mathbf{T}_e = 0$) and infinitely removed from each other ($\mathbf{V}_{ne} = \mathbf{V}_{ee} = \mathbf{V}_{nn} = 0$).

The operators may be collected according to the number of electron indices.

$$\begin{aligned}\mathbf{h}_i &= -\frac{1}{2} \nabla_i^2 - \sum_a^{N_{nuclei}} \frac{Z_a}{|\mathbf{R}_a - \mathbf{r}_i|} \\ \mathbf{g}_{ij} &= \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} \\ \mathbf{H}_e &= \sum_i^{N_{elec}} \mathbf{h}_i + \sum_{j>i}^{N_{elec}} \mathbf{g}_{ij} + \mathbf{V}_{nn}\end{aligned}\quad (3.24)$$

The one-electron operator \mathbf{h}_i describes the motion of electron i in the field of all the nuclei, and \mathbf{g}_{ij} is a two-electron operator giving the electron–electron repulsion.

The energy may be written in terms of the permutation operator as (using eqs (3.21) and (3.22))

$$\begin{aligned}
E &= \langle \Phi | \mathbf{H} | \Phi \rangle \\
&= \langle \mathbf{A} \Pi | \mathbf{H} | \mathbf{A} \Pi \rangle \\
&= \sqrt{N!} \langle \Pi | \mathbf{H} | \mathbf{A} \Pi \rangle \\
&= \sum_p (-1)^p \langle \Pi | \mathbf{H} | \mathbf{P} \Pi \rangle
\end{aligned} \tag{3.25}$$

The nuclear repulsion operator is independent of electron coordinates and can immediately be integrated to yield a constant.

$$\langle \Phi | \mathbf{V}_{nn} | \Phi \rangle = V_{nn} \langle \Phi | \Phi \rangle = V_{nn} \tag{3.26}$$

For the one-electron operator only the identity operator can give a non-zero contribution. For coordinate 1 this yields a matrix element over orbital 1.

$$\begin{aligned}
\langle \Pi | \mathbf{h}_1 | \Pi \rangle &= \langle \phi_1(1)\phi_2(2)\dots\phi_N(N) | \mathbf{h}_1 | \phi_1(1)\phi_2(2)\dots\phi_N(N) \rangle \\
&= \langle \phi_1(1) | \mathbf{h}_1 | \phi_1(1) \rangle \langle \phi_2(2) | \phi_2(2) \rangle \dots \langle \phi_N(N) | \phi_N(N) \rangle \\
&= \langle \phi_1(1) | \mathbf{h}_1 | \phi_1(1) \rangle = h_1
\end{aligned} \tag{3.27}$$

This follows since all the MOs ϕ_i are normalized. All matrix elements involving a permutation operator gives zero. Consider for example the permutation of electrons 1 and 2.

$$\begin{aligned}
\langle \Pi | \mathbf{h}_1 | \mathbf{P}_{12} \Pi \rangle &= \langle \phi_1(1)\phi_2(2)\dots\phi_N(N) | \mathbf{h}_1 | \phi_2(1)\phi_1(2)\dots\phi_N(N) \rangle \\
&= \langle \phi_1(1) | \mathbf{h}_1 | \phi_2(1) \rangle \langle \phi_2(2) | \phi_1(2) \rangle \dots \langle \phi_N(N) | \phi_N(N) \rangle
\end{aligned} \tag{3.28}$$

This is zero as the integral over electron 2 is an overlap of two different MOs, which are orthogonal (eq. (3.20)).

For the two-electron operator, only the identity and \mathbf{P}_{ij} operators can give non-zero contributions. A three-electron permutation will again give a least one overlap integral between two different MOs, which will be zero. The term arising from the identity operator is given by eq. (3.29).

$$\begin{aligned}
\langle \Pi | \mathbf{g}_{12} | \Pi \rangle &= \langle \phi_1(1)\phi_2(2)\dots\phi_N(N) | \mathbf{g}_{12} | \phi_1(1)\phi_2(2)\dots\phi_N(N) \rangle \\
&= \langle \phi_1(1)\phi_2(2) | \mathbf{g}_{12} | \phi_1(1)\phi_2(2) \rangle \dots \langle \phi_N(N) | \phi_N(N) \rangle \\
&= \langle \phi_1(1)\phi_2(2) | \mathbf{g}_{12} | \phi_1(1)\phi_2(2) \rangle = J_{12}
\end{aligned} \tag{3.29}$$

The J_{12} matrix element is called a *Coulomb* integral. It represents the classical repulsion between two charge distributions described by $\phi_1^2(1)$ and $\phi_2^2(2)$. The term arising from the \mathbf{P}_{ij} operator is given in eq. (3.30).

$$\begin{aligned}
\langle \Pi | \mathbf{g}_{12} | \mathbf{P}_{12} \Pi \rangle &= \langle \phi_1(1)\phi_2(2)\dots\phi_N(N) | \mathbf{g}_{12} | \phi_2(1)\phi_1(2)\dots\phi_N(N) \rangle \\
&= \langle \phi_1(1)\phi_2(2) | \mathbf{g}_{12} | \phi_2(1)\phi_1(2) \rangle \dots \langle \phi_N(N) | \phi_N(N) \rangle \\
&= \langle \phi_1(1)\phi_2(2) | \mathbf{g}_{12} | \phi_2(1)\phi_1(2) \rangle = K_{12}
\end{aligned} \tag{3.30}$$

The K_{12} matrix element is called an *exchange* integral, and has no classical analogy. Note that the order of the MOs in the J and K matrix elements is according to the electron indices. The energy can thus be written as in eq. (3.31).

$$\begin{aligned}
E &= \sum_{i=1}^{N_{\text{elec}}} \langle \phi_i(1) | \mathbf{h}_1 | \phi_i(1) \rangle + \\
&\quad \sum_{i=1}^{N_{\text{elec}}} \sum_{j>i}^{N_{\text{elec}}} (\langle \phi_i(1) \phi_j(2) | \mathbf{g}_{12} | \phi_i(1) \phi_j(2) \rangle - \langle \phi_i(1) \phi_j(2) | \mathbf{g}_{12} | \phi_j(1) \phi_i(2) \rangle) + V_{\text{nn}} \quad (3.31) \\
E &= \sum_{i=1}^{N_{\text{elec}}} h_i + \sum_{i=1}^{N_{\text{elec}}} \sum_{j>i}^{N_{\text{elec}}} (J_{ij} - K_{ij}) + V_{\text{nn}}
\end{aligned}$$

The minus sign for the exchange term comes from the factor of $(-1)^p$ in the antisymmetrizing operator, eq. (3.21). The energy may also be written in a more symmetrical form as in eq. (3.32).

$$E = \sum_{i=1}^{N_{\text{elec}}} h_i + \frac{1}{2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j=1}^{N_{\text{elec}}} (J_{ij} - K_{ij}) + V_{\text{nn}} \quad (3.32)$$

The factor of $1/2$ allows the double sum to run over all electrons (it is easily seen from eqs (3.29) and (3.30) that the Coulomb “self-interaction” J_{ii} is exactly cancelled by the corresponding “exchange” element K_{ii}).

For the purpose of deriving the variation of the energy, it is convenient to express the energy in terms of Coulomb (\mathbf{J}) and exchange (\mathbf{K}) operators.

$$\begin{aligned}
E &= \sum_i^{N_{\text{elec}}} \langle \phi_i | \mathbf{h}_i | \phi_i \rangle + \frac{1}{2} \sum_{ij}^{N_{\text{elec}}} (\langle \phi_j | \mathbf{J}_i | \phi_j \rangle - \langle \phi_j | \mathbf{K}_i | \phi_j \rangle) + V_{\text{nn}} \\
\mathbf{J}_i | \phi_j(2) \rangle &= \langle \phi_i(1) | \mathbf{g}_{12} | \phi_i(1) \rangle | \phi_j(2) \rangle \\
\mathbf{K}_i | \phi_j(2) \rangle &= \langle \phi_i(1) | \mathbf{g}_{12} | \phi_j(1) \rangle | \phi_i(2) \rangle
\end{aligned} \quad (3.33)$$

Note that the \mathbf{J} operator involves “multiplication” with a matrix element with the same orbital on both sides, while the \mathbf{K} operator “exchanges” the two functions on the right-hand side of the \mathbf{g}_{12} operator.

The objective is now to determine a set of MOs that makes the energy a minimum, or at least stationary with respect to a change in the orbitals. The variation, however, must be carried out in such a way that the MOs remain orthogonal and normalized. This is a constrained optimization, and can be handled by means of *Lagrange multipliers* (see Section 12.5). The condition is that a small change in the orbitals should not change the Lagrange function, i.e. the Lagrange function is stationary with respect to an orbital variation.

$$\begin{aligned}
L &= E - \sum_{ij}^{N_{\text{elec}}} \lambda_{ij} (\langle \phi_i | \phi_j \rangle - \delta_{ij}) \\
\delta L &= \delta E - \sum_{ij}^{N_{\text{elec}}} \lambda_{ij} (\langle \delta \phi_i | \phi_j \rangle - \langle \phi_i | \delta \phi_j \rangle) = 0
\end{aligned} \quad (3.34)$$

The variation of the energy is given by eq. (3.35).

$$\begin{aligned}
\delta E &= \sum_i^{N_{\text{elec}}} (\langle \delta \phi_i | \mathbf{h}_i | \phi_i \rangle + \langle \phi_i | \mathbf{h}_i | \delta \phi_i \rangle) + \\
&\quad \frac{1}{2} \sum_{ij}^{N_{\text{elec}}} \left(\langle \delta \phi_i | \mathbf{J}_j - \mathbf{K}_j | \phi_j \rangle + \langle \phi_i | \mathbf{J}_j - \mathbf{K}_j | \delta \phi_j \rangle + \right. \\
&\quad \left. \langle \delta \phi_j | \mathbf{J}_i - \mathbf{K}_i | \phi_i \rangle + \langle \phi_j | \mathbf{J}_i - \mathbf{K}_i | \delta \phi_i \rangle \right)
\end{aligned} \quad (3.35)$$

The third and fifth terms are identical (since the summation is over all i and j), as are the fourth and sixth terms. They may be collected to cancel the factor of $1/2$, and the variation can be written in terms of a *Fock operator*, \mathbf{F}_i .

$$\begin{aligned}\delta E &= \sum_i^{N_{\text{elec}}} (\langle \delta\phi_i | \mathbf{h}_i | \phi_i \rangle + \langle \phi_i | \mathbf{h}_i | \delta\phi_i \rangle) + \sum_{ij}^{N_{\text{elec}}} (\langle \delta\phi_i | \mathbf{J}_j - \mathbf{K}_j | \phi_i \rangle + \langle \phi_i | \mathbf{J}_j - \mathbf{K}_j | \delta\phi_i \rangle) \\ \delta E &= \sum_i^{N_{\text{elec}}} (\langle \delta\phi_i | \mathbf{F}_i | \phi_i \rangle + \langle \phi_i | \mathbf{F}_i | \delta\phi_i \rangle) \\ \mathbf{F}_i &= \mathbf{h}_i + \sum_j^{N_{\text{elec}}} (\mathbf{J}_j - \mathbf{K}_j)\end{aligned}\tag{3.36}$$

The Fock operator is an effective one-electron energy operator, describing the kinetic energy of an electron and the attraction to all the nuclei (\mathbf{h}_i), as well as the repulsion to all the other electrons (via the \mathbf{J} and \mathbf{K} operators). Note that the Fock operator is associated with the *variation* of the total energy, not the energy itself. The Hamiltonian operator (3.23) is *not* a sum of Fock operators.

The variation of the Lagrange function (eq. (3.34)) now becomes eq. (3.37).

$$\delta L = \sum_i^{N_{\text{elec}}} (\langle \delta\phi_i | \mathbf{F}_i | \phi_i \rangle + \langle \phi_i | \mathbf{F}_i | \delta\phi_i \rangle) - \sum_{ij}^{N_{\text{elec}}} \lambda_{ij} (\langle \delta\phi_i | \phi_j \rangle + \langle \phi_i | \delta\phi_j \rangle)\tag{3.37}$$

The variational principle states that the desired orbitals are those that make $\delta L = 0$. Making use of the complex conjugate properties in eq. (3.38) gives eq. (3.39).

$$\begin{aligned}\langle \phi | \delta\phi \rangle &= \langle \delta\phi | \phi \rangle^* \\ \langle \phi | \mathbf{F} | \delta\phi \rangle &= \langle \delta\phi | \mathbf{F} | \phi \rangle^*\end{aligned}\tag{3.38}$$

$$\delta L = \sum_i^{N_{\text{elec}}} \langle \delta\phi_i | \mathbf{F}_i | \phi_i \rangle - \sum_{ij}^{N_{\text{elec}}} \lambda_{ij} \langle \delta\phi_i | \phi_j \rangle + \sum_i^{N_{\text{elec}}} \langle \delta\phi_i | \mathbf{F}_i | \phi_i \rangle^* - \sum_{ij}^{N_{\text{elec}}} \lambda_{ij} \langle \delta\phi_j | \phi_i \rangle^* = 0\tag{3.39}$$

The variation of *either* $\langle \delta\phi |$ or $\langle \delta\phi |^*$ should make $\delta L = 0$, i.e. the first two terms in eq. (3.39) must cancel, and the last two terms must cancel. Taking the complex conjugate of the last two terms and subtracting them from the first two gives eq. (3.40).

$$\sum_{ij}^{N_{\text{elec}}} (\lambda_{ij} - \lambda_{ji}^*) \langle \delta\phi_i | \phi_j \rangle = 0\tag{3.40}$$

This means that the Lagrange multipliers are elements of a Hermitian matrix ($\lambda_{ij} = \lambda_{ji}^*$). The final set of *Hartree–Fock equations* may be written as in eq. (3.41).

$$\mathbf{F}_i \phi_i = \sum_j^{N_{\text{elec}}} \lambda_{ij} \phi_j\tag{3.41}$$

The equations may be simplified by choosing a unitary transformation (Section 16.2) that makes the matrix of Lagrange multipliers diagonal, i.e. $\lambda_{ij} = 0$ and $\lambda_{ii} = \varepsilon_i$. This special set of molecular orbitals (ϕ') is called *canonical* MOs, and transforms eq. (3.41) into a set of pseudo-eigenvalue equations.

$$\mathbf{F}_i \phi'_i = \varepsilon_i \phi'_i\tag{3.42}$$

The Lagrange multipliers are seen to have the physical interpretation of MO energies, i.e. they are the expectation value of the Fock operator in the MO basis (multiply eq. (3.42) by $\phi_i'^*$ from the left and integrate).

$$\langle \phi_i' | \mathbf{F} | \phi_i' \rangle = \varepsilon_i \langle \phi_i' | \phi_i' \rangle = \varepsilon_i \quad (3.43)$$

The Hartree–Fock equations form a set of pseudo-eigenvalue equations as the Fock operator depends on all the occupied MOs (via the Coulomb and exchange operators, eqs (3.36) and (3.33)). A specific Fock orbital can only be determined if all the other occupied orbitals are known, and iterative methods must therefore be employed for solving the problem. A set of functions that is a solution to eq. (3.42) is called *self-consistent field (SCF)* orbitals.

The canonical MOs may be considered as a convenient set of orbitals for carrying out the variational calculation. The total energy, however, depends only on the total wave function, which is a Slater determinant written in terms of the occupied MOs, eq. (3.20). The total wave function is unchanged by a unitary transformation of the occupied MOs among themselves (rows and columns in a determinant can be added and subtracted without affecting the determinant itself). After having determined the canonical MOs, other sets of MOs may be generated by forming linear combinations, such as localized MOs, or MOs displaying hybridization, which is discussed in more detail in Section 9.4.

The orbital energies can be considered as matrix elements of the Fock operator with the MOs (dropping the prime notation and letting ϕ be the canonical orbitals). The total energy can be written either as eq. (3.32) or in terms of MO energies (using the definition of \mathbf{F} in eqs (3.36) and (3.43)).

$$E = \sum_i^{N_{\text{elec}}} \varepsilon_i - \frac{1}{2} \sum_{ij}^{N_{\text{elec}}} (J_{ij} - K_{ij}) + V_{\text{nn}} \quad (3.44)$$

$$\varepsilon_i = \langle \phi_i | \mathbf{F} | \phi_i \rangle = h_i + \sum_j^{N_{\text{elec}}} (J_{ij} - K_{ij})$$

The total energy is *not* simply a sum of MO orbital energies. The Fock operator contains terms describing the repulsion to all other electrons (\mathbf{J} and \mathbf{K} operators), and the sum over MO energies therefore counts the electron–electron repulsion twice, which must be corrected for. It is also clear that the total energy cannot be exact, as it describes the repulsion between an electron and all the other electrons, assuming that their spatial distribution is described by a set of orbitals. The electron–electron repulsion is only accounted for in an average fashion, and the HF method is therefore also referred to as a *mean-field* approximation. As mentioned previously, this is due to the approximation of a single Slater determinant as the trial wave function.

3.4 Koopmans' Theorem

The canonical MOs are convenient for the physical interpretation of the Lagrange multipliers. Consider the energy of an N -electron system and the corresponding system with one electron removed from orbital number k , and assume that the MOs are identical for the two systems (eq. (3.32)).

$$\begin{aligned}
 E_N &= \sum_{i=1}^{N_{\text{elec}}} h_i + \frac{1}{2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j=1}^{N_{\text{elec}}} (J_{ij} - K_{ij}) + V_{\text{nn}} \\
 E_{N-1} &= \sum_{i=1}^{N_{\text{elec}}-1} h_i + \frac{1}{2} \sum_{i=1}^{N_{\text{elec}}-1} \sum_{j=1}^{N_{\text{elec}}-1} (J_{ij} - K_{ij}) + V_{\text{nn}}
 \end{aligned} \tag{3.45}$$

Subtracting the two total energies gives eq. (3.46).

$$E_N - E_{N-1} = h_k + \frac{1}{2} \sum_{i=1}^{N_{\text{elec}}} (J_{ik} - K_{ik}) + \frac{1}{2} \sum_{j=1}^{N_{\text{elec}}} (J_{kj} - K_{kj}) \tag{3.46}$$

The last two sums are identical and the energy difference becomes eq. (3.47).

$$E_N - E_{N-1} = h_k + \sum_{i=1}^{N_{\text{elec}}} (J_{ik} - K_{ik}) = \epsilon_k \tag{3.47}$$

As seen from eq. (3.44), this is exactly the orbital energy ϵ_k . The ionization energy within the “frozen MO” approximation is given simply as the orbital energy, a result known as *Koopmans’ theorem*.⁷ Similarly, the electron affinity of a neutral molecule is given as the orbital energy of the corresponding anion, or, since the MOs are assumed constant, as the energy of the k th unoccupied orbital energy in the neutral species.

$$E_{N+1}^k - E_N = \epsilon_k \tag{3.48}$$

Computationally, however, there is a significant difference between the eigenvalue of an occupied orbital for the anion and the eigenvalue corresponding to an unoccupied orbital in the neutral species when the orbitals are expanded in a set of basis functions (Section 3.5). Eigenvalues corresponding to occupied orbitals are well defined and they converge to a specific value as the size of the basis set is increased. In contrast, unoccupied orbitals in a sense are only the “left-over” functions in a given basis set, and their number increases as the basis set is made larger. The lowest unoccupied eigenvalue usually converges to zero, corresponding to a solution for a free electron, described by a linear combination of the most diffuse basis functions. Equating ionization potentials to occupied orbital energies is therefore justified based on the frozen MO approximation, but taking unoccupied orbital energies as electron affinities is questionable, since continuum solutions are mixed in.

3.5 The Basis Set Approximation

For small highly symmetric systems, such as atoms and diatomic molecules, the Hartree–Fock equations may be solved by mapping the orbitals on a set of grid points, and these are referred to as *numerical Hartree–Fock* methods.⁸ However, essentially all calculations use a basis set expansion to express the unknown MOs in terms of a set of known functions. Any type of basis functions may in principle be used:

exponential, Gaussian, polynomial, cube functions, wavelets, plane waves, etc. There are two guidelines for choosing the basis functions. One is that they should have a behaviour that agrees with the physics of the problem, since this ensures that the convergence as more basis functions are added is reasonably rapid. For bound atomic and molecular systems, this means that the functions should go toward zero as the distance between the nucleus and the electron becomes large. The second guideline is a practical one: the chosen functions should make it easy to calculate all the required integrals.

The first criterion suggests the use of exponential functions located on the nuclei, since such functions are known to be exact solutions for the hydrogen atom. Unfortunately, exponential functions turn out to be computationally difficult. Gaussian functions are computationally much easier to handle, and although they are poorer at describing the electronic structure on a one-to-one basis, the computational advantages more than make up for this. For periodic systems, the infinite nature of the problem suggests the use of plane waves as basis functions, since these are the exact solutions for a free electron. We will return to the precise description of basis sets in Chapter 5, but for now simply assume that a set of M_{basis} basis functions located on the nuclei has been chosen.

Each MO ϕ is expanded in terms of the basis functions χ , conventionally called *atomic orbitals* (MO = LCAO, *Linear Combination of Atomic Orbitals*), although they are generally not solutions to the atomic HF problem.

$$\phi_i = \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} \quad (3.49)$$

The Hartree–Fock equations (3.42) may be written as in eq. (3.50).

$$\mathbf{F}_i \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} = \epsilon_i \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} \quad (3.50)$$

Multiplying from the left by a specific basis function and integrating yields the *Roothaan–Hall* equations (for a closed shell system).⁹ These are the Hartree–Fock equations in the *atomic orbital basis*, and all the M_{basis} equations may be collected in a matrix notation.

$$\begin{aligned} \mathbf{FC} &= \mathbf{SC}\epsilon \\ F_{\alpha\beta} &= \langle \chi_{\alpha} | \mathbf{F} | \chi_{\beta} \rangle \\ S_{\alpha\beta} &= \langle \chi_{\alpha} | \chi_{\beta} \rangle \end{aligned} \quad (3.51)$$

The \mathbf{S} matrix contains the overlap elements between basis functions, and the \mathbf{F} matrix contains the Fock matrix elements. Each $F_{\alpha\beta}$ element contains two parts from the Fock operator (eq. (3.36)), integrals involving the one-electron operators, and a sum over occupied MOs of coefficients multiplied with two-electron integrals involving the electron–electron repulsion. The latter is often written as a product of a *density* matrix and two-electron integrals.

$$\begin{aligned}
\langle \chi_\alpha | \mathbf{F} | \chi_\beta \rangle &= \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle + \sum_j^{\text{occ.MO}} \langle \chi_\alpha | \mathbf{J}_j - \mathbf{K}_j | \chi_\beta \rangle \\
&= \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle + \sum_j^{\text{occ.MO}} (\langle \chi_\alpha \phi_j | \mathbf{g} | \chi_\beta \phi_j \rangle - \langle \chi_\alpha \phi_j | \mathbf{g} | \phi_j \chi_\beta \rangle) \\
&= \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle + \sum_j^{\text{occ.MO}} \sum_{\gamma\delta}^{M_{\text{basis}}} c_{\gamma j} c_{\delta j} (\langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\delta \chi_\beta \rangle) \quad (3.52) \\
&= \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle + \sum_{\gamma\delta}^{M_{\text{basis}}} D_{\gamma\delta} (\langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\delta \chi_\beta \rangle) \\
D_{\gamma\delta} &= \sum_j^{\text{occ.MO}} c_{\gamma j} c_{\delta j}
\end{aligned}$$

For use in Section 3.8, it can also be written in a more compact notation.

$$\begin{aligned}
F_{\alpha\beta} &= h_{\alpha\beta} + \sum_{\gamma\delta} G_{\alpha\beta\gamma\delta} D_{\gamma\delta} \\
\mathbf{F} &= \mathbf{h} + \mathbf{G} \cdot \mathbf{D}
\end{aligned} \quad (3.53)$$

Here $\mathbf{G} \cdot \mathbf{D}$ denotes the contraction of the \mathbf{D} matrix with the four-dimensional \mathbf{G} tensor.

The total energy (eq. (3.32)) in term of integrals over basis functions is given in eq. (3.54).

$$\begin{aligned}
E &= \sum_i^{N_{\text{elec}}} \langle \phi_i | \mathbf{h} | \phi_i \rangle + \frac{1}{2} \sum_{ij}^{N_{\text{elec}}} (\langle \phi_i \phi_j | \mathbf{g} | \phi_i \phi_j \rangle - \langle \phi_i \phi_j | \mathbf{g} | \phi_j \phi_i \rangle) + V_{\text{nn}} \\
&= \sum_i^{N_{\text{elec}}} \sum_{\alpha\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta i} \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle + \frac{1}{2} \sum_{ij}^{N_{\text{elec}}} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} c_{\alpha i} c_{\gamma j} c_{\beta i} c_{\delta j} \left(\langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\delta \chi_\beta \rangle \right) + V_{\text{nn}} \quad (3.54) \\
&= \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} h_{\alpha\beta} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} (\langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\delta \chi_\beta \rangle) + V_{\text{nn}}
\end{aligned}$$

The latter expression may also be written as in eq. (3.55).

$$E = \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} h_{\alpha\beta} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} (D_{\alpha\beta} D_{\gamma\delta} - D_{\alpha\delta} D_{\gamma\beta}) \langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\beta \chi_\delta \rangle + V_{\text{nn}} \quad (3.55)$$

The one- and two-electron integrals in the atomic basis are given as eq. (3.24).

$$\begin{aligned}
\langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle &= \int \chi_\alpha(1) \left(-\frac{1}{2} \nabla^2 \right) \chi_\beta(1) d\mathbf{r}_1 + \sum_a^{N_{\text{nuclei}}} \int \chi_\alpha(1) \left(\frac{Z_a}{|\mathbf{R}_a - \mathbf{r}_1|} \right) \chi_\beta(1) d\mathbf{r}_1 \\
\langle \chi_\alpha \chi_\gamma | \mathbf{g} | \chi_\beta \chi_\delta \rangle &= \int \chi_\alpha(1) \chi_\gamma(2) \left(\frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \right) \chi_\beta(1) \chi_\delta(2) d\mathbf{r}_1 d\mathbf{r}_2
\end{aligned} \quad (3.56)$$

The two-electron integrals are often written in a notation without electron coordinates or the \mathbf{g} operator present.

$$\int \chi_\alpha(1) \chi_\gamma(2) \left(\frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \right) \chi_\beta(1) \chi_\delta(2) d\mathbf{r}_1 d\mathbf{r}_2 = \langle \chi_\alpha \chi_\gamma | \chi_\beta \chi_\delta \rangle \quad (3.57)$$

This is known as the *physicist's* notation, where the ordering of the functions is given by the electron indices. They may also be written in an alternative order with both functions depending on electron 1 on the left, and the functions depending on electron 2 on the right, this is known as the *Mulliken* or *chemist's* notation.

$$\int \chi_\alpha(1)\chi_\beta(1)\left(\frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|}\right)\chi_\gamma(2)\chi_\delta(2)d\mathbf{r}_1d\mathbf{r}_2 = (\chi_\alpha\chi_\beta|\chi_\gamma\chi_\delta) \quad (3.58)$$

The bra-ket notation has the electron indices $\langle 12|12\rangle$, while the parenthesis notation has the order $(11|22)$. In many cases the integrals are written with only the indices given, i.e. $\langle \chi_\alpha\chi_\beta|\chi_\gamma\chi_\delta\rangle = \langle \alpha\beta|\gamma\delta\rangle$. Since Coulomb and exchange integrals often are used as their difference, the following double-bar notations are also used frequently.

$$\begin{aligned} \langle \chi_\alpha\chi_\beta||\chi_\gamma\chi_\delta\rangle &= \langle \chi_\alpha\chi_\beta|\chi_\gamma\chi_\delta\rangle - \langle \chi_\alpha\chi_\beta|\chi_\delta\chi_\gamma\rangle \\ (\chi_\alpha\chi_\beta||\chi_\gamma\chi_\delta) &= (\chi_\alpha\chi_\beta|\chi_\gamma\chi_\delta) - (\chi_\alpha\chi_\gamma|\chi_\beta\chi_\delta) \end{aligned} \quad (3.59)$$

The Roothaan–Hall equation (3.51) is a determination of the eigenvalues of the Fock matrix (see Section 16.2.3 for details). To determine the unknown MO coefficients c_{ai} , the Fock matrix must be diagonalized. However, the Fock matrix is only known if all the MO coefficients are known (eq. (3.52)). The procedure therefore starts off by some guess of the coefficients, forms the \mathbf{F} matrix, and diagonalizes it. The new set of coefficients is then used for calculating a new Fock matrix, etc. This is continued until the set of coefficients used for constructing the Fock matrix is equal to those resulting from the diagonalization (to within a certain threshold). This set of coefficients determines a self-consistent field solution.




Figure 3.3 Illustration of the SCF procedure

The potential (or field) generated by the SCF electron density is identical to that produced by solving for the electron distribution. The Fock matrix, and therefore the total energy, only depends on the occupied MOs. Solving the Roothaan–Hall equations produces a total of M_{basis} MOs, i.e. there are N_{elec} occupied and $M_{\text{basis}} - N_{\text{elec}}$ unoccupied, or *virtual*, MOs. The virtual orbitals are orthogonal to all the occupied orbitals,

but have no direct physical interpretation, except as electron affinities (via Koopmans' theorem).

In order to construct the Fock matrix in eq. (3.51), integrals between all pairs of basis functions and the one-electron operator \mathbf{h} are needed. For M_{basis} functions there are of the order of M_{basis}^2 such *one-electron integrals*. These one-electron integrals are also known as *core integrals*, as they describe the interaction of an electron with the whole frame of bare nuclei. The second part of the Fock matrix involves integrals over four basis functions and the \mathbf{g} two-electron operator. There are of the order of M_{basis}^4 of these *two-electron integrals*. In conventional HF methods, the two-electron integrals are calculated and saved before the SCF procedure is begun, and is then used in each SCF iteration. *Formally, in the large basis set limit the SCF procedure involves a computational effort that increases as the number of basis functions to the fourth power.* It will be shown below that the scaling may be substantially smaller in actual calculations.

For the two-electron integrals, the four basis functions may be located on one, two, three or four different atomic centres. It has already been mentioned that exponential-type basis functions ($\chi \propto \exp(-\alpha r)$) are fundamentally better suited for electronic structure calculations. However, it turns out that the calculation of especially three- and four-centre two-electron integrals is very time-consuming for exponential functions. Gaussian functions ($\chi \propto \exp(-\alpha r^2)$) are much easier for calculating two-electron integrals. This is due to the fact that the product of two Gaussians located at two different positions (\mathbf{R}_A and \mathbf{R}_B) with different exponents (α and β) can be written as a single Gaussian located at an intermediate position \mathbf{R}_C between the two original. This allows compact formulas for all types of one- and two-electron integrals to be derived.

$$\begin{aligned}
 G_A(\mathbf{r}) &= \left(\frac{2\alpha}{\pi} \right)^{3/4} e^{-\alpha(\mathbf{r}+\mathbf{R}_A)^2} \\
 G_B(\mathbf{r}) &= \left(\frac{2\beta}{\pi} \right)^{3/4} e^{-\beta(\mathbf{r}+\mathbf{R}_B)^2} \\
 G_A(\mathbf{r})G_B(\mathbf{r}) &= K e^{-\gamma(\mathbf{r}+\mathbf{R}_C)^2} \\
 \gamma &= \alpha + \beta \\
 \mathbf{R}_C &= \frac{\alpha\mathbf{R}_A + \beta\mathbf{R}_B}{\alpha + \beta} \\
 K &= \left(\frac{2}{\pi} \right)^2 (\alpha\beta)^{3/4} e^{-\frac{\alpha\beta}{\alpha+\beta}(\mathbf{R}_A-\mathbf{R}_B)^2}
 \end{aligned} \tag{3.60}$$

As the number of basis functions increases, the accuracy of the MOs improves. In the limit of a complete basis set (infinite number of basis functions), the results are identical to those obtained by a numerical HF method, and this is known as the *Hartree–Fock limit*. This is *not* the exact solution to the Schrödinger equation, only the best single-determinant wave function that can be obtained. In practical calculations, the HF limit is never reached, and the term Hartree–Fock is normally used also to cover SCF solutions with an incomplete basis set. *Ab initio* HF methods, where all the necessary integrals are calculated from a given basis set, are one-dimensional. As the

size of the basis set is increased, the variational principle ensures that the results become better (at least in an energetic sense). The quality of a result can therefore be assessed by running calculations with an increasingly larger basis set.

3.6 An Alternative Formulation of the Variational Problem

The objective is to minimize the total energy as a function of the molecular orbitals, subject to the orthogonality constraint. In the above formulation, this was handled by means of Lagrange multipliers. The final Fock matrix in the MO basis is diagonal, with the diagonal elements being the orbital energies. During the iterative sequence, i.e. before the orbitals have converged to an SCF solution, the Fock matrix is not diagonal. Starting from an initial set of molecular orbitals, the problem may also be formulated as a rotation of the orbitals (unitary transformation) in order to make the operator diagonal.¹⁰ Since the operator depends on the orbitals, the procedure again becomes iterative.

The orbital rotation is given by a unitary matrix \mathbf{U} , which can be written as an exponential transformation.

$$\phi' = \mathbf{U}\phi = e^{\mathbf{X}}\phi \quad (3.61)$$

The \mathbf{X} matrix contains the parameters describing the unitary transformation of the M_{basis} orbitals, being of the size of $M_{\text{basis}} \times M_{\text{basis}}$. The orthogonality is incorporated by requiring that the \mathbf{X} matrix is antisymmetric, $x_{ij} = -x_{ji}$.

$$\mathbf{U}^\dagger \mathbf{U} = (e^{\mathbf{X}})^\dagger (e^{\mathbf{X}}) = (e^{\mathbf{X}^\dagger})(e^{\mathbf{X}}) = (e^{-\mathbf{X}})(e^{\mathbf{X}}) = 1 \quad (3.62)$$

Normally the orbitals are real, and the unitary transformation becomes an orthogonal transformation. In the case of only two orbitals, the \mathbf{X} matrix contains the rotation angle α , and the \mathbf{U} matrix describes a 2×2 rotation. The connection between \mathbf{X} and \mathbf{U} is illustrated in Section 16.2 (Figure 16.3) and involves diagonalization of \mathbf{X} (to give eigenvalues of $\pm i\alpha$), exponentiation (to give complex exponentials that may be written as $\cos \alpha \pm i \sin \alpha$), followed by back-transformation.

$$\begin{aligned} \mathbf{X} &= \begin{pmatrix} 0 & \alpha \\ -\alpha & 0 \end{pmatrix} \\ \mathbf{U} = e^{\mathbf{X}} &= \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \end{aligned} \quad (3.63)$$

In the general case, the \mathbf{X} matrix contains rotational angles for rotating all pairs of orbitals.

It should be noted that the unoccupied orbitals do not enter the energy expression (eq. (3.32)), and a rotation between the virtual orbitals can therefore not change the energy. A rotation between the occupied orbitals corresponds to making linear combinations of these, but this does not change the total wave function or the total energy. The occupied–occupied and virtual–virtual blocks of the \mathbf{X} matrix can therefore be chosen as zero. The variational parameters are the elements in the \mathbf{X} matrix that describe the mixing of the occupied and virtual orbitals, i.e. there are a total of $N_{\text{occ}} \times (M_{\text{basis}} - N_{\text{occ}})$ parameters. The goal of the iterations is to make the off-diagonal elements in the occupied–virtual block of the Fock matrix zero. Alternatively stated, the

off-diagonal elements are the gradients of the energy with respect to the orbitals, and the stationary condition is that the gradient vanishes.

Using the concepts from Chapter 16, the variational problem can be considered as a rotation of the coordinate system. In the original function space, the basis functions, the Fock operator depends on all the M_{basis} functions, and the corresponding Fock matrix is non-diagonal. By performing a rotation of the coordinate system to the molecular orbitals, however, the matrix can be made diagonal, i.e. in this coordinate system the Fock operator only depends on N_{occ} functions.

3.7 Restricted and Unrestricted Hartree-Fock

So far there has not been any restriction on the MOs used to build the determinantal trial wave function. The Slater determinant has been written in terms of spin-orbitals, eq. (3.20), being products of a spatial orbital and a spin function (α or β). If there are no restrictions on the form of the spatial orbitals, the trial function is an *Unrestricted Hartree-Fock* (UHF) wave function.¹¹ The term *different orbitals for different spins* (DODS) is also sometimes used. If the interest is in systems with an even number of electrons and a singlet type of wave function (a *closed shell* system), the restriction that each spatial orbital should have two electrons, one with α and one with β spin, is normally made. Such wave functions are known as *Restricted Hartree-Fock* (RHF). Open-shell systems may also be described by restricted type wave functions, where the spatial part of the doubly occupied orbitals is forced to be the same, and this is known as *Restricted Open-shell Hartree-Fock* (ROHF). For open-shell species, a UHF treatment leads to well-defined orbital energies, which may be interpreted as ionization potentials (Section 3.4). For an ROHF wave function, it is not possible to choose a unitary transformation that makes the matrix of Lagrange multipliers in eq. (3.41) diagonal, and orbital energies from an ROHF wave function are consequently not uniquely defined and cannot be equated to ionization potentials by a Koopmans-type argument.




Figure 3.4 Illustrating an RHF singlet, and ROHF and UHF doublet states

The UHF wave function allows different spatial orbitals for the two electrons in an orbital. As restricted type wave functions put constraints on the variation parameters, the energy of a UHF wave function is always lower than or equal to a corresponding R(O)HF type wave function. For singlet states near the equilibrium geometry, it is usually not possible to lower the energy by allowing the α and β MOs to be different. For an open-shell system such as a doublet, however, it is clear that forcing the α and β MOs to be identical *is* a restriction. If the unpaired electron has α spin, it will interact differently with the other α electrons than with the β electrons, and consequently the optimum α and β orbitals will be different. The UHF description, however, has the disadvantage that the wave function is not an eigenfunction of the \mathbf{S}^2 operator (unless it is equal to the RHF solution), where the \mathbf{S}^2 operator evaluates the value of the total electron spin squared. This means that a “singlet” UHF wave function may also contain contributions from higher lying triplet, quintet, etc., states. Similarly, a “doublet” UHF wave function will contain spurious (non-physical) contributions from higher lying quartet, sextet, etc., states. This will be discussed in more detail in Section 4.4.

Semi-empirical methods (Section 3.10) sometimes employ the so-called *half-electron* method for describing open-shell systems, such as doublets and triplets. In this model a doublet state is described by putting two “half” electrons in the same orbitals with opposite spins, i.e. constructing an RHF type wave functions where all electron spins are paired. A triplet state may similarly be modelled as having two orbitals, each occupied by two half electrons with opposite spin. The main motivation behind this artificial construct is that open- and closed shell systems (such as a triplet and singlet state) will have different amounts of electron correlation. Since semi-empirical methods perform the parameterization based on single-determinant wave functions, the half-electron method cancels the difference in electron correlations, and allows open- and closed shell systems to be treated on an equal footing in terms of energy. It has the disadvantage that the open-shell nature is no longer present in the wave function; it is for example not possible to calculate spin densities (i.e. where the unpaired electron(s) is(are) most likely to be).

3.8 SCF Techniques

As discussed in Section 3.6, the Roothaan–Hall (or Pople–Nesbet for the UHF case) equations must be solved iteratively since the Fock matrix depends on its own solutions. The procedure illustrated in Figure 3.3 involves the following steps:

- (1) Calculate all one- and two-electron integrals.
- (2) Generate a suitable start guess for the MO coefficients.
- (3) Form the initial density matrix.
- (4) Form the Fock matrix as the core (one-electron) integrals + the density matrix times the two-electron integrals.
- (5) Diagonalize the Fock matrix. The eigenvectors contain the new MO coefficients.
- (6) Form the new density matrix. If it is sufficiently close to the previous density matrix, we are done, otherwise go to step 4.

There are several points hidden in this scheme. Will the procedure actually converge at all? Will the SCF solution correspond to the desired energy minimum (and not a maximum or saddle point)? Can the number of iterations necessary for convergence

be reduced? Does the most efficient method depend on the type of computer and/or the size of the problem?

Let us look at some of the SCF techniques used in practice.

3.8.1 SCF convergence

There is *no* guarantee that the above iterative scheme will converge. For geometries near equilibrium and using small basis sets, the straightforward SCF procedure often converges unproblematically. Distorted geometries (such as transition structures) and large basis sets containing diffuse functions, however, rarely converge, and metal complexes, where several states with similar energies are possible, are even more troublesome. There are various tricks that can be tried to help convergence:¹²

- (1) *Extrapolation*. This is a method for trying to make the convergence faster by extrapolating previous Fock matrices to generate a (hopefully) better Fock matrix than the one calculated directly from the current density matrix. Typically, the last three matrices are used in the extrapolation.
- (2) *Damping*. The reason for divergence, or very slow convergence, is often due to oscillations. A given density matrix \mathbf{D}_n gives a Fock matrix \mathbf{F}_n , which, upon diagonalization, gives a density matrix \mathbf{D}_{n+1} . The Fock matrix \mathbf{F}_{n+1} from \mathbf{D}_{n+1} gives a density matrix \mathbf{D}_{n+2} that is close to \mathbf{D}_n , but \mathbf{D}_n and \mathbf{D}_{n+1} are very different, as illustrated in Figure 3.5. The damping procedure tries to solve this by replacing the current density matrix with a weighted average, $\mathbf{D}'_{n+1} = \omega \mathbf{D}_n + (1 - \omega) \mathbf{D}_{n+1}$. The weighting factor ω may be chosen as a constant or changed dynamically during the SCF procedure.




Figure 3.5 An oscillating SCF procedure

- (3) *Level shifting*. This technique¹³ is perhaps best understood in the formulation of a rotation of the MOs that form the basis for the Fock operator (Section 3.6). At convergence, the Fock matrix elements in the MO basis between occupied and virtual orbitals are zero. The iterative procedure involves mixing (making linear combinations of) occupied and virtual MOs. During the iterative procedure, these mixings may be large, causing oscillations or making the total energy increase. The degree of mixing may be reduced by artificially increasing the energy of the virtual orbitals. If a sufficiently large constant is added to the virtual orbital energies, it can be shown that the total energy is guaranteed to decrease, thereby forcing convergence. The more the virtual orbitals are raised in energy, the more stable is the

convergence, but the *rate* of convergence also decreases with level shifting. For large enough shifts, convergence is guaranteed, but it is likely to occur very slowly, and may in some cases converge to a state that is not the ground state.

- (4) *Direct inversion in the iterative subspace* (DIIS). This procedure was developed by P. Pulay and is an extrapolation procedure.¹⁴ It has proved to be very efficient in forcing convergence and in reducing the number of iterations at the same time, and it is now one of the most commonly used methods for helping SCF convergence. The idea is as follows. As the iterative procedure runs, a sequence of Fock and density matrices ($\mathbf{F}_0, \mathbf{F}_1, \mathbf{F}_2, \dots$ and $\mathbf{D}_0, \mathbf{D}_1, \mathbf{D}_2, \dots$) are produced. At each iteration, it is also assumed that an estimate of the “error” ($\mathbf{E}_0, \mathbf{E}_1, \mathbf{E}_2, \dots$) is available, i.e. how far the current Fock/density matrix is from the converged solution. The converged solution has an error of zero, and the DIIS method forms a linear combination of the error indicators that, in a least squares sense, is a minimum (as close to zero as possible). In the function space generated by the previous iterations we try to find the point with lowest error, which is not necessarily one of the points actually calculated. It is common to use the trace (sum of diagonal elements) of the matrix product of the error matrix with itself as a scalar indicator of the error.

$$\begin{aligned} \text{ErrF}(\mathbf{c}) &= \text{trace}(\mathbf{E}_{n+1} \cdot \mathbf{E}_{n+1}) \\ \mathbf{E}_{n+1} &= \sum_{i=0}^n c_i \mathbf{E}_i \\ \sum_{i=0}^n c_i &= 1 \end{aligned} \quad (3.64)$$

Minimization of the ErrF subject to the normalization constraint is handled by the Lagrange method (Section 12.5), and leads to the following set of linear equations, where λ is the multiplier associated with the normalization.

$$\begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} & -1 \\ a_{21} & a_{22} & \cdots & a_{2n} & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} & -1 \\ -1 & -1 & \cdots & -1 & 0 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \\ -\lambda \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ -1 \end{pmatrix} \quad (3.65)$$

$$a_{ij} = \text{trace}(\mathbf{E}_i \cdot \mathbf{E}_j)$$

$$\mathbf{A}\mathbf{c} = \mathbf{b}$$

$$\mathbf{c} = \mathbf{A}^{-1}\mathbf{b}$$

In iteration n the \mathbf{A} matrix has dimension $(n+1) \times (n+1)$, where n usually is less than 20. The coefficients \mathbf{c} can be obtained by directly inverting the \mathbf{A} matrix and multiplying it onto the \mathbf{b} vector, i.e. in the “subspace” of the “iterations” the linear equations are solved by “direct inversion”, thus the name DIIS. Having obtained the coefficients that minimize the error function at iteration n , the same set of coefficients is used for generating an extrapolated Fock matrix (\mathbf{F}^*) at iteration n , which is used in place of \mathbf{F}_n for generating the new density matrix.

$$\mathbf{F}_n^* = \sum_{i=0}^n c_i \mathbf{F}_i \quad (3.66)$$

The only remaining question is the nature of the error function. Pulay suggested the difference $\mathbf{FDS} - \mathbf{SDF}$ (\mathbf{S} is the overlap matrix), which is related to the gradient of the SCF energy with respect to the MO coefficients, and this has been found to work well in practice. A closely related method uses the energy as the error indicator, and has the acronym EDIIS.¹⁵

- (5) “*Direct minimization*” techniques. The variational principle indicates that we want to minimize the energy as a function of the MO coefficients or the corresponding density matrix elements, as given by eq. (3.54). In this formulation, the problem is no different from other types of non-linear optimizations, and the same types of technique, such as steepest descent, conjugated gradient or Newton–Raphson methods can be used (see Chapter 12 for details).

As mentioned in Section 3.6, the variational procedure can be formulated in terms of an exponential transformation of the MOs, with the (independent) variational parameters contained in an \mathbf{X} matrix. Note that the \mathbf{X} -variables are preferred over the MO coefficients in eq. (3.54) for optimization, since the latter are not independent (the MOs must be orthonormal). The exponential may be written as a series expansion, and the energy expanded in terms of the \mathbf{X} -variables describing the occupied–virtual mixing of the orbitals.¹⁶

$$\begin{aligned} e^{\mathbf{X}} &= 1 + \mathbf{X} + \frac{1}{2} \mathbf{X} \mathbf{X} + \dots \\ E(\mathbf{X}) &= E(\mathbf{0}) + E'(\mathbf{0})\mathbf{X} + \frac{1}{2} \mathbf{X} E''(\mathbf{0})\mathbf{X} + \dots \end{aligned} \quad (3.67)$$

The first and second derivatives of the energy with respect to the \mathbf{X} -variables ($E'(\mathbf{0})$ and $E''(\mathbf{0})$) can be written in terms of Fock matrix elements and two-electron integrals in the MO basis.¹⁷ For an RHF type wave function these are given in eq. (3.68).

$$\begin{aligned} \frac{\partial E}{\partial x_{ia}} &= 4 \langle \phi_i | \mathbf{F} | \phi_a \rangle \\ \frac{\partial^2 E}{\partial x_{ia} \partial x_{jb}} &= 4 \left[\delta_{ij} \langle \phi_a | \mathbf{F} | \phi_b \rangle - \delta_{ab} \langle \phi_i | \mathbf{F} | \phi_j \rangle + \right. \\ &\quad \left. \langle \phi_i \phi_b | \phi_a \phi_j \rangle - \langle \phi_i \phi_j | \phi_a \phi_b \rangle - \langle \phi_i \phi_a | \phi_j \phi_b \rangle \right] \end{aligned} \quad (3.68)$$

The gradient of the energy is an off-diagonal element of the molecular Fock matrix, which is easily calculated from the atomic Fock matrix. The second derivative, however, involves two-electron integrals that require an AO to MO transformation (see Section 4.2.1), and is therefore computationally expensive.

In a density matrix formulation, the energy depends on the density matrix elements as variables, and can formally be written as the trace of the contraction of the density matrix with the one-electron matrix \mathbf{h} and the two-electron matrix \mathbf{G} , with the latter depending implicitly on \mathbf{D} .

$$E(\mathbf{D}) = \text{trace}(\mathbf{D}\mathbf{h}) + \text{trace}(\mathbf{D}\mathbf{G}(\mathbf{D})) \quad (3.69)$$

The density matrix elements cannot be varied freely, however, as the orbitals must remain orthonormal, and this constraint can be formulated as the density matrix having to be *idempotent*, $\mathbf{DSD} = \mathbf{D}$. It is difficult to ensure this during an

optimization step, but the non-idempotent density matrix derived from taking an optimization step can be “purified” by the McWeeny procedure.¹⁸

$$\mathbf{D}_{\text{purified}} = 3\mathbf{D}^2 - 2\mathbf{D}^3 \quad (3.70)$$

The idempotency condition ensures that each orbital is occupied by exactly one electron. E. Cancès has shown that relaxing this condition to allow fractional occupancy during the optimization improves the convergence, a procedure named *relaxed constraint algorithm* (RCA)¹⁹ and which was subsequently improved using ideas from the DIIS algorithm, leading to the *EDIIS* (*Energy DIIS*) method.¹⁵ The optimization in terms of density matrix elements has the potential advantage that the matrix becomes sparse for large systems, and can therefore be solved by techniques that scale linearly with the system’s size.²⁰

The Newton–Raphson method has the advantage of being quadratically convergent, i.e. sufficiently near the minimum it converges very fast. The main problem in using Newton–Raphson methods for wave function optimization is computational efficiency. The exact calculation of the second derivative matrix is somewhat demanding, and each iteration in a Newton–Raphson optimization therefore takes longer than the simple Roothaan–Hall iterative scheme. Owing to the fast convergence near the minimum, a Newton–Raphson approach normally takes fewer iterations than for example DIIS, but the overall computational time is still a factor of ~ 2 longer. Alternative schemes, where an approximation to the second derivative matrix is used (pseudo-Newton–Raphson), have also been developed, and they are often competitive with DIIS.²¹ It should be kept in mind that the simple Newton–Raphson is unstable, and requires some form of stabilization, for example by using the augmented Hessian techniques discussed in Section 12.2.²² Alternatively, for large system (thousands of basis functions) the optimization may be carried out by conjugate gradient methods, but the convergence characteristic of these methods is significantly poorer.²³ Direct minimization methods have the advantage of a more stable convergence for difficult systems, where DIIS may display problematic behaviour or converge to solutions that are not the global minimum.

3.8.2 Use of symmetry

From group theory it may be shown that an integral can only be non-zero if the integrand belongs to the totally symmetric representation. Furthermore, the product of two functions can only be totally symmetric if they belong to the same irreducible representation. As both the Hamiltonian and Fock operators are totally symmetric (otherwise the energy would change by a rotation of the coordinate system), integrals of the following type can only be non-zero if the basis functions involving the same electron coordinate belong to the same representation.

$$\int \chi_\alpha(1)\chi_\beta(1)d\mathbf{r}_1 \quad \int \chi_\alpha(1)\mathbf{F}\chi_\beta(1)d\mathbf{r}_1 \quad \int \chi_\alpha(1)\mathbf{H}\chi_\beta(1)d\mathbf{r}_1 \quad (3.71)$$

Similar considerations hold for the two-electron integrals.

By forming suitable linear combinations of basis functions (*symmetry-adapted* functions), many one- and two-electron integrals need not be calculated as they are known to be exactly zero owing to symmetry. Furthermore, the Fock (in an HF calculation)

or Hamiltonian matrix (in a configuration interaction (CI) calculation) will become block-diagonal, as only matrix elements between functions having the same symmetry can be non-zero. The saving depends on the specific system, but as a guideline the computational time is reduced by roughly a factor corresponding to the order of the point group (number of symmetry operations). Although the large majority of molecules do not have any symmetry, a sizeable proportion of the small molecules for which *ab initio* electronic structure calculations are possible are symmetric. Almost all *ab initio* programs employ symmetry as a tool for reducing the computational effort.

3.8.3 Ensuring that the HF energy is a minimum, and the correct minimum

The standard iterative procedure produces a solution where the variation of the HF energy is *stationary* with respect to all orbital variations, i.e. the first derivatives of the energy with respect to the MO coefficients are zero. In order to ensure that this corresponds to an energy minimum, the second derivatives should also be calculated.²⁴ This is a matrix the size of the number of occupied MOs multiplied by the number of virtual MOs (identical to that arising in quadratic convergent SCF methods (Section 3.8.1)), and the eigenvalues of this matrix should all be positive in order to be an energy minimum. Of course only the lowest eigenvalue is required to probe whether the solution is a minimum. A negative eigenvalue means that it is possible to get to a lower energy state by “exciting” an electron from an occupied to an unoccupied orbital, i.e. the solution is unstable. In practice, the stability is rarely checked – it is assumed that the iterative procedure has converged to a minimum. It should be noted that a positive definite second-order matrix only ensures that the solution is a *local* minimum; there may be other minima with lower energies.

The problem of convergence to saddle points in the wave function parameter space and the existence of multiple minima is rarely a problem for systems composed of elements from the first two rows in the periodic table. For systems having more than one metal atom with several partially filled d-orbitals, however, care must be taken to ensure that the iterative procedure converges to the desired solution. Consider for example the Fe_2S_2 system in Figure 3.6, where the d-electrons of two Fe atoms are coupled through the sulfur bridge atoms.




Figure 3.6 Two different singlet states generated by coupling either two high-spin or two low-spin states

Each of the two Fe atoms is formally in the +III oxidation state, and therefore has a d^5 configuration. A high-spin state corresponding to all the ten d-electrons being aligned can readily be described by a singledeterminant wave function, but the situation is more complicated for a low-spin singlet state. A singlet HF wave function must have an equal number of orbitals with α and β electron spin, but this can be obtained in several different ways. If each metal atom is in a high-spin state, an overall singlet state must have all the d-orbitals on one Fe atom occupied by electrons with α spin, while all the d-orbitals on the other Fe atom must be occupied by electrons with β spin. An alternative singlet state, however, can be generated by coupling the single unpaired electron from the two Fe centres in a low-spin configuration. Each of these two wave functions will be valid minima in the orbital parameter space, but clearly describe complexes with different properties. Note also that neither of these two singlet wave functions can be described by an RHF type wave function. UHF type wave functions with the above two types of spin coupling can be generated, but will often be severely spin contaminated. One can consider other spin coupling schemes to generate an overall singlet wave function, and the situation becomes more complicated if intermediate (triplet, pentet, etc.) spin states are desired, and for mixed valence states ($\text{Fe}^{2+}/\text{Fe}^{3+}$). The complications further increase when larger clusters are considered, as for example with the Fe_4S_4 moiety involved in electron transfer in the photosystem I and nitrogenase enzymes.

The question as to whether the energy is a minimum is closely related to the concept of wave function stability. If a lower energy RHF solution can be found, the wave function is said to possess a *singlet instability*. It is also possible that an RHF type wave function is a minimum in the coefficient space, but is a saddle point if the constraint of double occupancy of each MO is relaxed. This indicates that a lower energy wave function of the UHF type can be constructed, and this is called a *triplet instability*. It should be noted that in order to generate such UHF wave functions for a singlet state, an initial guess of the SCF coefficients must be specified that has the spatial parts of at least one set of α and β MOs different. There are other types of such instabilities, such as relaxing the constraint that the MOs should be real (allowing complex orbitals), or the constraint that a MO should only have a single spin function. Relaxing the latter produces the “general” HF method, where each MO is written as a spatial part having α spin plus another spatial part having β spin.²⁵ Such wave functions are no longer eigenfunctions of the \hat{S}_z operator, and are rarely used.

Another aspect of wave function instability concerns *symmetry breaking*, i.e. the wave function has a lower symmetry than the nuclear framework.²⁶ It occurs, for example, with the allyl radical with an ROHF type wave function. The nuclear geometry has C_{2v} symmetry, but the C_{2v} symmetric wave function corresponds to a (first-order) saddle point. The lowest energy ROHF solution has only C_s symmetry, and corresponds to a localized double bond and a localized electron (radical). Relaxing the double occupancy constraint, and allowing the wave function to become UHF, re-establishes the correct C_{2v} symmetry. Such symmetry breaking phenomena usually indicate that the type of wave function used is not flexible enough for even a qualitatively correct description.

3.8.4 Initial guess orbitals

The quality of the initial guess orbitals influences the number of iterations necessary for achieving convergence. As each iteration involves a computational effort proportional to M_{basis}^4 , it is of course desirable to generate as good a guess as possible. Different start orbitals may in some cases result in convergence to different SCF solutions, or make the difference between convergence and divergence. One possible way of generating a set of start orbitals is to diagonalize the Fock matrix consisting only of the one-electron contributions, the “core” matrix. This corresponds to initializing the density matrix as a zero matrix, totally neglecting the electron–electron repulsion in the first step. This is generally a poor guess, but it is available for all types of basis set and is easily implemented. Essentially all programs therefore have it as an option.

More sophisticated procedures involve taking the start MO coefficients from a semi-empirical calculation, such as Extended Hückel Theory (EHT) or Intermediate Neglect of Differential Overlap (INDO) (Sections 3.13 and 3.10). The EHT method has the advantage that it is readily parameterized for all elements, and it can provide start orbitals for systems involving elements from essentially the whole periodic table. An INDO calculation normally provides better start orbitals, but at a price. The INDO calculation itself is iterative, and it may suffer from convergence problems, just as the *ab initio* SCF itself.

Many systems of interest are symmetric. The MOs will transform as one of the irreducible representations in the point group, and most programs use this to speed up calculations. The initial guess for the start orbitals involves selecting how many MOs of each symmetry should be occupied, i.e. the electron configuration. Different start configurations produce different final SCF solutions. Many programs automatically select the start configuration based on the orbital energies of the starting MOs, which may be “wrong” in the sense that it does not produce the desired solution. Of course, a given solution may be checked to see if it actually corresponds to an energy minimum, but as stated above, this is rarely done. Furthermore, there may be several (local) minima, thus the verification that the found solution is an energy minimum is no guarantee that it is the global minimum. A particular case is open-shell systems having at least one element of symmetry, as the open-shell orbital(s) determine the overall wave function symmetry. An example is the N_2^+ radical cation, where two states of Σ_g and Π_u symmetry exist with a difference of only ~ 70 kJ/mol in energy.

The reason different initial electron configurations may generate different final solutions is because matrix elements between orbitals belonging to different representations are exactly zero, thus only orbitals belonging to the same representation can mix. Forcing the program to run the calculation without symmetry usually does not help. Although turning the symmetry off will make the program actually calculate all matrix elements, those between MOs of different symmetry will still be zero (except for numerical inaccuracies). It is therefore often necessary to specify manually which orbitals should be occupied initially to generate the desired solution.

3.8.5 Direct SCF

The number of two-electron integrals formally grows as the fourth power of the size of the basis set. Owing to permutation symmetry (the following integrals are identical $\langle \chi_1 \chi_2 | \chi_3 \chi_4 \rangle = \langle \chi_3 \chi_2 | \chi_1 \chi_4 \rangle = \langle \chi_1 \chi_4 | \chi_3 \chi_2 \rangle = \langle \chi_3 \chi_4 | \chi_1 \chi_2 \rangle = \langle \chi_2 \chi_1 | \chi_4 \chi_3 \rangle = \langle \chi_4 \chi_1 | \chi_2 \chi_3 \rangle = \langle \chi_2 \chi_3 | \chi_4 \chi_1 \rangle = \langle \chi_4 \chi_3 | \chi_2 \chi_1 \rangle$) the total number is approximately $1/8 M_{\text{basis}}^4$. Each integral is a floating point number associated with four indices indicating which basis functions are involved in the integral. Storing a floating point number in double precision (which is necessary for calculating the energy with an accuracy of ~ 14 digits) requires 64 bits = 8 bytes. A basis set with 100 functions thus generates $\sim 12 \times 10^6$ integrals, requiring ~ 100 Mbytes of disk space or memory. The disk space required for storing the integrals rises rapidly, thus a basis set with 1000 functions requires ~ 1000 Gbytes of disk space (or memory). This is out of reach for most computers. In practice, the storage requirement is somewhat less, since many of the integrals are small, and can be ignored. Typically, a cutoff around $\sim 10^{-10}$ is employed: if the integral is less than this value it is not stored, and consequently makes a zero contribution to the construction of the Fock matrix in the iterative procedure. However, the disk space requirement effectively limits conventional HF methods to basis sets smaller than ~ 500 functions.

Older computers had only very limited amounts of memory, and disk storage of the integrals was the only option. Modern machines often have quite significant amounts of memory – a few hundred Gbytes is not uncommon. For small- and medium-sized systems, it may be possible to store all the integrals in memory instead of on disk. Such “in-core” methods are very efficient for performing an HF calculation. The integrals are only calculated once, and each SCF iteration is just a multiplication of the integral tensor with a density matrix to form the Fock matrix (eq. (3.53)). Essentially all machines have optimized routines for doing matrix multiplication efficiently. The only limitation is the quartic (M_{basis}^4) growth of the memory requirement with basis set size, which in practice restricts such in-core methods to basis sets with less than ~ 200 functions.

The disk space (or memory) requirement can be reduced dramatically by performing the SCF in a *direct* fashion.²⁷ In the direct SCF method, the integrals are calculated from scratch in each iteration. At first this would appear to involve a computational effort that is larger than a conventional HF calculation by a factor close to the number of iterations. There are, however, a number of considerations that often makes direct SCF methods computationally quite competitive or even advantageous.

In disk-based methods, all the integrals are first calculated and written to disk. To reduce the disk space requirement, the four indices associated with each integral are “packed” into a single number, and written to disk. The whole set of integrals must be read in each iteration, and the indices “unpacked” before the integrals are multiplied with the proper density matrix elements and added to the Fock matrix. Typically, half the time in an SCF procedure is spent calculating the integrals and writing them to disk, the other half is spent reading, unpacking and forming the Fock matrix maybe 20 times. In a direct approach, there is no overhead due to packing/unpacking of indices, or writing/reading of integrals.

In disk-based methods, only integrals larger than a certain cutoff are saved. In direct methods, it is possible to ignore additional integrals. The contribution to a Fock matrix element is a product of density matrix elements and two-electron integrals. In disk-

based methods, the density matrix is not known when the integrals are calculated, and all integrals above the cutoff must be saved and processed in each iteration. In direct methods, however, the density matrix is known at the time when the integrals are calculated. Thus if the *product* of the density matrix elements and the integral is less than the cutoff, the integral can be ignored. Of course, this is only a saving if an estimate of the size of the integral is available before it is actually calculated. One such estimate is the *Schwarz inequality* (eq. (3.72)), but more advanced screening methods have also been developed.²⁸

$$|\langle\alpha\gamma|\beta\delta\rangle| \leq \sqrt{\langle\alpha\alpha|\beta\beta\rangle} \cdot \sqrt{\langle\gamma\gamma|\delta\delta\rangle} \quad (3.72)$$

The number of two-centre integrals on the right-hand side is quite small (of the order of M_{basis}^2) and can easily be calculated beforehand. Thus if the product of the density matrix elements and the upper limit of the integral is less than the cutoff, the integral does not need to be calculated. In practice, integrals are calculated in *batches*, where a batch is a collection of integrals having the same exponent. For a $\langle pp|pp\rangle$ type batch there are thus 81 individual integrals, a $\langle dd|dd\rangle$ type batch has 625 individual integrals, etc. The integral screening is normally done at the batch level, i.e. if the largest term is smaller than a given cutoff, the whole batch can be neglected.

The above integral screening is even more advantageous if the Fock matrix is formed incrementally. Consider two sequential density and Fock matrices in the iterative procedure (eq. (3.53)).

$$\begin{aligned} \mathbf{F}_n &= \mathbf{h} + \mathbf{G} \cdot \mathbf{D}_n \\ \mathbf{F}_{n+1} &= \mathbf{h} + \mathbf{G} \cdot \mathbf{D}_{n+1} \\ \mathbf{F}_{n+1} - \mathbf{F}_n &= \mathbf{G} \cdot (\mathbf{D}_{n+1} - \mathbf{D}_n) \\ \Delta \mathbf{F}_{n+1} &= \mathbf{G} \cdot \Delta \mathbf{D}_{n+1} \end{aligned} \quad (3.73)$$

The *change* in the Fock matrix depends only on the *change* in the density matrix. Combined with the above screening procedure, it is thus only necessary to calculate those integrals to be multiplied with density matrix elements that have changed significantly since the last iteration. As the SCF converges, there are fewer and fewer integrals that need to be calculated.

The formal scaling of HF methods is M_{basis}^4 , since the total number of two-electron integrals increases as M_{basis}^4 . As just seen, however, we do not need to calculate all the two-electron integrals – many can be neglected without affecting the final results. The observed scaling is therefore less than the quartic dependence, but the exact power depends on how the size of the problem is increased. If the number of atoms is increased for a fixed basis set per atom, the scaling depends on the dimensionality of the atomic arrangement and the size of the atomic basis. The most favourable case is a small compact basis set (such as a minimum basis) and an essential one-dimensional system, such as polyacetylene, $\text{H}-(\text{C}\equiv\text{C})_n-\text{H}$, or linear alkanes. In this case, the scaling is close to M_{basis}^2 once the number of functions exceeds ~ 100 . A two-dimensional arrangement of atoms (such as a slab of graphite) has a slightly larger exponent dependence, while a three-dimensional system (such as a diamond structure) has a power dependence close to $M_{\text{basis}}^{2.3}$.²⁹ It should be noted that most molecular systems have a dimensionality between two and three – the presence of “holes” in the structure reduces the effective dimensionality to below three. With a larger basis set,

especially if diffuse functions are present, the screening of integrals becomes much less efficient, or equivalently, the molecular system must be significantly larger to achieve the limiting scaling. In practice, however, the increase in the total number of basis functions is often not due to an enlargement of the molecular system, but rather to the use of an increasingly larger basis set per atom for a fixed sized molecule. For such cases, the observed scaling is often *worse* than the theoretical M_{basis}^4 dependence, since the integral screening becomes less and less efficient.

The combination of these effects means that the increase in computational time for a direct SCF calculation compared with a disk-based method is less than initially expected. For a medium-sized SCF calculation that requires say 20 iterations, the increase in CPU time may only be a factor of 2 or 3. Due to the more efficient screening, however, the direct method actually becomes more and more advantageous relative to disk-based methods as the size of the system increases. At some point, direct methods will therefore require *less* CPU time than a conventional method. Exactly where the cross-over point occurs depends on the way the number of basis functions is increased, the machine type and the efficiency of the integral code. Small compact basis sets in general experience the cross-over point quite early (perhaps around 100 functions) while it occurs later for large extended basis sets. Since conventional disk-based methods are limited to 200–300 basis functions, direct methods are normally the only choice for large calculations. Direct methods are essentially only limited by the available CPU time, and calculations involving up to several thousand basis functions have been reported.

Although direct methods for small- and medium-size systems require more CPU time than disk-based methods, this is in many cases irrelevant. For the user the determining factor is the time from submitting the calculation to the results being available. Over the years the speed of CPUs has increased much more rapidly than the speed of data transfer to and from disk. Most modern machines have very slow data transfer to disk compared with CPU speed. Measured by the elapsed wall clock time, disk-based HF methods are often the slowest in delivering the results, despite the fact that they require the least CPU time. Simply speaking, the CPU may be spending most of its time waiting for data to be transferred from disk. Direct methods, on the other hand, use the CPU with a near 100% efficiency. For machines without fast disk transfer (such as workstation-type machines) the cross-over point for direct versus conventional methods in terms of wall clock time may be so low that direct methods are always preferred.

Finally, it should be mentioned that there is a strong research effort towards designing computational chemistry programs to run on parallel computers. These types of machines have more than one CPU, typically in the range 10–1000. Making direct SCF calculations run efficiently in a parallel fashion is fairly easy: each processor is given the task of calculating a certain batch of integrals and the total Fock matrix is simply the sum of contributions from each individual CPU.

3.8.6 *Reduced scaling techniques*

The computational bottleneck in HF methods is the calculation of the two-electron Coulomb and exchange terms arising from the electron–electron repulsion. In non-metallic systems, the exchange term is quite short-ranged, while the Coulomb interac-

tion is long-ranged. In the large system limit, the Coulomb integrals thus dominate the computational cost. By using the screening techniques described in the previous section, the scaling in the large system limit will eventually be reduced from M_{basis}^4 to M_{basis}^2 . Similar considerations hold for DFT methods (Chapter 6). Although an M_{basis}^2 scaling is quite modest, it is clear that a reduction down to linear scaling will be advantageous in order to move the calculations into the thousand atoms regime.³⁰

The *Fast Multipole Moment* (FMM) method (Section 14.3) was originally developed for calculating interactions between point charges. A direct calculation involves a summation over all pairs, i.e. a computational effort that increases with M_{basis}^2 . The idea in FMM is to split the total interaction into a near- and a far-field. The near-field is evaluated directly, while the far-field is calculated by dividing the physical space into boxes, and the interaction between all the charges in one box and all the charges in another is approximated as interactions between multipoles located at the centre of the boxes. The further away from each other two boxes are, the larger the boxes can be for a given accuracy, thereby reducing the formal M_{basis}^2 behaviour to linear scaling, i.e. proportional to M_{basis} .

The original FMM has been refined by also adjusting the accuracy of the multipole expansion as a function of the distance between boxes, producing the *very Fast Multipole Moment* (vFMM) method.³¹ Both of these have been generalized to continuous charge distributions, as is required for calculating the Coulomb interaction between electrons in a quantum description.³² The use of FMM methods in electronic structure calculations enables the Coulomb part of the electron–electron interaction to be calculated with a computational effort that depends linearly on the number of basis functions, once the system becomes sufficiently large.

Instead of dividing the physical space into a near- and far-field, the Coulomb operator itself may be partitioned into a short- and long-ranged part.³³ The short-ranged operator is evaluated exactly, while the long-ranged part is evaluated for example by means of a Fourier transformation. The net effect is again that the total Coulomb interaction can be calculated with a computational effort that only scales linearly with system size.

Although the exchange term in principle is short-ranged, and thus should benefit significantly from integral screening, this is normally not observed in practical calculations. This has been attributed to basis set incompleteness,³⁴ and this insight allowed a formulation of a more aggressive screening technique that enables the exchange part of the electron–electron interaction also to be reduced to an order M_{basis} method.

Another approach for achieving linear scaling is to break the system into smaller parts, perform a calculation on each subsection, and subsequently piece these results together.³⁵ This *divide and conquer* strategy relies on the nearsightedness of molecular systems, i.e. the local electronic wave function is rarely sensitive to molecular features further away than 15–20 Å. There are of course exceptions, such as large conjugated systems, or long-range charge transfer, which are problematic to treat by these methods.

The use of methods with a reduced scaling does not necessarily lead to a reduced computational cost for systems that can be studied by the available resources. The cross-over point for when the linear scaling methods becomes competitive with traditional methods may be so high that is it of little practical use. At present, there is little

available data, but test calculations indicate that the cross-over point involves several thousand basis functions, at least for large systems with modest basis sets.

An alternative approach is to accept a relatively large formal scaling, and focus on reducing the prefactor in the computationally expensive step. Methods relying on the so-called resolution of the identity for splitting the calculation of four-index integrals into three- and two-index quantities belong to this class. At the HF level, the formal scaling is (only) reduced from M_{basis}^4 to M_{basis}^3 , but actual timings show that the total computational cost is reduced by roughly an order of magnitude, without compromising the accuracy.³⁶ Furthermore, the efficiency gain increases with the basis set size, i.e. these methods become very favourable even for small systems when large basis sets are used for achieving high accuracy.

Since the HF method may be formulated and implemented in several different way, a practical question is which of these methods will be the fastest computationally for a given problem. The scaling only determines which method will be the fastest for large systems, i.e. for $N \rightarrow \infty$. An equally important parameter is the prefactor, i.e. the proportionality constant between computational time and system size. A method with a favourable scaling will often have a larger prefactor than a method with a more demanding scaling behaviour. Figure 3.7 illustrates a quartic, quadratic and linear scaling algorithm with different prefactors.




Figure 3.7 Method scaling with system size

For systems smaller than N_1 , the most efficient method is the quartic one, the quadratic algorithm is the most efficient for systems sizes between N_1 and N_2 , while the linear scaling method becomes the most efficient beyond N_2 . Note that N_2 may be so large that the total computational resources may be exhausted before the cross-over point is reached.

With the advent of methods that enable the construction of the Fock matrix to be done with a computational effort that scales linearly with systems size, the

diagonalization step for solving the HF equations eventually becomes the computational bottleneck, since matrix diagonalization depends on the third power of the problem size, and this cross-over occurs for a few thousand basis functions. As discussed in Section 3.8.1, however, it is possible to reformulate the SCF problem in terms of a minimization of an energy functional that depends directly on the density matrix elements or orbital rotation parameters. This functional can then be minimized for example by conjugate gradient methods (Section 12.2.2), taking advantage of the fact that the density matrix becomes sparse for large systems. The HF method therefore appears to have reached the “holy grail” of quantum chemistry, i.e. linear scaling with system size.

3.9 Periodic Systems

Periodic systems can be described as a fundamental *unit cell* being repeated to form an infinite system. The periodicity can be in one dimension (e.g. a polymer), two dimensions (e.g. a surface) or three dimensions (e.g. a crystal), with the latter being the most common. The unit cell in three dimensions can be characterized by three vectors \mathbf{a}_1 , \mathbf{a}_2 and \mathbf{a}_3 spanning the physical space, with the length and the angles between them defining the shape.³⁷ There are seven possible shapes, the simplest of which is cubic, where all vector lengths are equal and all angles are 90° .




Figure 3.8 A cubic unit cell defined by three vectors

A unit cell can have atoms (or molecules) occupying various positions within the cell (corners, sides, centre), and the combination of a unit cell and its occupancy is called a *Bravais lattice*, of which there are fourteen possible forms. The periodic (infinite) system can then be generated by translation of the unit cell (Bravais lattice) by lattice vectors \mathbf{t} .

The *reciprocal cell* is defined by three vectors \mathbf{b}_1 , \mathbf{b}_2 and \mathbf{b}_3 derived from the \mathbf{a}_1 , \mathbf{a}_2 and \mathbf{a}_3 vectors of the direct cell, and obeying the orthonormality condition $\mathbf{a}_i \cdot \mathbf{b}_j = 2\pi \delta_{ij}$.

$$\mathbf{b}_1 = 2\pi \frac{\mathbf{a}_2 \times \mathbf{a}_3}{L^3} \quad \mathbf{b}_2 = 2\pi \frac{\mathbf{a}_3 \times \mathbf{a}_1}{L^3} \quad \mathbf{b}_3 = 2\pi \frac{\mathbf{a}_1 \times \mathbf{a}_2}{L^3} \quad (3.74)$$

The reciprocal cell of a cubic cell with side length L is also a cube, with the side length $2\pi/L$. The equivalent of a unit cell in reciprocal space is called the (first) *Brillouin zone*. Just as a point in real space may be described by a vector \mathbf{r} , a “point” in reciprocal space may be described by a vector \mathbf{k} . Since \mathbf{k} has units of inverse length, it is often called a wave vector. It is also closely related to the momentum and energy, e.g. the momentum and kinetic energy of a (free) particle described by a plane wave of the form $e^{i\mathbf{k} \cdot \mathbf{r}}$ is $\hbar \mathbf{k}$ and $\hbar^2 k^2 / 2m$, respectively.

The periodicity of the nuclei in the system means that the square of the wave function must display the same periodicity. This is inherent in the *Bloch theorem* (eq. (3.75)), which states that the wave function value at equivalent positions in different cells are related by a complex phase factor involving the lattice vector \mathbf{t} and a vector in the reciprocal space.

$$\phi(\mathbf{r} + \mathbf{t}) = e^{i\mathbf{k} \cdot \mathbf{t}} \phi(\mathbf{r}) \quad (3.75)$$

Alternatively stated, the Bloch theorem indicates that a *crystalline orbital* (ϕ) for the n th band in the unit cell can be written as a wave-like part and a cell-periodic part (φ), called a *Bloch orbital*.

$$\phi_{n,k}(\mathbf{r}) = e^{i\mathbf{k} \cdot \mathbf{r}} \varphi_n(\mathbf{r}) \quad (3.76)$$

The Bloch orbital can be expanded into a basis set of plane wave functions (χ^{PW}).

$$\begin{aligned} \varphi_n(\mathbf{r}) &= \sum_{\alpha}^{M_{\text{basis}}} c_{n\alpha} \chi_{\alpha}^{\text{PW}}(\mathbf{r}) \\ \varphi_{n,k}(\mathbf{r}) &= e^{i\mathbf{k} \cdot \mathbf{r}} \sum_{\alpha}^{M_{\text{basis}}} c_{n\alpha} \chi_{\alpha}^{\text{PW}}(\mathbf{r}) \end{aligned} \quad (3.77)$$

Alternatively, the basis set can be chosen as a set of nuclear-centred (Gaussian) basis functions, from which a set of Bloch orbitals can be constructed.

$$\begin{aligned} \varphi_{k\alpha}(\mathbf{r}) &= \sum_{\mathbf{t}} e^{i\mathbf{k} \cdot \mathbf{t}} \chi_{\alpha}^{\text{GTO}}(\mathbf{r} + \mathbf{t}) \\ \varphi_{n,k}(\mathbf{r}) &= \sum_{\alpha}^{M_{\text{basis}}} c_{n\alpha} \varphi_{k\alpha}(\mathbf{r}) = \sum_{\alpha}^{M_{\text{basis}}} \sum_{\mathbf{t}} c_{n\alpha} e^{i\mathbf{k} \cdot \mathbf{t}} \chi_{\alpha}^{\text{GTO}}(\mathbf{r} + \mathbf{t}) \end{aligned} \quad (3.78)$$

The problem has now been transformed from treating an infinite number of orbitals (electrons) to only treating those within the unit cell. The price is that the solutions become a function of the reciprocal space vector \mathbf{k} within the first Brillouin zone. For a system with M_{basis} functions, the variation problem can be formulated as a matrix equation analogous to eq. (3.51).

$$\mathbf{F}^{\mathbf{k}} \mathbf{C}^{\mathbf{k}} = \mathbf{S}^{\mathbf{k}} \mathbf{C}^{\mathbf{k}} \mathbf{e}^{\mathbf{k}} \quad (3.79)$$

The \mathbf{k} appears as a parameter in the equation similarly to the nuclear positions in molecular Hartree–Fock theory. The solutions are continuous as a function of \mathbf{k} , and provide a range of energies called a *band*, with the total energy per unit cell being calculated by integrating over \mathbf{k} space. Fortunately, the variation with \mathbf{k} is rather slow for non-metallic systems, and the integration can be done numerically by including relatively few points.³⁸ Note that the presence of the phase factors in eq. (3.76) means that the matrices in eq. (3.79) are complex quantities.

For a given value of \mathbf{k} , the solution of eq. (3.79) provides M_{basis} orbitals. In molecular systems, the molecular orbitals are filled with electrons according to the *aufbau* principle, i.e. according to energy. The same principle is used for periodic systems, and the equivalent of the molecular HOMO (highest occupied molecular orbital) is the *Fermi* energy level. Depending on the system, two situations can occur.

- The number of electrons is such that a certain number of (non-overlapping) bands are completely filled, while the rest are empty.
- The number of electrons is such that one (or more) band(s) are only partly filled.

The first situation is analogous to that for molecular systems having a closed-shell singlet state. The energy difference between the “top” of the highest filled band and the “bottom” of the lowest empty band is called the *band gap*, and is equivalent to the HOMO–LUMO gap in molecular systems. The second situation is analogous to an open-shell electronic structure for a molecular system, and corresponds to a band gap of zero. Systems with a band gap of zero are metallic, while those with a finite band gap are either insulators or semiconductors, depending on whether the band gap is large or small compared with the thermal energy kT .

As mentioned above, the basis functions within a unit cell can be either localized (Gaussian) or delocalized (plane wave) functions. For a Gaussian basis set, the computational problem of constructing the \mathbf{F}^k matrix is closely related to the molecular cases, involving multi-dimensional integrals over kinetic and potential energy operators. The periodic boundary condition means that the terms involving the potential energy operators in eq. (3.79) become infinite sums over \mathbf{t} vectors. Since the operators involve both positive and negative quantities and only decay as r^{-1} , they require special care to ensure convergence to a definite quantity, as for example Ewald sum³⁹ or fast multipole methods.⁴⁰ For a plane wave basis, the construction of the energy matrix can be done efficiently by using fast Fourier transform (FFT) methods for switching between the real and reciprocal space. All local potential operators are easily evaluated in real space, while the kinetic energy is just the square operator in reciprocal space. FFT methods have the big advantage that the computational cost only increases as $N \ln N$, with N being the number of grid points in the Fourier transform.

The solution of eq. (3.79) can be done by repeated diagonalization of the \mathbf{F}^k matrix, analogously to the situation for non-periodic systems. A plane wave basis, however, often involves several thousand functions, which means that alternative methods are used for solving the equation.

3.10 Semi-Empirical Methods

The cost of performing an HF calculation scales formally as the fourth power of the number of basis functions. This arises from the number of two-electron integrals necessary for constructing the Fock matrix. Semi-empirical methods reduce the computational cost by reducing the number of these integrals.⁴¹ Although linear scaling methods can reduce the scaling of *ab initio* HF methods to $\sim M_{\text{basis}}$, this is only the limiting behaviour in the large basis set limit, and *ab initio* methods will still require a significantly larger computational effort than semi-empirical methods.

The first step in reducing the computational problem is to consider only the valence electrons explicitly; the core electrons are accounted for by reducing the nuclear charge or introducing functions to model the combined repulsion due to the nuclei and core electrons. Furthermore, only a minimum basis set (the minimum number of functions necessary for accommodating the electrons in the neutral atom) is used for the valence electrons. Hydrogen thus has one basis function, and all atoms in the second and third

rows of the periodic table have four basis functions (one s - and one set of p -orbitals, p_x , p_y and p_z). The large majority of semi-empirical methods to date use only s - and p -functions, and the basis functions are taken to be Slater type orbitals (see Chapter 5), i.e. exponential functions.

The central assumption of semi-empirical methods is the *Zero Differential Overlap* (ZDO) approximation, which neglects all products of basis functions that depend on the same electron coordinates when located on different atoms. Denoting an atomic orbital on centre A as μ_A (it is customary to denote basis functions with μ , ν , λ and σ in semi-empirical theory, while we are using χ_α , χ_β , χ_γ and χ_δ for *ab initio* methods), the ZDO approximation corresponds to $\mu_A \nu_B = 0$. Note that it is the *product* of functions on different atoms that is set equal to zero, not the *integral* over such a product. This has the following consequences (eqs (3.51 and (3.56)):

- (1) The overlap matrix **S** is reduced to a unit matrix.
- (2) One-electron integrals involving three centres (two from the basis functions and one from the operator) are set to zero.
- (3) All three- and four-centre two-electron integrals, which are by far the most numerous of the two-electron integrals, are neglected.

To compensate for these approximations, the remaining integrals are made into parameters, and their values are assigned based on calculations or experimental data. Exactly how many integrals are neglected, and how the parameterization is done, defines the various semi-empirical methods.

Rewriting eq. (3.52) with semi-empirical labels gives the following expression for a Fock matrix element, where a two-electron integral is abbreviated as $\langle \mu\nu | \lambda\sigma \rangle$ (eq. (3.57)).

$$F_{\mu\nu} = h_{\mu\nu} + \sum_{\lambda\sigma}^{M_{\text{basis}}} D_{\lambda\sigma} (\langle \mu\nu | \lambda\sigma \rangle - \langle \mu\lambda | \nu\sigma \rangle) \quad (3.80)$$

$$h_{\mu\nu} = \langle \mu | \mathbf{h} | \nu \rangle$$

Approximations are made for the one- and two-electron parts as follows.

3.10.1 Neglect of Diatomic Differential Overlap Approximation (NDDO)

In the *Neglect of Diatomic Differential Overlap* (NDDO) approximation there are no further approximations than those mentioned above. Using μ and ν to denote either an s - or p -type (p_x , p_y or p_z) orbital, the NDDO approximation is defined by eqs (3.81)–(3.83).

Overlap integrals (eq. (3.51)):

$$S_{\mu\nu} = \langle \mu | \nu \rangle = \delta_{\mu\nu} \delta_{AB} \quad (3.81)$$

One-electron operator (eq. (3.24)):

$$\mathbf{h} = -\frac{1}{2} \nabla^2 - \sum_a^{N_{\text{nuclei}}} \frac{Z'_a}{|\mathbf{R}_a - \mathbf{r}|} = -\frac{1}{2} \nabla^2 - \sum_a^{N_{\text{nuclei}}} \mathbf{V}_a \quad (3.82)$$

Here Z'_a denotes that the nuclear charge has been reduced by the number of core electrons.

One-electron integrals (eq. (3.56)):

$$\begin{aligned}\langle \mu_A | \mathbf{h} | \nu_A \rangle &= \langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A | \nu_A \rangle - \sum_{a \neq A}^{N_{\text{nuclei}}} \langle \mu_A | \mathbf{V}_a | \nu_A \rangle \\ \langle \mu_A | \mathbf{h} | \nu_B \rangle &= \langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A - \mathbf{V}_B | \nu_B \rangle \\ \langle \mu_A | \mathbf{V}_C | \nu_B \rangle &= 0\end{aligned}\quad (3.83)$$

Due to the orthogonality of the atomic orbitals, the first one-centre matrix element in eq. (3.83) is zero unless the two functions are identical.

$$\langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A | \nu_A \rangle = \delta_{\mu\nu} \langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A | \mu_A \rangle \quad (3.84)$$

Two-electron integrals (eq. (3.57)):

$$\langle \mu_A \nu_B | \lambda_C \sigma_D \rangle = \delta_{AC} \delta_{BD} \langle \mu_A \nu_B | \lambda_A \sigma_B \rangle \quad (3.85)$$

3.10.2 Intermediate Neglect of Differential Overlap Approximation (INDO)

The *Intermediate Neglect of Differential Overlap* (INDO) approximation neglects all two-centre two-electron integrals that are not of the Coulomb type, in addition to those neglected by the NDDO approximations. Furthermore, in order to preserve rotational invariance, i.e. the total energy should be independent of a rotation of the coordinate system, integrals such as $\langle \mu_A | \mathbf{V}_a | \mu_A \rangle$ and $\langle \mu_A \nu_B | \mu_A \nu_B \rangle$ must be made independent of the orbital type (i.e. an integral involving a p-orbital must be the same as with an s-orbital). This has as a consequence that one-electron integrals involving two different functions on the same atom and a \mathbf{V}_a operator from another atom disappear. The INDO method involves the following additional approximations, beside those for NDDO.

One-electron integrals (eq. (3.83)):

$$\begin{aligned}\langle \mu_A | \mathbf{h} | \mu_A \rangle &= \langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A | \mu_A \rangle - \sum_{a \neq A}^{N_{\text{nuclei}}} \langle \mu_A | \mathbf{V}_a | \mu_A \rangle \\ \langle \mu_A | \mathbf{h} | \nu_A \rangle &= -\delta_{\mu\nu} \sum_{a \neq A}^{N_{\text{nuclei}}} \langle \mu_A | \mathbf{V}_a | \mu_A \rangle\end{aligned}\quad (3.86)$$

Two-electron integrals are approximated as in eq. (3.87), except that one-centre integrals $\langle \mu_A \lambda_A | \nu_A \sigma_A \rangle$ are preserved.

$$\langle \mu_A \nu_B | \lambda_C \sigma_D \rangle = \delta_{AC} \delta_{BD} \delta_{\mu\lambda} \delta_{\nu\sigma} \langle \mu_A \nu_B | \mu_A \nu_B \rangle \quad (3.87)$$

The surviving integrals are commonly denoted by γ .

$$\begin{aligned}\langle \mu_A \nu_A | \mu_A \nu_A \rangle &= \langle \mu_A \mu_A | \mu_A \mu_A \rangle = \gamma_{AA} \\ \langle \mu_A \nu_B | \mu_A \nu_B \rangle &= \gamma_{AB}\end{aligned}\quad (3.88)$$

The INDO method is intermediate between the NDDO and CNDO methods in terms of approximations.

3.10.3 Complete Neglect of Differential Overlap Approximation (CNDO)

In the *Complete Neglect of Differential Overlap* (CNDO) approximation all the Coulomb two-electron integrals are subjected to the condition in eq. (3.87), including

the one-centre integrals, and are again parameterized as in eq. (3.88). The approximations for the one-electron integrals in CNDO are the same as for INDO. The *Pariser–Pople–Parr* (PPP) method can be considered as a CNDO approximation where only π -electrons are treated.

The main difference between CNDO, INDO and NDDO is in the treatment of the two-electron integrals. While CNDO and INDO reduce these to just two parameters (γ_{AA} and γ_{AB}), all the one- and two-centre integrals are retained in the NDDO approximation. Within an *sp*-basis, however, there are only 27 different types of one- and two-centre integrals, while the number rises to over 500 for a basis containing both *s*-, *p*- and *d*-functions.

3.11 Parameterization

An *ab initio* HF calculation with a minimum basis set is rarely able to give more than a qualitative picture of the MOs, and it is of very limited value for predicting quantitative features. Introducing the ZDO approximation decreases the quality of the (already poor) wave function, i.e. a direct employment of the above NDDO/INDO/CNDO schemes is not useful. To “repair” the deficiencies due to the approximations, parameters are introduced in place of some or all of the integrals.

There are three methods that can be used for transforming the NDDO/INDO/CNDO approximations into working computational models.

- (1) The remaining integrals can be calculated from the functional form of the atomic orbitals.
- (2) The remaining integrals can be made into parameters, which are assigned values based on a few (usually atomic) experimental data.
- (3) The remaining integrals can be made into parameters, which are assigned values based on *fitting* to many (usually molecular) experimental data.

Method (2) derives specific atomic properties, such as ionization potentials and excitation energies, in terms of the parameters, and assigns their values accordingly. Method (3) takes the parameters as fitting constants, and assign their values based on a least squares fit to a large set of experimental data, analogously to the fitting of force field parameters (Section 2.3).

The CNDO, INDO and NDDO *methods* use a combination of methods (1) and (2) for assigning parameters.⁴² Some of the non-zero integrals are calculated from the atomic orbitals, and others are assigned values based on atomic ionization potentials and electron affinities. Many different versions exist; they differ in the exact way in which the parameters have been derived. Some of the names associated with these methods are CNDO/1, CNDO/2, CNDO/S, CNDO/FK, CNDO/BW, INDO/1, INDO/2, INDO/S and SINDO1. These methods are rarely used in modern computational chemistry, mainly because the “modified” methods described below usually perform better. Exceptions are INDO-based methods, such as SINDO1⁴³ and INDO/S.⁴⁴ SINDO (*Symmetric orthogonalized INDO*) methods employ the INDO approximations described above, but not the ZDO approximation for the overlap matrix. The INDO/S method (INDO parameterized for *Spectroscopy*) is especially designed for calculating electronic spectra of large molecules or systems involving heavy atoms.

The group centred around M. J. S. Dewar has used a combination of methods (2) and (3) for assigning parameter values, resulting in a class of commonly used methods. The molecular data used for parameterization are geometries, heats of formation, dipole moments and ionization potentials. These methods are denoted “modified” as their parameters have been obtained by fitting.

3.11.1 *Modified Intermediate Neglect of Differential Overlap (MINDO)*

Three versions of *Modified Intermediate Neglect of Differential Overlap* (MINDO) models exist, MINDO/1, MINDO/2 and MINDO/3. The first two attempts at parameterizing INDO gave quite poor results, but MINDO/3, introduced in 1975,⁴⁵ produced the first general purpose quantum chemical method that could successfully predict molecular properties at a relatively low computational cost. The parameterization of MINDO contains *diatomic* variables in the two-centre one-electron term, thus the β_{AB} parameters must be derived for all *pairs* of bonded atoms. The I_μ parameters are ionization potentials.

$$\begin{aligned}\langle \mu_A | \mathbf{h} | \nu_B \rangle &= \langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A - \mathbf{V}_B | \nu_B \rangle \\ &= S_{\mu\nu} \beta_{AB} (I_\mu + I_\nu) \\ S_{\mu\nu} &= \langle \mu_A | \nu_B \rangle\end{aligned}\quad (3.89)$$

MINDO/3 has been parameterized for H, B, C, N, O, F, Si, P, S and Cl, although certain combinations of these elements have been omitted. MINDO/3 is rarely used in modern computational chemistry, having been succeeded in accuracy by the NDDO methods below. Since there are parameters in MINDO that depend on two atoms, the number of parameters rises as the square of the number of elements. It is unlikely that MINDO will be parameterized beyond the abovementioned in the future.

3.11.2 *Modified NDDO models*

The MNDO, AM1 and PM3 methods⁴⁶ are parameterizations of the NDDO model where the parameterization is in terms of *atomic* variables, i.e. referring only to the nature of a single atom. MNDO, AM1 and PM3 are derived from the same basic approximations (NDDO), and differ only in the way in which the core–core repulsion is treated and in how the parameters are assigned. Each method considers only the valence *s*- and *p*-functions, which are taken as Slater type orbitals with corresponding exponents ζ_s and ζ_p .

The *one-centre one-electron integrals* have a value corresponding to the energy of a single electron experiencing the nuclear charge (U_s or U_p) plus terms from the potential due to all the other nuclei in the system (eq. (3.83)). The latter is parameterized in terms of the (reduced) nuclear charges Z' and a two-electron integral.

$$\begin{aligned}h_{\mu\nu} &= \langle \mu_A | \mathbf{h} | \nu_B \rangle = \delta_{\mu\nu} U_\mu - \sum_{a \neq A}^{N_{\text{nuclei}}} Z'_a \langle \mu_A \mu_A | \nu_A \nu_A \rangle \\ U_\mu &= \langle \mu_A | -\frac{1}{2} \nabla^2 - \mathbf{V}_A | \nu_A \rangle\end{aligned}\quad (3.90)$$

The *two-centre one-electron integrals* given by the second equation in eq. (3.83) are written as a product of the corresponding overlap integral multiplied by the average of two atomic “resonance” parameters, β .

$$\langle \mu_A | \mathbf{h} | \nu_B \rangle = \frac{1}{2} S_{\mu\nu} (\beta_\mu + \beta_\nu) \quad (3.91)$$

The overlap element $S_{\mu\nu}$ is calculated explicitly (note that this is not consistent with the ZDO approximation, and the inclusion is the origin of the “Modified” label).

There are only five types of *one-centre two-electron integrals* surviving the NDDO approximation within a *sp*-basis (eq. (3.85)).

$$\begin{aligned} \langle ss | ss \rangle &= G_{ss} \\ \langle sp | sp \rangle &= G_{sp} \\ \langle ss | pp \rangle &= H_{sp} \\ \langle pp | pp \rangle &= G_{pp} \\ \langle pp' | pp' \rangle &= G_{p2} \end{aligned} \quad (3.92)$$

The *G*-type parameters are Coulomb terms, while the *H* parameter is an exchange integral. The G_{p2} integral involves two different types of *p*-functions (i.e. p_x , p_y or p_z).

There are a total of 22 different *two-centre two-electron integrals* arising from an *sp*-basis, and these are modelled as interactions between multipoles. Electron 1 in an $\langle s\mu | s\mu \rangle$ type integral, for example, is modelled as a monopole, in an $\langle s\mu | p\mu \rangle$ type integral as a dipole and in a $\langle p\mu | p\mu \rangle$ type integral as a quadrupole. The dipole and quadrupole moments are generated as fractional charges located at specific points away from the nuclei, where the distance is determined by the orbital exponents ζ_s and ζ_p . The main reason for adapting a multipole expansion of these integrals was the limited computational resources available when these methods were developed initially. In the limit of the two nuclei being placed on top of each other, a two-centre two-electron integral becomes a one-centre two-electron integral, which puts boundary conditions on the functional form of the multipole interaction. The bottom line is that all two-centre two-electron integrals are written in terms of the orbital exponents and the one-centre two-electron parameters given in eq. (3.92).

The *core-core repulsion* is the repulsion between nuclear charges, properly reduced by the number of core electrons. The “exact” expression for this term is simply the product of the charges divided by the distance, $Z'_A Z'_B / R_{AB}$. Due to the inherent approximations in the NDDO method, however, this term is not cancelled by electron–electron terms at long distances, resulting in a net repulsion between uncharged molecules or atoms even when their wave functions do not overlap. The core–core term must consequently be modified to generate the proper limiting behaviour, which means that two-electron integrals must be involved. The specific functional form depends on the exact method, and is given below.

Each of the MNDO, AM1 and PM3 methods involves at least 12 parameters per atom, orbital exponents: $\zeta_{s/p}$; one-electron terms: $U_{s/p}$ and $\beta_{s/p}$; two-electron terms: G_{ss} , G_{sp} , G_{pp} , G_{p2} , H_{sp} ; and parameters used in the core–core repulsion, α , and for the AM1 and PM3 methods also a , b and c constants, as described below.

3.11.3 Modified Neglect of Diatomic Overlap (MNDO)

The core–core repulsion of the *Modified Neglect of Diatomic Overlap* (MNDO) model⁴⁷ has the form given in eq. (3.93).

$$V_{nn}^{\text{MNDO}}(\text{A}, \text{B}) = Z'_\text{A} Z'_\text{B} \langle s_\text{A} s_\text{A} | s_\text{B} s_\text{B} \rangle (1 + e^{-\alpha_\text{A} R_\text{AB}} + e^{-\alpha_\text{B} R_\text{AB}}) \quad (3.93)$$

The α exponents are taken as fitting parameters.

Interactions involving O—H and N—H bonds are treated differently.

$$V_{nn}^{\text{MNDO}}(\text{A}, \text{H}) = Z'_\text{A} Z'_\text{H} \langle s_\text{A} s_\text{A} | s_\text{H} s_\text{H} \rangle (1 + R_\text{AH} e^{-\alpha_\text{A} R_\text{AH}} + e^{-\alpha_\text{H} R_\text{AH}}) \quad (3.94)$$

In addition, MNDO uses the approximation $\zeta_\text{s} = \zeta_\text{p}$ for some of the lighter elements. The G_ss , G_sp , G_pp , G_p2 and H_sp parameters are taken from atomic spectra, while the others are fitted to molecular data. Although MNDO has been succeeded by the AM1 and PM3 methods, it is still used for some types of calculations where MNDO is known to give better results.

Some known limitations of the MNDO model are:

- (1) Branched and sterically crowded hydrocarbons (such as neopentane) are predicted to be too unstable, relative to their straight-chain analogues.
- (2) Four-membered rings are too stable.
- (3) Weak interactions are unreliable, for example MNDO does not predict hydrogen bonds.
- (4) Hypervalent molecules, such as sulfoxides and sulfones, are too unstable.
- (5) Activation energies for bond breaking/forming reactions are too high.
- (6) Non-classical structures are predicted to be unstable relative to classical structures (for example ethyl radical).
- (7) Proton affinities are poorly predicted.
- (8) Oxygen-containing substituents on aromatic rings are out-of-plane (for example nitrobenzene).
- (9) Peroxide bonds are too short by $\sim 0.17 \text{ \AA}$
- (10) The C—X—C angle in ethers and sulfides is too large by $\sim 9^\circ$.

MNDOC⁴⁸ (C for correlation) has the same functional form as MNDO, however, electron correlation is *explicitly* calculated by second-order perturbation theory. The derivation of the MNDOC parameters is done by fitting the correlated MNDOC results to experimental data. Electron correlation in MNDO is only included implicitly via the parameters, from fitting to experimental results. Since the training set only includes ground state stable molecules, MNDO has problems treating systems where the importance of electron correlation is substantially different from “normal” molecules. MNDOC consequently performs significantly better for system where this is not the case, such as transition structures and excited states.

3.11.4 Austin Model 1 (AM1)

After some experience with MNDO, it became clear that there were certain systematic errors. For example, the repulsion between two atoms that are 2–3 \AA apart is too high. This has as a consequence that activation energies in general are too large. The source was traced to a too repulsive interaction in the core–core potential. In order to

remedy this, the core–core function was modified by adding Gaussian functions, and the whole model was re-parameterized. The result was called *Austin Model 1* (AM1)⁴⁹, in honour of Dewar’s move to the University of Austin at the time. The core–core repulsion of AM1 has the form given in eq. (3.95).

$$V_{nn}^{\text{AM1}}(\text{A}, \text{B}) = V_{nn}^{\text{MNDO}}(\text{A}, \text{B}) + \frac{Z'_\text{A} Z'_\text{B}}{R_{\text{AB}}} \sum_k \left(a_{k\text{A}} e^{-b_{k\text{A}}(R_{\text{AB}} - c_{k\text{A}})^2} + a_{k\text{B}} e^{-b_{k\text{B}}(R_{\text{AB}} - c_{k\text{B}})^2} \right) \quad (3.95)$$

Here k is between 2 and 4, depending on the atom. It should be noted that the Gaussian functions were added more or less as patches onto the underlying parameters, which explains why a different number of Gaussians is used for each atom. As for MNDO, the G_{ss} , G_{sp} , G_{pp} , G_{p2} and H_{sp} parameters are taken from atomic spectra, while the rest, including the a_k , b_k and c_k constants, are fitted to molecular data.

Some known improvements and limitations of the AM1 model are:

- (1) AM1 does predict hydrogen bonds with an approximately correct strength, but the geometry is often wrong.
- (2) Activation energies are much improved over MNDO.
- (3) Hypervalent molecules are improved over MNDO, but still have significantly larger errors than other types of compounds.
- (4) Alkyl groups are systematically too stable by $\sim 8 \text{ kJ/mol}$ per CH_2 group.
- (5) Nitro compounds are systematically too unstable.
- (6) Peroxide bonds are too short by $\sim 0.17 \text{ \AA}$.
- (7) Phosphor compounds have problems when atoms are $\sim 3 \text{ \AA}$ apart, producing wrong geometries. P_4O_6 for example is predicted to have P—P bonds differing by 0.4 \AA , although experimentally they are identical.
- (8) The *gauche* conformation in ethanol is predicted to be more stable than the *trans*.

3.11.5 Modified Neglect of Diatomic Overlap, Parametric Method number 3 (PM3)

The parameterization of MNDO and AM1 had been done essentially by hand, taking the G_{ss} , G_{sp} , G_{pp} , G_{p2} and H_{sp} parameters from atomic data and varying the rest until a satisfactory fit had been obtained. Since the optimization was done by hand, only relatively few reference compounds could be included. J. J. P. Stewart made the optimization process automatic by deriving and implementing formulas for the derivative of a suitable error function with respect to the parameters.⁵⁰ All parameters could then be optimized simultaneously, including the two-electron terms, and a significantly larger training set with several hundred data could be employed. In this re-parameterization, the AM1 expression for the core–core repulsion (eq. (3.85)) was kept, except that only two Gaussians were assigned to each atom. These Gaussian parameters were included as an integral part of the model, and allowed to vary freely. The resulting method was denoted *Modified Neglect of Diatomic Overlap, Parametric Method Number 3* (MNDO-PM3 or PM3 for short), and is essentially AM1 with all the parameters fully optimized. In a sense, it is *the* best set of parameters (or at least a good local minimum) for the given set of experimental data. The optimization process, however, still requires some human intervention in selecting the experimental data and assigning appropriate weight factors to each set of data.

Some known limitations of the PM3 model are:

- (1) Almost all sp^3 -nitrogens are predicted to be pyramidal, which is contrary to experimental data.
- (2) Hydrogen bonds are too short by ~ 0.1 Å.
- (3) The *gauche* conformation in ethanol is predicted to be more stable than the *trans*.
- (4) Bonds between Si and Cl, Br and I are underestimated, the Si—I bond in H_3SiI , for example, is too short by ~ 0.4 Å.
- (5) H_2NNH_2 is predicted to have a C_{2h} structure, while the experimental result is C_{2v} , and ClF_3 is predicted to have a D_{3h} structure, while the experimental result is C_{2v} .
- (6) The charge on nitrogen atoms is often of “incorrect” sign and “unrealistic” magnitude.

Some common limitations of MNDO, AM1 and PM3 are:

- (1) Rotational barriers for bonds that have partly double bond character are significantly too low. The barrier for rotation around the central bond in butadiene is calculated to be only 2–8 kJ/mol, in contrast to the experimental value of 25 kJ/mol.⁵¹ Similarly, the rotational barrier around the C—N bond in amides is calculated to be 30–50 kJ/mol, which is roughly a factor of two smaller than the experimental value. A purely *ad hoc* fix has been made by adding a force field rotational term to the C—N bond that raises the value to ~ 100 kJ/mol and brings it into better agreement with experimental data.
- (2) Weak interactions, such as van der Waals complexes or hydrogen bonds, are poorly predicted. Either the interaction is too weak, or the minimum energy geometry is wrong.
- (3) Conformational energies for peptides are poorly reproduced.⁵²
- (4) The bond length to nitrosyl groups is underestimated. The N—N bond in N_2O_3 , for example, is ~ 0.7 Å too short.
- (5) Although MNDO, AM1 and PM3 have parameters for some metals, these are often based on only a few experimental data. Calculations involving metals should thus be treated with care.

The MNDO, AM1 and PM3 methods have been parameterized for most of the main group elements,⁵³ and parameters for many of the transition metals are also being developed under the name PM3(tm), which includes d-orbitals. The PM3(tm) set of parameters are determined exclusively from geometrical data (X-ray) since there are very few reliable energetic data available for transition metal compounds

3.11.6 Parametric Method number 5 (PM5) and PDDG/PM3 methods

Two approaches have appeared that try to further improve on the performance of the PM3 method. The PM5 (PM4 being an unpublished experimental version) method reintroduces diatomic parameters for the core–core repulsion, and the published results suggest that PM5 represent a slight improvement on the PM3 results.⁵⁴ No details of the methodology and parameterization have been published so far.

A similar approach using *Pairwise Distance Directed Gaussian* (PDDG) in connection with the MNDO and PM3 methods has also been reported.⁵⁵ The idea is related to the concept used in AM1 and PM3 by introducing parameterized Gaussian func-

tions for describing the core–core repulsion, except that the modification is based on interatomic distances, although the parameters are still purely atomic. The latter prevents the exponential increase in parameters with the number of atoms. The available results suggest a slight improvement over the regular MNDO or PM3 methods, but it is difficult to assess whether the improvement is simply due to more fitting parameters or to fundamentally better modelling of the underlying physical problem.

3.11.7 The MNDO/d and AM1/d methods

With only *s*- and *p*-functions included, the MNDO/AM1/PM3 methods are unable to treat a large part of the periodic table. Furthermore, from *ab initio* calculations it is known that d-orbitals significantly improve the results for compounds involving second row elements, especially hypervalent species. The main problem in extending the NDDO formalism to include d-orbitals is the significant increase in distinct two-electron integrals that ultimately must be assigned suitable values. For an *sp*-basis there are only five one-centre two-electron integrals, while there are 17 in an *spd*-basis. Similarly, the number of two-centre two-electron integrals rises from 22 to 491 when *d*-functions are included.

Thiel and Voityuk have constructed a workable NDDO model that also includes d-orbitals for use in connection with MNDO, called *MNDO/d*.⁵⁶ With reference to the above description for MNDO/AM1/PM3, it is clear that there are immediately three new parameters: ζ_d , U_d and β_d (eqs (3.90) and (3.91)). Of the 12 new one-centre two-electron integrals, only one (G_{dd}) is taken as a freely varied parameter. The other 11 are calculated analytically based on pseudo-orbital exponents, which are assigned such that the analytical formulas regenerate G_{ss} , G_{pp} and G_{dd} .

With only *s*- and *p*-functions present, the two-centre two-electron integrals can be modelled by multipoles up to order 4 (quadrupoles), however, with *d*-functions present multipoles up to order 16 must be included. In MNDO/d all multipoles beyond order 4 are neglected. The resulting MNDO/d method typically employs 15 parameters per atom, and it currently contains parameters for the following elements (beyond those already present in MNDO): Na, Mg, Al, Si, P, S, Cl, Br, I, Zn, Cd and Hg. Recently this technology has been used in connection with the AM1 model as well, which at least for phosphorous yields a further improvement.⁵⁷

3.11.8 Semi Ab initio Method 1

The philosophy behind the *Semi Ab Initio Method 1* (SAM1 and SAM1D) model is slightly different from the other “modified” methods.⁵⁸ It is again based on the NDDO approximation, but instead of replacing all integrals with parameters, the one- and two-centre two-electron integrals are calculated directly from the atomic orbitals. These integrals are then scaled by a function containing adjustable parameters to fit experimental data (R_{AB} being the interatomic distance).

$$\langle \mu_A \nu_B | \mu_A \nu_B \rangle \rightarrow f(R_{AB}) \langle \mu_A \nu_B | \mu_A \nu_B \rangle \quad (3.96)$$

The advantage is that basis sets involving d-orbitals are readily included (defining the SAM1D method), making it possible to perform calculations on a larger fraction of the periodic table. The SAM1 method explicitly uses the minimum STO-3G basis set,

but it is in principle also possible to use extended basis sets with this model. The actual calculation of the integrals makes the SAM1 method somewhat slower than MNDO/AM1/PM3, but only by a factor of ~2. The SAM1/SAM1D methods have been parameterized for these elements: H, Li, C, N, O, F, Si, P, S, Cl, Fe, Cu, Br and I. Although the SAM1 method was proposed in 1993, no details of the functional form or parameterization have been published, and there do not appear to have been any recent developments.

3.12 Performance of Semi-Empirical Methods

The electronic energy (including the core–core repulsion) calculated by MINDO, MNDO, MNDO/d, AM1 and PM3 is, in analogy with *ab initio* methods, the total energy relative to a situation where the nuclei (with their core electrons) and the valence electrons are infinitely separated. The electronic energy is normally converted to a heat of formation by subtracting the electronic energy of the isolated atoms that make up the system, and adding the experimental atomic heat of formation. It should be noted that thermodynamic corrections (e.g. zero point energies, see Section 13.5.5) should *not* be added to the ΔH_f values, as these are included implicitly by the parameterization.

$$\Delta H_f(\text{molecule}) = E_{\text{elec}}(\text{molecule}) - \sum^{M_{\text{atoms}}} E_{\text{elec}}(\text{atoms}) - \sum^{M_{\text{atoms}}} \Delta H_f(\text{atoms}) \quad (3.97)$$

Some typical errors in heat of formation for the MNDO, AM1 and PM3 methods are given in Table 3.1.⁵⁹ The exact numbers of course depend on which, and how many, compounds have been selected for comparison, thus the numbers should only be taken as a guideline for the accuracy expected.

Table 3.1 Average heat of formation error (kJ/mol)

Compounds	Number of compounds	MNDO	AM1	PM3
H, C, N, O	276	77	44	33
F	133	352	207	47
Si	78	96	87	59
All normal valent	607	102	62	47
Hypervalent	106	437	261	72
All	713	193	116	49

Some typical errors in bond distances are given in Table 3.2.

Table 3.2 Average errors in bond distances (Å)

Bond to:	MNDO	AM1	PM3
H	0.015	0.006	0.005
C	0.002	0.002	0.002
N	0.015	0.014	0.012
O	0.017	0.011	0.006
F	0.023	0.017	0.011
Si	0.030	0.019	0.045

Angles are typically predicted with an accuracy of a few degrees. The average errors for MNDO, AM1 and PM3 are 4.3° , 3.3° and 3.9° , respectively. Ionization potentials are typically accurate to 0.5–1.0 eV. Average errors for MNDO, AM1 and PM3 are 0.78, 0.61 and 0.57 eV, respectively. Average errors for dipole moments are 0.45, 0.35 and 0.38 debye, respectively.

Since AM1 contains more adjustable parameters than MNDO, and since PM3 can be considered to be a version of AM1 with all the parameters fully optimized, it is expected that the error decreases in the order MNDO > AM1 > PM3. This is indeed what is observed in the above tables. It should be noted, however, that the data in the tables refer to averages, thus for specific compounds or classes of compounds the ordering may be different. Bonds between silicon and iodine with PM3 are examples where a specific compound may be poorly described, although the average description for all compounds is better. It is clear that the PM3 method will perform better than AM1 in an average sense since the two-electron integrals are optimized to give a better fit to the given molecular data set. This does not mean, however, that PM3 necessarily will perform better than AM1 (or MNDO) for properties not included in the training set. Indeed it has been argued that the AM1 method tends to give more “realistic” values for atomic charges than PM3, especially for compounds involving nitrogen. An often quoted example is formamide, and the Mulliken population analysis by different methods is given in Table 3.3.

Table 3.3 Mulliken charges in formamide with different methods

	MNDO	AM1	PM3	HF/6–31G(d,p)	MP2/6–31G(d,p)
C	0.37	0.26	0.16	0.56	0.40
N	–0.49	–0.62	–0.13	–0.73	–0.63
O	–0.39	–0.40	–0.38	–0.56	–0.43

The negative charge on nitrogen produced by PM3 is significantly smaller than by the other methods, but it should be noted that atomic charges are not well-defined quantities, as discussed in Chapter 9. Nevertheless, it may indicate that the electrostatic potential generated by a PM3 wave function is of lower quality than one generated by the AM1 method.

Table 3.4 shows a comparison for some of the elements that have been parameterized for the MNDO, MNDO/d, AM1, PM3, SAM1 and SAM1D methods.

Considering that the parameters for the MNDO/d method for all first row elements (which are present in most of the training set of compounds) are identical to MNDO, the improvement by addition of *d*-functions is quite impressive. It should also be noted that MNDO/d only contains 15 parameters, compared with 18 for PM3, and that some of the 15 parameters are taken from atomic data (analogous to the MNDO/AM1 parameterization), and are not used in the molecular data fitting as in PM3.

The apparent accuracy of 20–40 kJ/mol for calculating heats of formation with semi-empirical methods is slightly misleading. Normally the interest is in relative energies of different species, and since the heat of formation errors are essentially random, relative energies may not be predicted as well (two random errors of 40 kJ/mol may

Table 3.4 Average heat of formation error (kJ/mol)

Compounds	Number of compounds	MNDO	AM1	PM3	MNDO/d	SAM1	SAM1D
Al	29	92	44	69	21		
Si	84	50	36	25	26	33	47
P	43	162	61	72	32	60	63
S	99	203	43	31	23	35	33
Cl	85	165	122	44	16	46	20
Br	51	68	64	34	14	36	22
I	42	106	91	56	17	28	28
Zn	18	88	71	62	21		
Hg	37	57	38	32	9		
Si, P, S, Cl, Br, I	404	132	67	40	21	39	34
Al, Si, P, S, Cl, Br, I, Zn, Hg	488	122	64	42	21		

add up to an error of 80 kJ/mol). This is in contrast to *ab initio* methods, which are usually better at predicting relative rather than absolute energies, since errors in this case tend to be systematic and at least partly cancel out when comparing similar systems.

3.13 Hückel Theory

3.13.1 Extended Hückel theory

The Hückel methods perform the parameterization on the Fock matrix elements (eq. (3.50)), and not at the integral level, as do NDDO/INDO/CNDO. This means that Hückel methods are non-iterative and they only require a single diagonalization of the Fock (Hückel) matrix. The *Extended Hückel Theory* (EHT) or *Method* (EHM), developed primarily by R. Hoffmann, again only consider the valence electrons.⁶⁰ It makes use of Koopmans' theorem (eq. (3.47)) and assigns the diagonal elements in the **F** matrix to be atomic ionization potentials. The off-diagonal elements are parameterized as averages of the diagonal elements, weighted by an overlap integral. The overlap integrals are actually calculated, i.e. the ZDO approximation is not invoked. The basis functions are taken as Slater type orbitals, with the exponents assigned according to the rules of Slater.⁶¹

$$\begin{aligned}
 F_{\mu\mu} &= -I_{\mu} \\
 F_{\mu\nu} &= -\frac{1}{2} K (I_{\mu} + I_{\nu}) S_{\mu\nu}
 \end{aligned}
 \tag{3.98}$$

The *K* constant is usually taken as 1.75, as this value reproduces the rotational barrier in ethane. An essentially identical approach has been used for periodic systems within the physics community, where it is here known as the *tight binding* model.⁶² Recent work in this area has used an approach to parameterize against density functional results, thereby providing a computationally very efficient model capable of yielding fairly accurate results.⁶³

Since the diagonal elements only depend on the nature of the atom (i.e. the nuclear charge), this means for example that all carbon atoms have the same ability to attract electrons. After having performed a Hückel calculation, the actual number of electrons associated with atom A, ρ_A , can be calculated according to eq. (3.99) (see Section 9.1, eqs (9.5) and (9.4)).

$$\rho_A = \sum_i^{\text{MO}} n_i \sum_{\alpha \in A}^{\text{AO}} \sum_{\beta}^{\text{AO}} c_{\alpha i} c_{\beta i} S_{\alpha \beta} \quad (3.99)$$

The effective (net) atomic charge Q_A is given as the (reduced) nuclear charge minus the electronic contribution.

$$Q_A = Z'_A - \rho_A \quad (3.100)$$

In general, it is unlikely that all carbon atoms have the exact same charge, i.e. owing to the different environments their ability to attract electrons is no longer equal. This may be argued to be inconsistent with the initial assumption of all carbons having the same diagonal elements in the Hückel matrix. In order to achieve “self-consistency”, a diagonal element $F_{\mu\mu}$ belonging to atom A may be modified by the calculated atomic charge.

$$F_{\mu\mu} = -I_{\mu} + \omega Q_A \quad (3.101)$$

The ω parameter determines the weight of the charge on the diagonal elements. Since Q_A is calculated from the *results* (MO coefficients, eq. (3.99)) but enters the Hückel matrix that *produces* the results (by diagonalization), such schemes become iterative. Methods where the matrix elements are modified by the calculated charge are often called *charge iteration* or *self-consistent* (Hückel) methods. Similar self-consistent charge models are used within the tight binding formalism.⁶⁴

The main advantage of extended Hückel theory is that only atomic ionization potentials are required, and it is easily parameterized to the whole periodic table. Extended Hückel theory can be used for large systems involving transition metals, where it often is the only possible computational model. The very approximate method of extended Hückel theory makes it unsuitable for geometry optimizations without additional modifications,⁶⁵ or for calculations of energetic features at any reasonable level of accuracy. It is primarily used for obtaining qualitatively correct MOs, which can for example be used as an initial guess of the density matrix for *ab initio* SCF calculations, or for use in connection with qualitative theories, as discussed in Chapter 15. Orbital energies (and thereby the total energy), however, in many cases show the correct trend for geometry perturbations corresponding to bond bending or torsional changes, and thus qualitative features regarding molecular shapes may often be predicted or rationalized from EHT calculations.

3.13.2 Simple Hückel theory

In the simple Hückel model the approximations are taken to the limit.⁶⁶ Only planar conjugated systems are considered. The σ -orbitals, which are symmetric with respect to a reflection in the molecular plane, are neglected. Only the π -electrons (antisymmetric with respect to the molecular mirror plane) are considered. The overlap matrix

is taken as a unit matrix and a diagonal element of the \mathbf{F} matrix is assigned a value of α , which depends on the atom type. Off-diagonal elements are taken either as β (depending on the two atom types) or zero, conditioned on whether the two atoms are “neighbours” (i.e. connected by a σ -bond) or not.

$$\begin{aligned} F_{\mu_A\mu_A} &= \alpha_A \\ F_{\mu_A\mu_B} &= \beta_{AB} \quad (\text{A and B are neighbours}) \\ F_{\mu_A\mu_B} &= 0 \quad (\text{A and B are not neighbours}) \end{aligned} \quad (3.102)$$

Atoms are assigned “types”, much as in force field methods, i.e. the parameters depend on the nuclear charge and the bonding situation. The α_A and β_{AB} parameters for atom types A and B are related to the corresponding parameters for sp^2 -hybridized carbon by means of the dimensionless constants h_A and k_{AB} .

$$\begin{aligned} \alpha_A &= \alpha_C + h_A \beta_{CC} \\ \beta_{AB} &= k_{AB} \beta_{CC} \end{aligned} \quad (3.103)$$

The carbon parameters α_C and β_{CC} are normally just denoted α and β , and are rarely assigned numerical values. Simple Hückel theory thus only considers the *connectivity* of the π -atoms: there is no information about the molecular geometry entering the calculation (e.g. whether some bonds are shorter or longer than others, or differences in bond angles).

In analogy to extended Hückel theory, there are also charge iterative methods for simple Hückel theory. The equivalent of eq. (3.99) is given in eq. (3.104).

$$\rho_A = \sum_i^{\text{MO}} n_i c_{Ai}^2 \quad (3.104)$$

Eq. (3.101) becomes eq. (3.105).

$$\alpha'_A = \alpha_A + \omega(n_A - \rho_A)\beta \quad (3.105)$$

where n_A is the number of π -electrons involved from atom A.

The Hückel method is essentially only used for educational purposes or for very qualitative orbital considerations. It has the ability to produce qualitatively correct MOs, involving a computational effort that is within reach of doing by hand.

3.14 Limitations and Advantages of Semi-Empirical Methods

The neglect of all three- and four-centre two-electron integrals reduces the construction of the Fock matrix from a formal order of M_{basis}^4 to M_{basis}^2 . However, the time required for diagonalization of the \mathbf{F} matrix grows as the cube of the matrix size, thus *semi-empirical methods formally scale as the cube of the number of basis functions in the limit of large molecules*. Diagonalization of a matrix becomes significant when the size exceeds $\sim 10000 \times 10000$. Several iterations are required for solving the SCF equations, and usually the geometry is also optimized, requiring several calculations at different geometries. This places the current limit of semi-empirical methods at around 1000 atoms. It should be noted that the conventional method of solving the HF equations by diagonalizing the Fock matrix rapidly becomes the rate-limiting step in semi-

empirical methods. Recent developments have therefore concentrated on formulating alternative methods for obtaining the SCF orbitals without the need for diagonalization.⁶⁷ Such methods display linear scaling with the number of atoms, allowing calculations to be performed for systems containing several thousand atoms.

The parameterization of MNDO/AM1/PM3 is performed by adjusting the constants involved in the different methods such that the results of HF calculations fit experimental data as closely as possible. This is in a sense wrong. We know that the HF method cannot give the correct result, even in the limit of an infinite basis set and without approximations. The HF results lack electron correlation, as will be discussed in Chapter 4, but the experimental data of course include such effects. This may be viewed as an advantage, the electron correlation effects are implicitly taken into account in the parameterization, and we need not perform complicated calculations to improve deficiencies in the HF procedure. However, it becomes problematic when the HF wave function cannot describe the system even qualitatively correctly, as with for example biradicals and excited states. In such cases, additional flexibility can be introduced in the trial wave function by adding more Slater determinants, for example by means of a CI procedure (see Chapter 4 for details). But electron correlation is then taken into account twice, once in the parameterization at the HF level, and once explicitly by the CI calculation.

Semi-empirical methods share the advantages and disadvantages of force field methods: they perform best for systems where much experimental information is already available but they are unable to predict totally unknown compound types. The dependence on experimental data is not as severe as for force field methods, owing to the more complex functional form of the model. The NDDO methods require only atomic parameters, not di-, tri- and tetra-atomic parameters as do force field methods. Once a given atom has been parameterized, all possible compound types involving this element can be calculated. The smaller number of parameters and the more complex functional form has the disadvantage compared with force field methods that it is very difficult to “repair” a specific problem by re-parameterization. The lack of a reasonable rotational barrier in amides, for example, cannot be attributed to an “improper” value for a single (or a few) parameter(s). Too low a rotational barrier in a force field model can easily be fixed by increasing the values of the corresponding torsional parameters. The clear advantage of semi-empirical methods over force field techniques is the ability to describe bond breaking and bond forming reactions.

Semi-empirical methods are zero-dimensional, just as force field methods are. There is no way of assessing the reliability of a given result within the method. This is due to the selection of a minimum basis set. The only way of judging results is by calibration, i.e. by comparing the accuracy of other calculations on similar systems with experimental data.

Semi-empirical models provide a method for calculating the electronic wave function, which may be used for predicting a variety of properties. There is nothing to hinder the calculation of say the polarizability of a molecule (the second derivative of the energy with respect to an external electric field), although it is known from *ab initio* calculations that good results require a large polarized basis set including diffuse functions, and the inclusion of electron correlation. Semi-empirical methods such as AM1 or PM3 only have a minimum basis (lacking polarization and diffuse functions), electron correlation is only included implicitly by the parameters, and no polarizability

data have been used for deriving the parameters. Whether such calculations can produce reasonable results, as compared with experimental data, is questionable, and careful calibration is certainly required. Again it should be emphasized: *The ability to perform a calculation is no guarantee that the results can be trusted!*

References

1. A. Szabo, N. S. Ostlund, *Modern Quantum Chemistry*, McGraw-Hill, 1982; R. McWeeny, *Methods of Molecular Quantum Mechanics*, Academic Press, 1992; W. J. Hehre, L. Radom, J. A. Pople, P. v. R. Schleyer, *Ab Initio Molecular Orbital Theory*, Wiley, 1986; J. Simons, *J. Phys. Chem.*, **95** (1991), 1017; J. Simons, J. Nichols, *Quantum Mechanics in Chemistry*, Oxford University Press, 1997; T. Helgaker, P. Jørgensen, J. Olsen, *Molecular Electronic Structure Theory*, Wiley, 2000.
2. W. Kolos, L. Wolniewicz, *J. Chem. Phys.*, **41** (1964), 3663; B. T. Sutcliffe, *Adv. Quant. Chem.*, **28** (1997), 65. B. T. Sutcliffe, *Adv. Chem. Phys.*, **114** (2000), 1.
3. N. C. Handy, A. M. Lee, *Chem. Phys. Lett.*, **252** (1996), 425.
4. L. J. Butler, *Ann. Rev. Phys. Chem.*, **49** (1998), 125.
5. E. F. Valeev, C. D. Sherrill, *J. Chem. Phys.*, **118** (2003), 3921.
6. A. D. Bochevarov, E. F. Valeev, D. C. Sherrill, *Mol. Phys.*, **102** (2004), 111.
7. T. A. Koopmans, *Physica*, **1** (1933), 104.
8. J. Kobus, *Adv. Quant. Chem.*, **28** (1997), 1.
9. C. C. J. Roothaan, *Rev. Mod. Phys.*, **23** (1951), 69; G. G. Hall, *Proc. R. Soc. (London)*, **A205** (1951), 541.
10. M. Head-Gordon, J. A. Pople, *J. Phys. Chem.*, **92** (1988), 3063.
11. J. A. Pople, R. K. Nesbet, *J. Chem. Phys.*, **22** (1954), 571.
12. C. Kollmar, *Int. J. Quant. Chem.*, **62** (1997), 617.
13. V. R. Saunders, I. H. Hillier, *Mol. Phys.*, **28** (1974), 819.
14. P. Pulay, *J. Comp. Chem.*, **3** (1982), 556.
15. K. N. Kudin, G. E. Scuseria, E. Cancès, *J. Chem. Phys.*, **116** (2002), 8255.
16. M. Head-Gordon, J. A. Pople, *J. Phys. Chem.*, **92** (1988), 3063. H. Larsen, J. Olsen, P. Jørgensen, T. Helgaker, *J. Chem. Phys.*, **115** (2001), 9685.
17. J. Douady, Y. Ellinger, R. Subra, B. Levy, *J. Chem. Phys.*, **72** (1980), 1452.
18. R. McWeeny, *Rev. Mod. Phys.*, **32** (1960), 335.
19. E. Cancès, *J. Chem. Phys.*, **114** (2001), 10616.
20. J. M. Millan, G. E. Scuseria, *J. Chem. Phys.*, **105** (1996), 5569; H. Larsen, J. Olsen, P. Jørgensen, T. Helgaker, *J. Chem. Phys.*, **115** (2001), 9685.
21. G. Chaban, M. W. Schmidt, M. S. Gordon, *Theor. Chem. Acc.*, **97** (1997), 88.
22. L. Thøgersen, J. Olsen, A. Köhn, P. Jørgensen, P. Salek, T. Helgaker, *J. Chem. Phys.*, **123** (2005), 074103.
23. H. Larsen, J. Olsen, P. Jørgensen, T. Helgaker, *J. Chem. Phys.*, **115** (2001), 9685.
24. R. Seeger, J. A. Pople, *J. Chem. Phys.*, **66** (1977), 3045; R. Bauernschmitt, R. Aldrichs, *J. Chem. Phys.*, **104** (1996), 9047.
25. P.-O. Löwdin, I. Mayer, *Adv. Quant. Chem.*, **24** (1992), 79.
26. E. R. Davidson, W. T. Borden, *J. Phys. Chem.*, **87** (1983), 4783.
27. J. Almlöf, K. Faegri Jr, K. Korsell, *J. Comp. Chem.*, **3** (1982), 385; J. Almlöf, *Modern Electronic Structure Theory*, Part I, D. Yarkony, Ed., World Scientific 1995, pp. 110–151.
28. D. S. Lambrecht, C. Ochsenfeld, *J. Chem. Phys.*, **123** (2005), 184101.
29. D. L. Strout, G. E. Scuseria, *J. Chem. Phys.*, **102** (1995), 8448.
30. S. Goedecker, *Rev. Mod. Phys.*, **71** (1999), 1085.
31. H. G. Petersen, D. Soelvason, J. W. Perram, E. R. Smith, *J. Chem. Phys.*, **101** (1994), 8870.

32. C. A. White, M. Head-Gordon, *J. Chem. Phys.*, **104** (1996), 2620; M. C. Strain, G. E. Scuseria, M. J. Frisch, *Science*, **271** (1996), 51.
33. A. M. Lee, S. W. Taylor, J. P. Dombroski, P. M. W. Gill, *Phys. Rev. A.*, **55** (1997), 3233.
34. E. Schwegler, M. Challacombe, M. Head-Gordon, *J. Chem. Phys.*, **106** (1997), 9708.
35. A. van der Vaart, V. Gogonea, S. L. Dixon, K. M. Merz Jr, *J. Comp. Chem.*, **21** (2000), 1494.
36. F. Weigend, *Phys. Chem. Chem. Phys.*, **4** (2002), 4285.
37. R. Dovesi, B. Civalieri, R. Orlando, C. Roetti, V. R. Saunders, *Rev. Comp. Chem.*, **21** (2005), 1.
38. H. J. Monkhorst, J. D. Pack., *Phys. Rev. B*, **13** (1976), 5188.
39. J. E. Jaffe, A. C. Hess, *J. Chem. Phys.*, **105** (1996), 10983.
40. K. N. Kudin, G. E. Scuseria, *Phys. Rev. B*, **61** (2000), 16440.
41. J. Sadley, *Semi-Empirical Methods of Quantum Chemistry*, Wiley, 1985; M. C. Zerner, *Rev. Comp. Chem.*, **2** (1991), 313, T. Bredow, K. Jug, *Theor. Chem. Acc.*, **113** (2005), 1.
42. W. P. Anderson, T. R. Cundari, M. C. Zerner, *Int. J. Quant. Chem.*, **39** (1991), 31.
43. J. Li, P. C. de Mello, K. Jug, *J. Comp. Chem.*, **13** (1992), 85.
44. M. Kotzian, N. Rösch, M. C. Zerner, *Theo. Chem. Acta*, **81** (1992), 201.
45. R. C. Bingham, M. J. S. Dewar, D. H. Lo, *J. Am. Chem. Soc.*, **97** (1975), 1294.
46. J. J. P. Stewart, *Rev. Comp. Chem.*, **1** (1990), 45.
47. M. J. S. Dewar, W. Thiel, *J. Am. Chem. Soc.*, **99** (1977), 4899.
48. W. Thiel, *J. Am. Chem. Soc.*, **103** (1981), 1413 and 1421; A. Schweig, W. Thiel, *J. Am. Chem. Soc.*, **103** (1981), 1425.
49. M. J. S. Dewar, E. G. Zoebisch, E. F. Healy, J. J. P. Stewart, *J. Am. Chem. Soc.*, **107** (1985), 3902.
50. J. J. P. Stewart, *J. Comp. Chem.*, **10** (1989), 209 and 221.
51. R. Engeln, D. Consalvo, J. Reuss, *Chem. Phys.*, **160** (1992), 427.
52. K. Möhle, H.-J. Hofmann, W. Thiel, *J. Comp. Chem.*, **22** (2001), 509.
53. J. J. P. Stewart, *J. Mol. Model.*, **10** (2004), 155.
54. J. J. P. Stewart, *J. Mol. Model.*, **10** (2004), 6.
55. M. P. Repasky, J. Chandrasekhar, W. L. Jorgensen, *J. Comp. Chem.*, **23** (2002), 1601.
56. W. Thiel, A. A. Voityuk, *J. Phys. Chem.*, **100** (1996), 616; W. Thiel, *Adv. Chem. Phys.*, **93** (1996), 703.
57. X. Lopez, Xabier, D. M. York, *Theor. Chem. Acc.*, **109** (2003), 149.
58. M. J. S. Dewar, C. Jie, J. Yu, *Tetrahedron*, **49** (1993), 5003.
59. J. J. P. Stewart, *J. Comp.-Aid. Mol. Design*, **4** (1990), 1.
60. R. Hoffmann, *J. Chem. Phys.*, **39** (1963), 1397.
61. J. C. Slater, *Phys. Rev.*, **36** (1930), 57.
62. D. Porezag, Th. Frauenheim, Th. Köhler, G. Seifert, R. Kaschner, *Phys. Rev. B.*, **51** (1995), 12947.
63. T. Krüger, M. Elstner, P. Schifffels, T. Frauenheim, *J. Chem. Phys.*, **122** (2005), 114110.
64. M. Elstner, D. Porezag, G. Jungnickel, J. Elsner, M. Haugk, Th. Frauenheim, S. Suhai, G. Seifert, *Phys. Rev. B.*, **58** (1998), 7260.
65. S. L. Dixon, P. C. Jurs, *J. Comp. Chem.*, **15** (1994), 733.
66. K. Yates, *Hückel Molecular Orbital Theory*, Academic Press, 1978.
67. J. J. P. Stewart, *Int. J. Quant. Chem.*, **58** (1996), 133; A. D. Daniels, J. M. Millam, G. E. Scuseria, *J. Chem. Phys.*, **107** (1997), 425.

4 Electron Correlation Methods

The Hartree–Fock method generates solutions to the Schrödinger equation where the real electron–electron interaction is replaced by an average interaction (Chapter 3). In a sufficiently large basis set, the HF wave function is able to account for ~99% of the total energy, but the remaining ~1% is often very important for describing chemical phenomena. The difference in energy between the HF and the lowest possible energy in the given basis set is called the *Electron Correlation* (EC) energy.¹ Physically, it corresponds to the motion of the electrons being correlated, i.e. on the average they are further apart than described by the HF wave function. As shown below, an unrestricted Hartree–Fock (UHF) type of wave function is, to a certain extent, able to include electron correlation. The proper reference for discussing electron correlation is therefore a restricted (RHF) or restricted open-shell (ROHF) wave function, although many authors use a UHF wave function for open-shell species. In the RHF case, all the electrons are paired in molecular orbitals. The two electrons in an MO occupy the same physical space, and differ only in the spin function. The spatial overlap between the orbitals of two such “pair”-electrons is (exactly) one, while the overlap between two electrons belonging to different pairs is (exactly) zero, owing to the orthonormality of the MOs. The latter is not the same as saying that there is no repulsion between electrons in different MOs, since the electron–electron repulsion integrals involve products of MOs ($\langle \phi_i | \phi_j \rangle = 0$ for $i \neq j$, but $\langle \phi_i \phi_j | g | \phi_i \phi_j \rangle$ and $\langle \phi_i \phi_j | g | \phi_j \phi_i \rangle$ are not necessarily zero).

Naively it may be expected that the correlation between pairs of electrons belonging to the same spatial MO would be the major part of the electron correlation. However, as the size of the molecule increases, the number of electron pairs belonging to different spatial MOs grows faster than those belonging to the same MO. Consider for example the valence orbitals for CH₄. There are four *intraorbital* electron pairs of opposite spins, but there are twelve *interorbital* pairs of *opposite* spins, and twelve *interorbital* pairs of *same* spin. A typical value for the intraorbital pair correlation of a single bond is ~80 kJ/mol, while that of an interorbital pair (where the two

MOs are spatially close, as in CH_4) is $\sim 8\text{kJ/mol}$. The interpair correlation is therefore often comparable to the intrapair contribution.

Since the correlation between opposite spins has both intra- and interorbital contributions, it will be larger than the correlation between electrons having the same spin. The Pauli principle (or, equivalently, the antisymmetry of the wave function) has the consequence that there is no intraorbital correlation from electron pairs with the same spin. The opposite spin correlation is sometimes called the *Coulomb correlation*, while the same spin correlation is called the *Fermi correlation*, i.e. the Coulomb correlation is the largest contribution. Another way of looking at electron correlation is in terms of the electron density. In the immediate vicinity of an electron, there is a reduced probability of finding another electron. For electrons of opposite spin, this is often referred to as the *Coulomb hole*, and the corresponding phenomenon for electrons of same spin is the *Fermi hole*. This hole picture is discussed in more detail in connection with density functional theory in Chapter 6.

Another distinction is between *dynamic* and *static* electron correlation. The dynamic contribution is associated with the “instant” correlation between electrons, such as between those occupying the same spatial orbital. The static part is associated with electrons avoiding each other on a more “permanent” basis, such as those occupying different spatial orbitals. The latter is also sometimes called a near-degeneracy effect, as it becomes important for systems where different orbitals (configurations) have similar energies. The electron correlation in a helium atom is almost purely dynamic, while the correlation in the H_2 molecule at the dissociation limit is purely static (here the bonding and antibonding MOs become degenerate). At the equilibrium distance for H_2 the correlation is mainly dynamic (resembles the He atom), but this gradually changes to static correlation as the bond distance is increased. Similarly, the Be atom contains both static (near degeneracy of the $1s^2 2s^2$ and $1s^2 2p^2$ configurations) and dynamical correlation. There is therefore no clear-cut way of separating the two types of correlation, although they form a conceptually useful way of thinking about correlation effects.

The HF method determines *the* energetically best one-determinant trial wave function (within the given basis set). It is therefore clear that, in order to improve on HF results, the starting point must be a trial wave function that contains more than one Slater determinant (SD) Φ . This also means that the mental picture of electrons residing in orbitals has to be abandoned and the more fundamental property, the electron density, should be considered. As the HF solution usually gives $\sim 99\%$ of the correct answer, electron correlation methods normally use the HF wave function as a starting point for improvements.

A generic multi-determinant trial wave function can be written as in eq. (4.1), where a_0 is usually close to one.

$$\Psi = a_0 \Phi_{\text{HF}} + \sum_{i=1} a_i \Phi_i \quad (4.1)$$

Electron correlation methods differ in how they calculate the coefficients in front of the other determinants, with a_0 being determined by the normalization condition.

As mentioned in Chapter 5, one can think of the expansion of an unknown MO in terms of basis functions as describing the MO “function” in the “coordinate system” of the basis functions. The multi-determinant wave function (eq. (4.1)) can similarly be

considered as describing the total wave function in a “coordinate” system of Slater determinants. The *basis set* determines the size of the *one-electron basis* (and thus limits the description of the one-electron functions, the MOs), while the *number of determinants* included determines the size of the *many-electron basis* (and thus limits the description of electron correlation).

$$\begin{aligned}\chi &\rightarrow \phi \rightarrow \Phi \rightarrow \Psi \\ \text{AO} &\rightarrow \text{MO} \rightarrow \text{SD} \rightarrow \text{ME} \\ \phi &= \sum_{\alpha=1} c_{\alpha} \chi_{\alpha} \\ \Psi &= \sum_{i=1} a_i \Phi_i\end{aligned}$$

Figure 4.1 Progression from atomic orbitals (AO) (basis functions), to molecular orbitals (MO), to Slater determinants (SD) and to a many-electron (ME) wave function

4.1 Excited Slater Determinants

The starting point is usually an RHF calculation, where a solution of the Roothaan–Hall equations for a system with N electrons and M basis functions will yield $1/2 N_{\text{elec}}$ occupied MOs and $M_{\text{basis}} - 1/2 N_{\text{elec}}$ unoccupied (virtual) MOs. Except for a minimum basis, there will always be more virtual than occupied MOs. A Slater determinant is constructed from $1/2 N_{\text{elec}}$ spatial MOs multiplied by the two spin functions to yield N_{elec} spin-orbitals. A whole series of determinants may be generated by replacing MOs that are occupied in the HF determinant by MOs that are unoccupied. These can be denoted according to how many occupied HF MOs have been replaced by unoccupied MOs, i.e. Slater determinants that are *singly*, *doubly*, *triply*, *quadruply*, etc., *excited* relative to the HF determinant, up to a maximum of N_{elec} excited electrons. These determinants are often referred to as *Singles* (S), *Doubles* (D), *Triples* (T), *Quadruples* (Q), etc.

The total number of determinants that can be generated depends on the size of the basis set: the larger the basis, the more virtual MOs, and the more excited




Figure 4.2 Excited Slater determinants generated from an HF reference

determinants can be constructed. If all possible determinants in a given basis set are included, all the electron correlation (in the given basis) is (or can be) recovered. For an infinite basis, the Schrödinger equation is then solved exactly. Note that “exact” in this context is not the same as the experimental value, as the nuclei are assumed to have infinite masses (Born–Oppenheimer approximation) and relativistic effects are neglected. *Methods that include electron correlation are thus two-dimensional*, the larger the one-electron expansion (basis set size) and the larger the many-electron expansion (number of determinants), the better the results. This is illustrated in Figure 4.3.




Figure 4.3 Convergence to the exact solution

In order to calculate total energies with a “chemical accuracy” of ~ 4 kJ/mol (~ 1 kcal/mol), it is necessary to use sophisticated methods for including electron correlation and large basis sets, which only is computationally feasible for small systems. The focus is therefore on calculating relative energies, where error cancellation can improve the accuracy of the calculated results. The important chemical changes take place in the valence orbitals, with the core orbitals being almost independent of the molecular environment. In many cases, the interest is therefore only in calculating the correlation energy associated with the valence electrons. Limiting the number of determinants to only those that can be generated by exciting the valence electrons is known as the *frozen-core* approximation. In some cases, the highest virtual orbitals corresponding to the antibonding combinations of the core orbitals are also removed from the correlation treatment (*frozen virtuals*). The frozen-core approximation is not justified in terms of total energy; the correlation of the core electrons gives a substantial energy contribution. However, it is essentially a constant factor, which drops out when calculating relative energies. Furthermore, if we really want to calculate the core electron correlation, the standard basis sets are insufficient. In order to represent the angular correlation, higher angular momentum functions with the same radial size as the filled orbitals are needed, e.g. *p*- and *d*-functions with large exponents for correlating the 1s-electrons, as discussed in Section 5.4.6. Just allowing excitations of the core electrons in a standard basis set does not “correlate” the core electrons.

There are three main methods for calculating electron correlation: *Configuration Interaction* (CI), *Many-Body Perturbation Theory* (MBPT) and *Coupled Cluster* (CC). A word of caution before we describe these methods in more details. The Slater determinants are composed of spin-MOs, but since the Hamiltonian operator is independent of spin, the spin dependence can be factored out. Furthermore, to facilitate notation, it is often assumed that the HF determinant is of the RHF type, rather than the more general UHF type. Finally, many of the expressions below involve double summations over identical sets of functions. To ensure only the unique terms are included, one of the summation indices must be restricted. Alternatively, both indices can be allowed to run over all values, and the overcounting corrected by a factor of $1/2$. Various combinations of these assumptions result in final expressions that differ by factors of $1/2$, $1/4$, etc., from those given here. In the present chapter, the MOs are always spin-MOs, and conversion of a restricted summation to unrestricted is always noted explicitly.

Finally a comment on notation. The quality of a calculation is given by the level of theory (i.e. how much electron correlation is included) and the size of the basis set. In a commonly used “/”-notation, introduced by J. A. Pople, this is denoted as “*level/basis*”. If nothing further is specified, this implies that the geometry is optimized at this level of theory. As discussed in Section 5.7, the geometry is usually much less sensitive to the theoretical level than relative energies, and high-level calculations are therefore often carried out using geometries optimized at a lower level. This is denoted as “*level2/basis2//level1/basis1*”, where the notation after the “//” indicates the level at which the geometry is optimized.

4.2 Configuration Interaction

This is the oldest and perhaps the easiest method to understand, and is based on the variational principle (Appendix B), analogous to the HF method. The trial wave function is written as a linear combination of determinants with the expansion coefficients determined by requiring that the energy should be a minimum (or at least stationary), a procedure known as *Configuration Interaction* (CI).² The MOs used for building the excited Slater determinants are taken from a Hartree–Fock calculation and held fixed. Subscripts S, D, T, etc., indicate determinants that are Singly, Doubly, Triply, etc., excited relative to the HF configuration.

$$\Psi_{\text{CI}} = a_0 \Phi_{\text{HF}} + \sum_{\text{S}} a_{\text{S}} \Phi_{\text{S}} + \sum_{\text{D}} a_{\text{D}} \Phi_{\text{D}} + \sum_{\text{T}} a_{\text{T}} \Phi_{\text{T}} + \cdots = \sum_{i=0} a_i \Phi_i \quad (4.2)$$

This is an example of a constrained optimization, the energy should be minimized under the constraint that the total CI wave function is normalized. Introducing a Lagrange multiplier (Section 12.5), this can be written as

$$L = \langle \Psi_{\text{CI}} | \mathbf{H} | \Psi_{\text{CI}} \rangle - \lambda (\langle \Psi_{\text{CI}} | \Psi_{\text{CI}} \rangle - 1) \quad (4.3)$$

The first bracket is the energy of the CI wave function and the second bracket is the norm of the wave function. In terms of determinants (eq. (4.2)), these can be written as in eq. (4.4).

$$\begin{aligned}
\langle \Psi_{\text{CI}} | \mathbf{H} | \Psi_{\text{CI}} \rangle &= \sum_{i=0} \sum_{j=0} a_i a_j \langle \Phi_i | \mathbf{H} | \Phi_j \rangle = \sum_{i=0} a_i^2 E_i + \sum_{i=0} \sum_{j \neq i} a_i a_j \langle \Phi_i | \mathbf{H} | \Phi_j \rangle \\
\langle \Psi_{\text{CI}} | \Psi_{\text{CI}} \rangle &= \sum_{i=0} \sum_{j=0} a_i a_j \langle \Phi_i | \Phi_j \rangle = \sum_{i=0} a_i^2 \langle \Phi_i | \Phi_i \rangle = \sum_{i=0} a_i^2
\end{aligned} \tag{4.4}$$

The diagonal elements in the sum involving the Hamiltonian operator are energies of the corresponding determinants. The overlap elements between different determinants are zero as they are built from orthogonal MOs (eq. (3.20)). The variational procedure corresponds to setting all the derivatives of the Lagrange function (eq. (4.3)) with respect to the a_i expansion coefficients equal to zero.

$$\begin{aligned}
\frac{\partial L}{\partial a_i} &= 2 \sum_j a_j \langle \Phi_i | \mathbf{H} | \Phi_j \rangle - 2 \lambda a_i = 0 \\
a_i (E_i - \lambda) + \sum_{j \neq 0} a_j \langle \Phi_i | \mathbf{H} | \Phi_j \rangle &= 0
\end{aligned} \tag{4.5}$$

If there is only one determinant in the expansion ($a_0 = 1$, $a_{i \neq 0} = 0$), the latter equation shows that the Lagrange multiplier λ is the (CI) energy.

As there is one equation (4.5) for each i , the variational problem is transformed into solving a set of CI *secular equations*. Introducing the notation $H_{ij} = \langle \Phi_i | \mathbf{H} | \Phi_j \rangle$, the stationary conditions in eq. (4.5) can be written as in eq. (4.6)

$$\begin{pmatrix} H_{00} - E & H_{01} & \cdots & H_{0j} & \cdots \\ H_{10} & H_{11} - E & \cdots & H_{1j} & \cdots \\ \vdots & \vdots & \ddots & \vdots & \cdots \\ H_{j0} & \vdots & \cdots & H_{jj} - E & \cdots \\ \vdots & \vdots & \cdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_j \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \vdots \end{pmatrix} \tag{4.6}$$

This can also be written as a matrix equation.

$$\begin{aligned}
(\mathbf{H} - E\mathbf{I})\mathbf{a} &= \mathbf{0} \\
\mathbf{H}\mathbf{a} &= E\mathbf{a}
\end{aligned} \tag{4.7}$$

Solving the secular equations is equivalent to diagonalizing the CI matrix (see Section 16.2.4). The CI energy is obtained as the lowest eigenvalue of the CI matrix, and the corresponding eigenvector contains the a_i coefficients in front of the determinants in eq. (4.2). The second lowest eigenvalue corresponds to the first excited state, etc.

4.2.1 CI Matrix elements

The CI matrix elements H_{ij} can be evaluated by the strategy employed for calculating the energy of a single determinant used for deriving the Hartree–Fock equations (Section 3.3). This involves expanding the determinants in a sum of products of MOs, thereby making it possible to express the CI matrix elements in terms of MO integrals. There are, however, some general features that make many of the CI matrix elements equal to zero.

The Hamiltonian operator (eq. (3.23)) does not contain spin, thus if two determinants have different total spin the corresponding matrix element is zero. This situation

occurs if an electron is excited from an α spin-MO to a β spin-MO, such as the second S-type determinant in Figure 4.2. When the HF wave function is a singlet, this excited determinant is (part of) a triplet. The corresponding CI matrix element can be written in terms of integrals over MOs, and the spin dependence can be separated out. If there is a different number of α and β spin-MOs, there will always be at least one integral $\langle\alpha|\beta\rangle = 0$. That matrix elements between different spin states are zero may be fairly obvious. If we are interested in a singlet wave function, only singlet determinants can enter the expansion with non-zero coefficients. However, if the Hamiltonian operator includes for example the spin-orbit operator, matrix elements between singlet and triplet determinants are not necessarily zero, and the resulting CI wave function will be a mixture of singlet and triplet determinants.

Consider now the case where an electron with α spin is moved from orbital i to orbital a . The first S-type determinant in Figure 4.2 is of this type. Alternatively, the electron with β spin could be moved from orbital i to orbital a . Both of these excited determinants will have an S_z value of 0, but neither are eigenfunctions of the S^2 operator. The difference and sum of these two determinants describe a singlet state and the $S_z = 0$ component of a triplet, as illustrated in Figure 4.4.




Figure 4.4 Forming configurational state functions from Slater determinants

Such linear combinations of determinants, which are proper spin eigenfunctions, are called *Spin-Adapted Configurations* (SAC) or *Configurational State Functions* (CSF). The construction of proper CSFs may involve several determinants for higher excited states. The first D-type determinant in Figure 4.2 is already a proper CSF, but the second D-type excitation must be combined with five other determinants corresponding to rearrangement of the electron spins to make a singlet CSF (actually there are two linearly independent CSFs that can be made). By making suitable linear combinations of determinants the number of non-zero CI matrix elements can therefore be reduced.

If the system contains symmetry, there are additional CI matrix elements that become zero. The symmetry of a determinant is given as the direct product of the symmetries of the MOs. The Hamiltonian operator always belongs to the totally symmetric representation, thus if two determinants belong to different irreducible representations, the CI matrix element is zero. This is again fairly obvious: if the interest is in a state of a specific symmetry, only those determinants that have the correct symmetry can contribute.

The Hamiltonian operator consists of a sum of one-electron and two-electron operators, eq. (3.24). If two determinants differ by more than two (spatial) MOs there will always be an overlap integral between two different MOs that is zero (same argument as in eq. (3.28)). CI matrix elements can therefore only be non-zero if the two

determinants differ by 0, 1 or 2 MOs, and they may be expressed in terms of integrals of one- and two-electron operators over MOs. These connection are known as the *Slater–Condon* rules. If the two determinants are identical, the matrix element is simply the energy of a single-determinant wave function, as given by eq. (3.32). For matrix elements between determinants differing by 1 (exciting an electron from orbital i to a) or 2 (exciting two electrons from orbitals i and j to orbitals a and b) MOs, the results are given in eq. (4.8) (compare with eq. (3.33), where the \mathbf{g} operator is implicit in the notation for the two-electron integrals (eq. (3.57)).

$$\begin{aligned}\langle \Phi_0 | \mathbf{H} | \Phi_i^a \rangle &= \langle \phi_i | \mathbf{h} | \phi_a \rangle + \sum_j (\langle \phi_i \phi_j | \phi_a \phi_j \rangle - \langle \phi_i \phi_j | \phi_j \phi_a \rangle) \\ \langle \Phi_0 | \mathbf{H} | \Phi_{ij}^{ab} \rangle &= \langle \phi_i \phi_j | \phi_a \phi_b \rangle - \langle \phi_i \phi_j | \phi_b \phi_a \rangle\end{aligned}\quad (4.8)$$

The matrix element between the HF and a singly excited determinant is a matrix element of the Fock operator (eq. (3.36)) between two different MOs.

$$\langle \phi_i | \mathbf{h} | \phi_a \rangle = \sum_j (\langle \phi_i \phi_j | \phi_a \phi_j \rangle - \langle \phi_i \phi_j | \phi_j \phi_a \rangle) = \langle \phi_i | \mathbf{F} | \phi_a \rangle \quad (4.9)$$

This is an occupied–virtual off-diagonal element of the Fock matrix in the MO basis, and is identical to the gradient of the energy with respect to an occupied–virtual mixing parameter (except for a factor of 4), see eq. (3.68). If the determinants are constructed from optimized canonical HF MOs, the gradient is zero, and the matrix element is zero. This may also be realized by noting that the MOs are eigenfunctions of the Fock operator, eq. (3.42).

$$\begin{aligned}\mathbf{F}\phi_a &= \varepsilon_a \phi_a \\ \langle \phi_i | \mathbf{F} | \phi_a \rangle &= \varepsilon_a \langle \phi_i | \phi_a \rangle = \varepsilon_a \delta_{ia}\end{aligned}\quad (4.10)$$

The disappearance of matrix elements between the HF reference and singly excited states is known as *Brillouin’s theorem*. The HF reference state therefore only has non-zero matrix elements with doubly excited determinants, and the full CI matrix acquires a block diagonal structure.

In order to evaluate the CI matrix elements we need one- and two-electron integrals over MOs. These can be expressed in terms of the corresponding AO integrals and the MO coefficients.

$$\begin{pmatrix} \Phi_{\text{HF}} & \Phi_{\text{S}} & \Phi_{\text{D}} & \Phi_{\text{T}} & \Phi_{\text{Q}} & \Phi_{\text{5}} & \dots \end{pmatrix} \begin{pmatrix} \Phi_{\text{HF}} \\ \Phi_{\text{S}} \\ \Phi_{\text{D}} \\ \Phi_{\text{T}} \\ \Phi_{\text{Q}} \\ \Phi_{\text{5}} \\ \vdots \end{pmatrix} \begin{pmatrix} E_{\text{HF}} & 0 & x & 0 & 0 & 0 & 0 \\ 0 & E_{\text{S}} & x & x & 0 & 0 & 0 \\ x & x & E_{\text{D}} & x & x & 0 & 0 \\ 0 & x & x & E_{\text{T}} & x & x & 0 \\ 0 & 0 & x & x & E_{\text{Q}} & x & x \\ 0 & 0 & 0 & x & x & E_{\text{5}} & x \\ 0 & 0 & 0 & 0 & x & x & \ddots \end{pmatrix}$$

Figure 4.5 Structure of the CI matrix

$$\begin{aligned}
\langle \phi_i | \mathbf{h} | \phi_a \rangle &= \sum_{\alpha}^{M_{\text{basis}}} \sum_{\beta}^{M_{\text{basis}}} c_{a\alpha} c_{\beta i} \langle \chi_{\alpha} | \mathbf{h} | \chi_{\beta} \rangle \\
\langle \phi_i \phi_j | \phi_k \phi_l \rangle &= \sum_{\alpha}^{M_{\text{basis}}} \sum_{\beta}^{M_{\text{basis}}} \sum_{\gamma}^{M_{\text{basis}}} \sum_{\delta}^{M_{\text{basis}}} c_{a\alpha} c_{\beta j} c_{\gamma k} c_{\delta l} \langle \chi_{\alpha} \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle
\end{aligned} \tag{4.11}$$

Such MO integrals are required for all electron correlation methods. The two-electron AO integrals are the most numerous and the above equation appears to involve a computational effort proportional to M_{basis}^8 (M_{basis}^4 AO integrals each multiplied by four sets of M_{basis} MO coefficients). However, by performing the transformation one index at a time, the computational effort can be reduced to M_{basis}^5 .

$$\langle \phi_i \phi_j | \phi_k \phi_l \rangle = \sum_{\alpha}^{M_{\text{basis}}} c_{a\alpha} \left(\sum_{\beta}^{M_{\text{basis}}} c_{\beta j} \left(\sum_{\gamma}^{M_{\text{basis}}} c_{\gamma k} \left(\sum_{\delta}^{M_{\text{basis}}} c_{\delta l} \langle \chi_{\alpha} \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle \right) \right) \right) \tag{4.12}$$

Each step now only involves multiplication of M_{basis}^4 integrals with M_{basis} coefficients, i.e. the M_{basis}^8 dependence is reduced to four M_{basis}^5 operations. *In the large basis set limit, all electron correlation methods formally scale as at least M_{basis}^5* , since this is the scaling for the AO to MO integral transformation. The transformation is an example of a “rotation” of the “coordinate” system consisting of the AOs, to one where the Fock operator is diagonal, the MOs (see Section 16.2). The diagonal system allows a much more compact representation of the matrix elements needed for the electron correlation treatment. The coordinate change is also known as a *four index transformation*, since it involves four indices associated with the basis functions.

4.2.2 Size of the CI matrix

The excited Slater determinants are generated by removing electrons from occupied orbitals, and placing them in virtual orbitals. The number of excited SDs is thus a combinatorial problem, and therefore increases *factorially* with the number of electrons and basis functions. Consider for example a system such as H₂O with a 6-31G(d) basis. For the purpose of illustration, let us for a moment return to the spin-orbital description. There are 10 electrons and 38 spin-MOs, of which 10 are occupied and 28 are empty. There are $K_{10,n}$ possible ways of selecting n electrons out of the 10 occupied orbitals, and $K_{28,n}$ ways of distributing them in the 28 empty orbitals. The number of excited states for a given excitation level is thus $K_{10,n} \cdot K_{28,n}$, and the total number of excited determinants will be a sum over 10 such terms. This is also equivalent to $K_{38,10}$, the total number of ways 10 electrons can be distributed in 38 orbitals.

$$\text{Number of SDs} = \sum_{n=0}^{10} K_{10,n} \cdot K_{28,n} = K_{38,10} = \frac{38!}{10!(38-10)!} \tag{4.13}$$

Many of these excited determinants will of course have different spin multiplicity (triplet, pentet, etc., states for a singlet HF determinant), and can therefore be left out in the calculation. Generating only the singlet CSFs, the number of configurations at each excitation level is shown in Table 4.1.

The number of determinants (or CSFs) that can be generated grows wildly with the excitation level! Even if the C_{2v} symmetry of H₂O is employed, there is still a total of

Table 4.1 Number of singlet CSFs as a function of excitation level for H₂O with a 6-31G(d) basis

Excitation level n	Number of n th excited CSF	Total number of CSFs
1	71	71
2	2 485	2 556
3	40 040	42 596
4	348 530	391 126
5	1 723 540	2 114 666
6	5 033 210	7 147 876
7	8 688 680	15 836 556
8	8 653 645	24 490 201
9	4 554 550	29 044 751
10	1 002 001	30 046 752

7 536 400 singlet CSFs with A₁ symmetry. If all possible determinants are included, we have a *full CI* wave function and there is no truncation in the many-electron expansion besides that generated by the finite one-electron expansion (size of the basis set). This is the best possible wave function within the limitations of the basis set, i.e. it recovers 100% of the electron correlation in the given basis. For the water case with a medium basis set, this corresponds to diagonalizing a matrix of size 30 046 752 \times 30 046 752, which is impossible. However, normally the interest is only in the lowest (or a few of the lowest) eigenvalue(s) and eigenvector(s), and there are special iterative methods (Section 4.2.4) for determining one (or a few) eigenvector(s) of a large matrix.

In the general case of N electrons and M basis functions the total number of singlet CSFs that can be generated is given by eq. (4.14).

$$\text{Number of CSFs} = \frac{M!(M+1)!}{\left(\frac{N}{2}\right)! \left(\frac{N}{2}+1\right)! \left(M-\frac{N}{2}\right)! \left(M-\frac{N}{2}+1\right)!} \quad (4.14)$$

For H₂O with the above 6-31G(d) basis there are $\sim 30 \times 10^6$ CSFs ($N = 10$, $M = 19$), and with the larger 6-311G(2d,2p) basis there are $\sim 106 \times 10^9$ CSFs ($N = 10$, $M = 41$). For H₂C=CH₂ with the 6-31G(d) basis there are $\sim 334 \times 10^{12}$ CSFs ($N = 16$, $M = 38$).

One of the recent large-scale full CI calculations considered H₂O in a DZP type basis with 24 functions. Allowing all possible excitations of the 10 electrons generates 451 681 246 determinants.³ The variational wave function thus contains roughly half a billion parameters, i.e. the formal size of the CI matrix is of the order of half a billion squared. Although a determination of the lowest eigenvalue of such a problem can be done in a matter of hours on a modern computer, the result is a single number, the ground state energy of the H₂O molecule. Due to basis set limitations, however, it is still some 0.2 au (~ 500 kJ/mol) larger than the experimental value. The computational effort for extracting a single eigenvalue and eigenvector scales essentially linearly with the number of CSFs, and it is possible to handle systems with up to a few billion

determinants. The factorial growth of the number of determinants with the size of the basis set, however, makes the full CI method unfeasible for all but the very smallest systems. Full CI calculations are thus *not* a routine computational procedure for including electron correlation, but they are a very useful reference for developing more approximate methods, as the full CI gives the best results that can be obtained in the given basis.

4.2.3 Truncated CI methods

In order to develop a computationally tractable model, the number of excited determinants in the CI expansion (eq. (4.2)) must be reduced. Truncating the excitation level at one (*CI with Singles* (CIS)) does not give any improvement over the HF result as all matrix elements between the HF wave function and singly excited determinants are zero. CIS is equal to HF for the ground state energy, although higher roots from the secular equations may be used as approximations to excited states. It has already been mentioned that only doubly excited determinants have matrix elements with the HF wave function different from zero, thus the lowest CI level that gives an improvement over the HF result is to include only doubly excited states, yielding the *CI with Doubles* (CID) model. Compared with the number of doubly excited determinants, there are relatively few singly excited determinants (see for example Table 4.1), and including these gives the CISD method. Computationally, this is only a marginal increase in effort over CID. Although the singly excited determinants have zero matrix elements with the HF reference, they enter the wave function indirectly as they have non-zero matrix elements with the doubly excited determinants. In the large basis set limit the CISD method scales as M_{basis}^6 .

The next level in improvement is inclusion of the triply excited determinants, giving the CISDT method, which is an M_{basis}^8 method. Taking into account also quadruply excited determinants yields the CISDTQ method which is an M_{basis}^{10} method. As shown below, the CISDTQ model in general gives results close to the full CI limit, but even truncating the excitation level at four produces so many configurations that it can only be applied to small molecules and small basis sets. The only CI method that is generally applicable for a large variety of systems is CISD. For computationally feasible systems (i.e. medium-size molecules and basis sets), it typically recovers 80–90% of the available correlation energy. The percentage is highest for small molecules; as the molecule gets larger the CISD method recovers less and less of the correlation energy, which is discussed in more detail in Section 4.5.

Since only doubly excited determinants have non-zero matrix elements with the HF state, these are the most important. This may be illustrated by considering a full CI calculation for the Ne atom in a [5s4p3d] basis, where the 1s-electrons are omitted from the correlation treatment.⁴ The contribution to the full CI wave function from each level of excitation is given in Table 4.2.

The weight is the sum of a_i^2 coefficients at the given excitation level, eq. (4.2). The CI method determines the coefficients from the variational principle, thus Table 4.2 shows that the doubly excited determinants are by far the most important in terms of energy. The singly excited determinants are the second most important, followed by the quadruples and triples. Excitations higher than four make only very small

Table 4.2 Weights of excited configurations for the neon atom

Excitation level	Weight
0	9.6×10^{-1}
1	9.8×10^{-4}
2	3.4×10^{-2}
3	3.7×10^{-4}
4	4.5×10^{-4}
5	1.9×10^{-5}
6	1.7×10^{-6}
7	1.4×10^{-7}
8	1.1×10^{-9}

contributions, although there are actually many more of these highly excited determinants than the triples and quadruples, as illustrated in Table 4.1.

The relative importance of the different excitations may qualitatively be understood by noting that the doubles provide electron correlation for electron pairs. Quadruply excited determinants are important as they primarily correspond to products of double excitations. The singly excited determinants allow inclusion of multi-reference character in the wave function, i.e. they allow the orbitals to “relax”. Although the HF orbitals are optimum for the single-determinant wave function, that is no longer the case when many determinants are included. The triply excited determinants are doubly excited relative to the singles, and can then be viewed as providing correlation for the “multi-reference” part of the CI wave function.

While singly excited states make relatively small contributions to the correlation energy of a CI wave function, they are very important when calculating properties (Chapter 10). Molecular properties measure how the wave function changes when a perturbation, such as an external electric field, is added. The change in the wave function introduced by the perturbation makes the MOs no longer optimal in the variational sense. The first-order change in the MOs is described by the off-diagonal elements in the Fock matrix, as these are essentially the gradient of the HF energy with respect to the MOs. In the absence of a perturbation, these are zero, as the HF energy is stationary with respect to an orbital variation (eq. (3.39)). As shown in eqs (4.8) and (4.9), the Fock matrix off-diagonal elements are CI matrix elements between the HF and singly excited states. For molecular properties, the singly excited states thus allow the CI wave function to “relax” the MOs, i.e. letting the wave function respond to the perturbation.

4.2.4 Direct CI methods

As illustrated above, even quite small systems at the CISD level result in millions of CSFs. The variational problem is to extract one or possibly a few of the lowest eigenvalues and eigenvectors of a matrix the size of millions squared. This cannot be done by standard diagonalization methods where all the eigenvalues are found. There are,

however, iterative methods for extracting one, or a few, eigenvalues and eigenvectors of a large matrix. The CI problem eq. (4.7) may be written as in eq. (4.15).

$$(\mathbf{H} - E\mathbf{I})\mathbf{a} = \mathbf{0} \quad (4.15)$$

The \mathbf{H} matrix contains the matrix element between the CSFs in the CI expansion, and the \mathbf{a} vector the expansion coefficients. The idea in iterative methods is to generate a suitable guess for the coefficient vector and calculate $(\mathbf{H} - E\mathbf{I})\mathbf{a}$. This will in general not be zero, and the deviation may be used for adding a correction to \mathbf{a} , forming an iterative algorithm. If the interested is in the lowest eigenvalue, a suitable start eigenvector may be one that only contains the HF configuration, i.e. $\{1,0,0,0, \dots\}$. Since the \mathbf{H} matrix elements are essentially two-electron integrals in the MO basis (eq. (4.8)), the iterative procedure may be formulated as integral driven, i.e. a batch of integrals are read in (or generated otherwise) and used directly in the multiplication with the corresponding \mathbf{a} -coefficients. The CI matrix is therefore not needed explicitly, only the effect of its multiplication with a vector containing the variational parameters, and storage of the entire matrix is avoided. This is the basis for being able to handle CI problems of almost any size, and is known as *direct CI*. Note that it is not “direct” in the sense used to describe the direct SCF method, where all the AO integrals are calculated as needed. The direct CI approach just assumes that the CI matrix elements (e.g. two-electron integrals in the MO basis) are available as required, traditionally stored in a file on a disk. There are several variations on how the \mathbf{a} -vector is adjusted in each iteration, and the most commonly used versions are based on the *Davidson algorithm*.⁵

4.3 Illustrating how CI Accounts for Electron Correlation, and the RHF Dissociation Problem

Consider the H_2 molecule in a minimum basis consisting of one s -function on each centre, χ_A and χ_B . An RHF calculation will produce two MOs, ϕ_1 and ϕ_2 , being the sum and difference of the two AOs. The sum of the two AOs is a bonding MO, with increased probability of finding the electrons between the two nuclei, while the difference is an antibonding MO, with decreased probability of finding the electrons between the two nuclei.




Figure 4.6 Molecular orbitals for H_2

The HF wave function will have two electrons in the lowest energy (bonding) MO.

$$\begin{aligned}\phi_1 &= (\chi_A + \chi_B) \\ \phi_2 &= (\chi_A - \chi_B) \\ \Phi_0 &= \begin{vmatrix} \phi_1(1) & \overline{\phi_1(1)} \\ \phi_1(2) & \overline{\phi_1(2)} \end{vmatrix}\end{aligned}\quad (4.16)$$

We have here neglected the normalization constants for both the MOs and the determinantal wave function. The bar above the MO indicates that the electron has a β spin function, no bar indicates an α spin function. In this basis, there are one doubly (Φ_1) and four singly excited Slater determinants (Φ_{2-5}).

$$\begin{aligned}\Phi_1 &= \begin{vmatrix} \phi_2(1) & \overline{\phi_2(1)} \\ \phi_2(2) & \overline{\phi_2(2)} \end{vmatrix} \\ \Phi_2 &= \begin{vmatrix} \phi_1(1) & \overline{\phi_2(1)} \\ \phi_1(2) & \overline{\phi_2(2)} \end{vmatrix} \\ \Phi_3 &= \begin{vmatrix} \overline{\phi_1(1)} & \phi_1(1) \\ \overline{\phi_1(2)} & \phi_1(2) \end{vmatrix} \\ \Phi_4 &= \begin{vmatrix} \phi_1(1) & \phi_2(1) \\ \phi_1(2) & \phi_2(2) \end{vmatrix} \\ \Phi_5 &= \begin{vmatrix} \overline{\phi_1(1)} & \overline{\phi_2(1)} \\ \overline{\phi_1(2)} & \overline{\phi_2(2)} \end{vmatrix}\end{aligned}\quad (4.17)$$

Configurations Φ_4 and Φ_5 are clearly the $S_z = 1$ and $S_z = -1$ components of a triplet state. The plus combination of Φ_2 and Φ_3 is the $S_z = 0$ component of the triplet, and the minus combination is a singlet configuration, Figure 4.4. The H_2 molecule belongs to the $D_{\infty h}$ point group, and the two MOs transform as the σ_g (ϕ_1) and σ_u (ϕ_2) representations. The singly excited CSF ($\Phi_2 - \Phi_3$) has overall Σ_u symmetry, while the HF (Φ_0) and doubly excited determinant (Φ_1) have Σ_g . The full 6×6 CI problem therefore blocks into a 2×2 block of singlet Σ_g states, a 1×1 block of singlet Σ_u , and a 3×3 block of triplet Σ_u states. Owing to the orthogonality of the spin functions, the triplet block is already diagonal.

The full CI for the $^1\Sigma_g$ states involves only two configurations, the reference HF and the doubly excited determinant.

$$\begin{pmatrix} ^1\Phi_{0,\Sigma_g} & ^1\Phi_{1,\Sigma_g} & ^1(\Phi_2 - \Phi_3)_{\Sigma_u} & ^3\Phi_{4,\Sigma_u} & ^3(\Phi_2 + \Phi_3)_{\Sigma_u} & ^3\Phi_{5,\Sigma_u} \end{pmatrix} \begin{pmatrix} ^1\Phi_{0,\Sigma_g} \\ ^1\Phi_{1,\Sigma_g} \\ ^1(\Phi_2 - \Phi_3)_{\Sigma_u} \\ ^3\Phi_{4,\Sigma_u} \\ ^3(\Phi_2 + \Phi_3)_{\Sigma_u} \\ ^3\Phi_{5,\Sigma_u} \end{pmatrix} \begin{pmatrix} E_0 & H_{01} & 0 & 0 & 0 & 0 \\ H_{01} & E_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & E_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & E_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & E_4 & 0 \\ 0 & 0 & 0 & 0 & 0 & E_5 \end{pmatrix}$$

Figure 4.7 Structure of the full CI matrix for the H_2 system in a minimum basis

$$\begin{aligned}\Phi_0 &= \phi_1(1)\overline{\phi_1(2)} - \overline{\phi_1(1)}\phi_1(2) = \phi_1\phi_1(\alpha\beta - \beta\alpha) \\ \Phi_1 &= \phi_2(1)\overline{\phi_2(2)} - \overline{\phi_2(1)}\phi_2(2) = \phi_2\phi_2(\alpha\beta - \beta\alpha)\end{aligned}\quad (4.18)$$

In eq. (4.18) the electron coordinate is given implicitly by the order in which the orbitals are written, i.e. $\phi_1\phi_1[\alpha\beta - \beta\alpha] = \phi_1(1)\phi_1(2)[\alpha(1)\beta(2) - \beta(1)\alpha(2)]$. Ignoring the spin functions (which may be integrated out since \mathbf{H} is spin independent), the determinants can be expanded in AOs.

$$\begin{aligned}\Phi_0 &= (\chi_A(1) + \chi_B(1))(\chi_A(2) + \chi_B(2)) = \chi_A\chi_A + \chi_B\chi_B + \chi_A\chi_B + \chi_B\chi_A \\ \Phi_1 &= (\chi_A(1) - \chi_B(1))(\chi_A(2) - \chi_B(2)) = \chi_A\chi_A + \chi_B\chi_B - \chi_A\chi_B - \chi_B\chi_A\end{aligned}\quad (4.19)$$

The first two terms on the right-hand side have both electrons on the same nuclear centre, and they describe *ionic* contributions to the wave function, H^+H^- . The later two terms describe *covalent* contributions to the wave function, $\text{H}\cdot\text{H}\cdot$. The HF wave function thus contains equal amounts of ionic and covalent contributions.

The full CI wave function may be written in terms of AOs as in eq. (4.20), with the optimum values of the a_0 and a_1 coefficients determined by the variational procedure.

$$\begin{aligned}\Psi_{\text{CI}} &= a_0\Phi_0 + a_1\Phi_1 \\ \Psi_{\text{CI}} &= (a_0 + a_1)(\chi_A\chi_A + \chi_B\chi_B) + (a_0 - a_1)(\chi_A\chi_B + \chi_B\chi_A)\end{aligned}\quad (4.20)$$

The HF wave function constrains both electrons to move in the same bonding orbital. By allowing the doubly excited state to enter the wave function, the electrons can better avoid each other, as the antibonding MO is now also available. The antibonding MO has a nodal plane (where $\phi_2 = 0$) perpendicular to the molecular axis (Figure 4.6), and the electrons are able to correlate their movements by being on opposite sides of this plane. This *left-right* correlation is a molecular equivalent of the atomic radial correlation discussed in Section 5.2.

Consider now the behaviour of the HF wave function Φ_0 (eq. (4.19)) as the distance between the two nuclei is increased toward infinity. Since the HF wave function is an equal mixture of ionic and covalent terms, the dissociation limit is 50% H^+H^- and 50% $\text{H}\cdot\text{H}\cdot$. In the gas phase, all bonds dissociate homolytically, and the ionic contribution should be 0%. The HF dissociation energy is therefore much too high. This is a general problem with RHF type wave functions: the constraint of doubly occupied MOs is inconsistent with breaking bonds to produce radicals. In order for an RHF wave function to dissociate correctly, an even-electron molecule must break into two even-electron fragments, each being in the lowest electronic state. Furthermore, the orbital symmetries must match. There are only a few covalently bonded systems that obey these requirements (the simplest example is HHe^+). The wrong dissociation limit for RHF wave functions has several consequences:

- (1) The energy for stretched bonds is too high. Most transition structures have partly formed/broken bonds, thus activation energies are too high at the RHF level.
- (2) The too steep increase in energy as a function of the bond length causes the minimum on a potential energy curve to occur too “early” for covalently bonded systems, and equilibrium bond lengths are too short at the RHF level.

- (3) The too steep increase in energy as a function of the bond length causes the curvature of the potential energy surface near the equilibrium to be too large, and vibrational frequencies, especially those describing bond stretching, are in general too high.
- (4) The wave function contains too much “ionic” character, and RHF dipole moments (and also atomic charges) are in general too large.

It should be noted that dative bonds, such as in metal complexes and charge transfer species, in general have RHF wave functions that dissociate correctly, and the equilibrium bond lengths in these cases are normally too long.

The dissociation problem is solved in the case of a full CI wave function in this minimum basis. As seen from eq. (4.20), the ionic term can be made to disappear by setting $a_1 = -a_0$. The full CI wave function generates the lowest possible energy (within the limitations of the chosen basis set) at all distances, with the optimum weights of the HF and doubly excited determinant determined by the variational principle. In the general case of a polyatomic molecule and large basis set, correct dissociation of all bonds can be achieved if the CI wave function contains all determinants generated by a full CI in the valence orbital space. The latter corresponds to a full CI if a minimum basis is employed, but is much smaller than a full CI if an extended basis is used.

4.4 The UHF Dissociation, and the Spin Contamination Problem

The dissociation problem can also be “solved” by using a wave function of the UHF type. Here the α and β bonding MOs are allowed to “localize”, thereby reducing the MO symmetries to $C_{\infty v}$.

$$\begin{aligned}\phi_1 &= (\chi_A + c\chi_B)\alpha \\ \bar{\phi}_1 &= (c\chi_A + \chi_B)\beta \\ \Phi_0^{\text{UHF}} &= \begin{vmatrix} \phi_1(1) & \bar{\phi}_1(1) \\ \phi_1(2) & \bar{\phi}_1(2) \end{vmatrix}\end{aligned}\tag{4.21}$$

The optimum value of c is determined by the variational principle. If $c = 1$, the UHF wave function is identical to RHF. This will normally be the case near the equilibrium distance. As the bond is stretched, the UHF wave function allows each of the electrons to localize on a nucleus, causing c to go towards 0. The point where the RHF and UHF descriptions start to differ is often referred to as the RHF/UHF *instability* point. This is an example of symmetry breaking, as discussed in Section 3.8.3. The UHF wave function correctly dissociates into two hydrogen atoms, but the symmetry breaking of the MOs has two other, closely connected, consequences: introduction of *electron correlation* and *spin contamination*. To illustrate these concepts, we need to look at the Φ_0 UHF determinant, and the six RHF determinants in eqs (4.16) and (4.17) in more detail. We will again ignore all normalization constants. The six RHF determinants can be expanded in terms of the AOs as in eq. (4.22).

$$\begin{aligned}
\Phi_0 &= [\chi_A \chi_A + \chi_B \chi_B + \chi_A \chi_B + \chi_B \chi_A](\alpha\beta - \beta\alpha) \\
\Phi_1 &= [\chi_A \chi_A + \chi_B \chi_B - \chi_A \chi_B - \chi_B \chi_A](\alpha\beta - \beta\alpha) \\
\Phi_2 &= [\chi_A \chi_A - \chi_B \chi_B](\alpha\beta - \beta\alpha) - [\chi_A \chi_B](\alpha\beta + \beta\alpha) \\
\Phi_3 &= [\chi_A \chi_A - \chi_B \chi_B](\alpha\beta - \beta\alpha) + [\chi_A \chi_B - \chi_B \chi_A](\alpha\beta + \beta\alpha) \\
\Phi_4 &= [\chi_A \chi_B - \chi_B \chi_A](\alpha\alpha) \\
\Phi_5 &= [\chi_A \chi_B - \chi_B \chi_A](\beta\beta)
\end{aligned} \tag{4.22}$$

Subtracting and adding Φ_2 and Φ_3 produces a pure singlet ($^1\Phi_-$) and the $S_z = 0$ component of the triplet ($^3\Phi_+$) wave function.

$$\begin{aligned}
^1\Phi_- &= \Phi_2 - \Phi_3 = [\chi_A \chi_A - \chi_B \chi_B](\alpha\beta - \beta\alpha) \\
^3\Phi_+ &= \Phi_2 + \Phi_3 = [\chi_A \chi_B - \chi_B \chi_A](\alpha\beta + \beta\alpha)
\end{aligned} \tag{4.23}$$

Performing the expansion of the Φ_0^{UHF} determinant (eq. (4.21)) gives eq. (4.24).

$$\begin{aligned}
\Phi_0^{\text{UHF}} &= c[\chi_A \chi_A + \chi_B \chi_B](\alpha\beta - \beta\alpha) + \\
&\quad [\chi_A \chi_B \alpha\beta - c^2 \chi_B \chi_A \beta\alpha] + \\
&\quad [c^2 \chi_B \chi_A \alpha\beta - \chi_A \chi_B \beta\alpha]
\end{aligned} \tag{4.24}$$

Adding and subtracting factors of $\chi_A \chi_B \alpha\beta$ and $\chi_B \chi_A \beta\alpha$ allow this to be written as in eq. (4.25).

$$\begin{aligned}
\Phi_0^{\text{UHF}} &= [c(\chi_A \chi_A + \chi_B \chi_B) + (\chi_A \chi_B + \chi_B \chi_A)](\alpha\beta - \beta\alpha) + \\
&\quad (1 - c^2)[\chi_A \chi_B \beta\alpha - \chi_B \chi_A \alpha\beta]
\end{aligned} \tag{4.25}$$

Since $0 \leq c \leq 1$, the first term shows that UHF orbitals reduce the ionic contribution relative to the covalent structures, compared with the RHF case, eq. (4.19). This is the same effect as for the CI procedure (eq. (4.20)), i.e. the first term shows that the UHF wave function partly includes electron correlation.

The first term in eq. (4.25) can be written as a linear combination of the Φ_0 and Φ_1 determinants, and describes a pure singlet state. The last part of the UHF determinant, however, has terms identical to two of those in the triplet $^3\Phi_+$ combination, eq. (4.23). If we had chosen the alternative set of UHF orbital with the α spin being primarily on centre B in eq. (4.21), we would have obtained the other two terms in $^3\Phi_+$, i.e. the last term in eq. (4.25) breaks the symmetry. The UHF determinant is therefore not a pure spin state, it contains both singlet and triplet spin states. This feature is known as *spin contamination*. For $c = 1$, the UHF wave function is identical to RHF, and Φ_0^{UHF} is a pure singlet. For $c = 0$, the UHF wave function only contains the covalent terms, which is the correct dissociation behaviour, but also contains equal amounts of singlet and triplet character. When the bond distance is very large, the singlet and triplet states have identical energies, and the spin contamination has no consequence for the energy. In the intermediate region where the bond is not completely broken, however, spin contamination is important.

Compared with full CI, the UHF energy is too high as the higher lying triplet state is mixed into the wave function. The variational principle guarantees that the UHF

energy is lower than or equal to the RHF energy since there are more variational parameters. The full CI energy is the lowest possible (for the given basis set) as it recovers 100% of the correlation energy. The UHF wave function thus lowers the energy by introducing some electron correlation, but at the same time raises the energy by including higher energy spin states. At the single-determinant level, the variational principle guarantees that the first effect dominates. If the second effect dominated, the UHF would collapse to the RHF solution. The correlation energy in general increases as a bond is stretched, and the instability point can thus be viewed as the geometry where the correlation effect becomes larger than the spin contamination. Pictorially, the dissociation curves appear as shown in Figure 4.8.




Figure 4.8 Bond dissociation curves for H_2

Another way of viewing spin contamination is to write the UHF wave function as a linear combination of pure R(O)HF determinants, e.g. for a singlet state.

$$^1\Phi^{\text{UHF}} = a_1^1\Phi^{\text{RHF}} + a_3^3\Phi^{\text{ROHF}} + a_5^5\Phi^{\text{ROHF}} + \dots \quad (4.26)$$

Since the UHF wave function is multi-determinantal in terms of R(O)HF determinants, it follows that it to some extent includes electron correlation (relative to the RHF reference).

The amount of spin contamination is given by the expectation value of the \mathbf{S}^2 operator, $\langle \mathbf{S}^2 \rangle$. The theoretical value for a pure spin state is $S_z(S_z + 1)$, i.e. 0 for a singlet ($S_z = 0$), 0.75 for a doublet ($S_z = 1/2$), 2.00 for a triplet ($S_z = 1$), etc. A UHF “singlet” wave function will contain some amounts of triplet, quintet, etc., states, increasing the $\langle \mathbf{S}^2 \rangle$ value from its theoretical value of zero for a pure spin state. Similarly, a UHF “doublet” wave function will contain some amounts of quartet, sextet, etc., states. Usually the contribution from the next higher spin state than the desired is the most

important. The $\langle S^2 \rangle$ value for a UHF wave function is operationally calculated from the spatial overlap between all pairs of α and β spin-orbitals.

$$\langle S^2 \rangle = S_z(S_z + 1) + N_\beta - \sum_{ij}^{N_{MO}} \langle \phi_i^\alpha | \phi_j^\beta \rangle^2 \quad (4.27)$$

If the α and β orbitals are identical, there is no spin contamination, and the UHF wave function is identical to RHF.

By including electron correlation in the wave function, the UHF method introduces more biradical character into the wave function than RHF. The spin contamination part is also purely biradical in nature, i.e. a UHF treatment in general will overestimate the biradical character of the wave function. Most singlet states are well described by a closed shell wave function near the equilibrium geometry and, in those cases, it is not possible to generate a UHF solution that has a lower energy than the RHF. There are systems, however, for which this does not hold. An example is the ozone molecule, where two types of resonance structures can be drawn.




Figure 4.9 Resonance structures for ozone

The biradical resonance structure for ozone requires two singly occupied MOs, and it is clear that an RHF type wave function, which requires all orbitals to be doubly occupied, cannot describe this. A UHF type wave function, however, allows the α and β orbitals to be spatially different, and can to a certain extent incorporate both resonance structures. Systems with biradical character will in general have a (singlet) UHF wave function different from an RHF.

As mentioned above, spin contamination in general increases as a bond is stretched. This has important consequences for transition structures, which contain elongated bonds. While most singlet systems have identical RHF and UHF descriptions near the equilibrium geometry, it will normally be possible to find a lower energy UHF solution in the TS region. However, since the spin contamination is not constant along the reaction coordinate, and since the UHF overestimates the biradical character, it is possible that the TS actually becomes a minimum on the UHF energy surface. In other words, the spin contamination may severely distort the shape of the potential energy surface. This may qualitatively be understood by considering the “singlet” UHF wave function as a linear combination of a singlet and a triplet states, as shown in Figure 4.10.

The degree of mixing is determined by the energy difference between the pure singlet and triplet states (as shown for example by second-order perturbation theory, see Section 4.8), which in general decreases along the reaction coordinate. Even if the mixing is not large enough to actually transform a TS to a minimum, it is clear that the UHF energy surface will be much too flat in the TS region. Activation energies calculated at the UHF level will always be lower than the RHF value, but may be either higher or lower than the “correct” value, depending on the amount of spin contamination, since RHF normally overestimates activation energies.




Figure 4.10 Mixing of pure singlet and triplet states may generate artificial minima on the UHF energy surface

From the above it should be clear that UHF wave functions that are spin contaminated (more than a few percent deviation of $\langle \mathbf{S}^2 \rangle$ from the theoretical value of $S_z(S_z + 1)$) have disadvantages. For closed shell systems, an RHF procedure is therefore normally preferred. For open-shell systems, however, the UHF method has been widely used. It is possible to use an ROHF type wave function for open-shell systems, but this leads to computational procedures that are somewhat more complicated than for the UHF case when electron correlation is introduced.

The main problem with the UHF method is the spin contamination, and there have been several proposals on how to remove these unwanted states. There are three strategies that can be considered for removing the contamination:

- During the SCF procedure.
- After the SCF has converged.
- After electron correlation has been added to the UHF solution.

A popular method of removing unwanted states is to project them out with a suitable projection operator (in the picture of the wave function being described in the coordinate system consisting of determinants, the components of the wave function along the higher spin states is removed). As mentioned above, the next higher spin state is usually the most important, and in many cases it is a quite good approximation to only remove this state. After projection, the wave function is then renormalized. If only the first contaminant is removed, this may in extreme cases actually increase the $\langle \mathbf{S}^2 \rangle$ value.

Performing the projection *during* the SCF procedure produces a wave function for which it is difficult to formulate a satisfactory theory for including electron correlation by means of perturbation or coupled cluster methods (Sections 4.8 and 4.9). Projections of the *converged* UHF wave function will lower the energy (although the projected UHF (PUHF) energy is no longer variational), since the contributions of the higher lying states are removed, and only the correlation effect remains. However, the problem of artificial distortion of the energy surface is even more pronounced at the PUHF level than with the UHF method itself. For example, it is often found that

a false minimum is generated just after the RHF/UHF instability point on a bond dissociation curve. Furthermore, the derivatives of the PUHF energy are not continuous at the RHF/UHF instability point. Projection of the wave function *after* electron correlation has been added, however, turns out to be a viable pathway. This has mainly been used in connection with perturbation methods, to be described in Section 4.8.2.

4.5 Size Consistency and Size Extensivity

As mentioned above, full CI is impossible except for very small system. The only generally applicable method is CISD. Consider now a series of CISD calculations in order to construct the interaction potential between two H_2 molecules as a function of the distance between them. Relative to the HF wave function, there will be determinants that correspond to single excitations on only one of the H_2 fragments (S-type determinants), single excitations on both (D-type determinants), and double excitations only on one of the H_2 fragments (also D-type determinants). This will be the case at all intermolecular distances, even when the separation is very large. In that case, however, the system is just two H_2 molecules, and we could consider calculating the energy instead as twice the energy of one H_2 molecule. A CISD calculation on one H_2 molecule would generate singly and doubly excited determinants, and multiplying this by two would generate determinants that are triply and quadruply excited for the combined H_4 system. A CISD calculation of two H_2 molecules separated by say 100 \AA will therefore not give the same energy as twice the results from a CISD calculation on one H_2 molecule (the latter will be lower). This problem is referred to a *Size Inconsistency*. A very similar, but not identical concept, is *Size Extensivity*. Size consistency is only defined if the two fragments are non-interacting (separated by say 100 \AA), while size extensivity implies that the method scales properly with the number of particles, i.e. the fragments can be interacting (separated by say 5 \AA). Full CI is size consistent (and extensive), but all forms of truncated CI are not. The lack of size extensivity is the reason why CISD recovers less and less electron correlation as the systems grow larger.

4.6 Multi-Configuration Self-Consistent Field

The *Multi-Configuration Self-Consistent Field* (MCSCF) method can be considered as a CI where not only are the coefficients in front of the determinants (eq. (4.2)) optimized by the variational principle, but the MOs used for constructing the determinants are also optimized.⁶ The MCSCF optimization is iterative like the SCF procedure (if the “multi configuration” is only one, it is simply HF). Since the number of MCSCF iterations required for achieving convergence tends to increase with the number of configurations included, the size of MCSCF wave functions that can be treated is somewhat smaller than for CI methods.

When deriving the HF equations only the *variation* of the energy with respect to an orbital variation was required to be zero, which is equivalent to the first derivatives of the energy with respect to the MO expansion coefficients being equal to zero. The HF equations can be solved by an iterative SCF method, and there are many techniques for helping the iterative procedure to converge (Section 3.8). There is, however, no

guarantee that the solution found by the SCF procedure is a *minimum* of the energy as a function of the MO coefficients. In order to ensure that a minimum has been found, the matrix of second derivatives of the energy with respect to the MO coefficients can be calculated and diagonalized, with a minimum having only positive eigenvalues. This is rarely checked for SCF wave functions; in the large majority of cases the SCF procedure converges to a minimum without problems. MCSCF wave functions, on the other hand, are much harder to converge, and much more prone to converge on solutions that are not minima. MCSCF wave function optimizations are therefore normally carried out by expanding the energy to second order in the variational parameters (orbital and configurational coefficients), analogously to the second-order SCF procedure described in Section 3.8.1, and using the Newton–Raphson-based methods described in Section 12.2.3 to force convergence to a minimum.

MCSCF methods are rarely used for calculating large fractions of the correlation energy. The orbital relaxation usually does not recover much electron correlation, and it is more efficient to include additional determinants and keep the MOs fixed (CI) if the interest is just in obtaining a large fraction of the correlation energy. Single-determinant HF wave functions normally give a qualitatively correct description of the electron structure, but there are many examples where this is not the case. MCSCF methods can be considered as an extension of single-determinant methods to give a qualitatively correct description.

Consider again the ozone molecule with the two resonance structures shown in Figure 4.9. Each type of resonance structure essentially translates into a different determinant. If more than one *non-equivalent* resonance structure is important, this means that the wave function cannot be described even qualitatively correctly at the RHF single-determinant level (benzene, for example, has two *equivalent* cyclohexatriene resonance structures, and is adequately described by an RHF wave function). A UHF wave function allows some biradical character, with the disadvantages discussed in Section 4.4. Alternatively, a second restricted type CSF (consisting of two determinants) with two singly occupied MOs may be included in the wave function. The simplest MCSCF for ozone contains two configurations (often denoted TCSCF), with the optimum MOs and configurational weights determined by the variational principle. The CSFs entering an MCSCF expansion are pure spin states, and MCSCF wave functions therefore do not suffer from the problem of spin contamination.

Our definition of electron correlation uses the RHF energy as the reference. For ozone, both the UHF and the TCSCF wave functions have lower energies, and include some electron correlation. This type of “electron correlation” is somewhat different from the picture presented at the start of this chapter. In a sense it is a consequence of our chosen zero point for the correlation energy, the RHF energy. The energy lowering introduced by adding enough flexibility in the wave function to be able to qualitatively describe the system is sometimes called the *static* electron correlation. This is essentially the effect of allowing orbitals to become (partly) singly occupied instead of forcing double occupation, i.e. describing near-degeneracy effects (two or more configurations having almost the same energy). The remaining energy lowering by correlating the motion of the electrons is called *dynamical* correlation. The problem is that there is no rigorous way of separating these effects. In the ozone example the energy lowering by going from RHF to UHF, or to a TCSCF, is almost pure static correlation. Increasing the number of configurations in an MCSCF will recover more and more of

the dynamical correlation, until, at the full CI limit, the correlation treatment is exact. As mentioned above, MCSCF methods are mainly used for generating a qualitatively correct wave function, i.e. recovering the “static” part of the correlation.

The major problem with MCSCF methods is selecting which configurations are necessary to include for the property of interest. One of the most popular approaches is the *Complete Active Space Self-Consistent Field* (CASSCF) method (also called *Full Optimized Reaction Space* (FORS)). Here the selection of configurations is done by partitioning the MOs into *active* and *inactive* spaces. The active MOs will typically be some of the highest occupied and some of the lowest unoccupied MOs from an RHF calculation. The inactive MOs have either 2 or 0 electrons, i.e. always either doubly occupied or empty. Within the active MOs a full CI is performed and all the proper symmetry-adapted configurations are included in the MCSCF optimization. Which MOs to include in the active space must be decided manually, by considering the problem at hand and the computational expense. If several points on the potential energy surface are desired, the MCSCF active space should include all those orbitals that change significantly, or for which the electron correlation is expected to change. A common notation is $[n,m]$ -CASSCF, which indicates that n electrons are distributed in all possible ways in m orbitals.

As for any full CI expansion, the CASSCF becomes unmanageably large even for quite small active spaces. A variation of the CASSCF procedure is the *Restricted Active Space Self-Consistent Field* (RASSCF) method.⁷ Here the active MOs are divided into three sections, RAS1, RAS2 and RAS3, each having restrictions on the occupation numbers (excitations) allowed. A typical model consists of the configurations in the RAS2 space being generated by a full CI (analogously to CASSCF), or perhaps limited to SDTQ excitations. The RAS1 space consists of MOs that are doubly occupied in the HF reference determinant, and the RAS3 space consists of MOs that are empty in the HF. Configurations additional to those from the RAS2 space are generated by allowing for example a maximum of two electrons to be excited *from* the RAS1 and




Figure 4.11 Illustrating the CAS and RAS orbital partitions

a maximum of two electrons to be excited *to* the RAS3 space. In essence, a typical RASSCF procedure thus generates configurations by a combination of a full CI in a small number of MOs (RAS2) and a CISD in a somewhat larger MO space (RAS1 and RAS3).

The full CI expansion within the active space severely restricts the number of orbitals and electrons that can be treated by CASSCF methods. Table 4.3 shows how many singlet CSFs are generated for an $[n,n]$ -CASSCF wave function (eq. (4.13)), without reductions arising from symmetry.

Table 4.3 Number of configurations generated in an $[n,n]$ -CASSCF wave function

n	Number of CSFs
2	3
4	20
6	175
8	1764
10	19404
12	226512
14	2760615

The factorial increase in the number of CSFs effectively limits the active space for CASSCF wave functions to fewer than 10–12 electrons/orbitals. Selecting the “important” orbitals to correlate therefore becomes very important. The goal of MCSCF methods is usually not to recover a large fraction of the total correlation energy, but rather to recover all the *changes* that occur in the correlation energy for the given process. Selecting the active space for an MCSCF calculation requires some insight into the problem. There are a few rules of thumb that may be of help in selecting a proper set of orbitals for the active space:

- (1) For each occupied orbital, there will typically be one corresponding virtual orbital. This leads naturally to $[n,m]$ -CASSCF wave functions where n and m are identical or nearly so.
- (2) Including all the valence orbitals, i.e. the space spanned by a minimum basis set, leads to a wave function that can correctly describe all dissociation pathways. Unfortunately, a full valence CASSCF wave function rapidly becomes unmanageably large for realistic-sized systems.
- (3) The orbital energies from an RHF calculation may be used for selecting the important orbitals. The highest occupied and lowest unoccupied are usually the most important orbitals to include in the active space. This can be partly justified by the formula for the second-order perturbation energy correction (Section 4.8.1): the smaller the orbital energy difference, the larger contribution to the correlation energy. Using RHF orbital energies for selecting the active space may be problematic in two situations. The first is when extended basis sets are used, where there will be many virtual orbitals with low energies, and the exact order is more or less accidental. Furthermore, RHF virtual orbitals basically describe electron attachment (via Koopmans’ theorem, Section 3.4), and are therefore not particularly well

suited for describing electron correlation. An inspection of the form of the orbitals may reveal which to choose: they should be the ones that resemble the occupied orbitals in terms of basis function contribution. The second problem is more fundamental. If the real wave function has significant multi-configurational character, then the RHF may be qualitatively wrong, and selecting the active orbitals based on a qualitatively wrong wave function may lead to erroneous results. The problem is that we wish to include the important orbitals for describing the multi-determinant nature, but these are not known until the final wave function is known.

- (4) An attempt to overcome this self-referencing problem is to use the concept of natural orbitals. The natural orbitals are those that diagonalize the density matrix, and the eigenvalues are the occupation numbers. Orbitals with occupation numbers significantly different from 0 or 2 (for a closed shell system) are usually those that are the most important to include in the active space. An RHF wave function will have occupation numbers of exactly 0 or 2, and some electron correlation must be included to obtain orbitals with non-integer occupation numbers. This may for example be done by running a preliminary MP2 or CISD calculation prior to the MCSCF. Alternatively, a UHF (when different from RHF) type wave function may also be used. The total UHF density, which is the sum of the α and β density matrices, will also provide fractional occupation numbers since UHF includes some electron correlation. The procedure may still fail. If the underlying RHF wave function is poor, the MP2 correction may also give poor results, and selecting the active MCSCF orbitals based on MP2 occupation number may again lead to erroneous results. In practice, however, selecting active orbitals based on for example MP2 occupation numbers appears to be quite efficient, and better than using RHF orbital energies.

In a CASSCF type wave function the CI coefficients do not have the same significance as for a single-reference CI based on HF orbitals. In a full CI (as in the active space of the CASSCF), the orbitals may be rotated among themselves without affecting the total wave function. A rotation of the orbitals, however, influences the magnitude of the coefficients in front of each CSF. While the HF coefficient in a single-reference CISD gives some indication of the “multi-reference” nature of the wave function, this is not the case for a CASSCF wave function, where the corresponding CI coefficient is arbitrary.

It should be noted that CASSCF methods inherently tend to give an unbalanced description, since all the electron correlation recovered is in the active space, with none in the inactive space, or between the active and inactive electrons.⁸ This is not a problem if all the valence electrons are included in the active space, but this is only possible for small systems. If only part of the valence electrons are included in the active space, the CASSCF method tends to overestimate the importance of “biradical” structures. Consider for example acetylene where the hydrogens have been bent 60° away from linearity (this may be considered a model for ortho-benzyne). The in-plane “ π -orbital” now acquires significant biradical character. The true structure may be described as a linear combination of the following three configurations.

The structure on the left is biradical, while the two others are ionic, corresponding to both electrons being at the same carbon. The simplest CASSCF wave function that can qualitatively describe this system has two electrons in two orbitals, giving the three




Figure 4.12 Important configurations for a bend acetylene model

configurations shown above. The dynamical correlation between the two active electrons will tend to keep them as far apart as possible, i.e. favouring the biradical structure. Consider now a full valence CASSCF wave function with ten electrons in ten orbitals. This will analogously tend to separate the two electrons in each bond with one being at each end. The correlation of the electrons in the C—H bonds, for example, will place more electron density on the carbon atoms. This in turn favours the ionic structures in Figure 4.12 and disfavors the biradical, i.e. the dynamical correlation of the other electrons may take advantage of the empty orbital in the ionic structures but not in the biradical structure. These general considerations may be quantified by considering the natural orbital occupancies for increasingly large CASSCF wave functions, as shown in Table 4.4 with the 6-31G(d,p) basis.

Table 4.4 Natural orbital occupation numbers for the distorted acetylene model in Figure 4.12; only the occupation numbers for the six “central” orbitals are shown

	n_5	n_6	n_7	n_8	n_9	n_{10}
RHF	2.00	2.00	2.00	0.00	0.00	0.00
UHF	2.00	1.72	1.30	0.70	0.28	0.01
[2,2]-CASSCF	2.00	2.00	1.62	0.38	0.00	0.00
[4,4]-CASSCF	2.00	1.85	1.67	0.33	0.14	0.00
[10,10]-CASSCF	1.97	1.87	1.71	0.30	0.13	0.02

The [4,4]-CASSCF also includes the two out-of-plane π -orbitals in the active space, while the [10,10]-CASSCF generates a full-valence CI wave function. The unbalanced description for the [2,2]-CASSCF is reminiscent of the spin contamination problem for UHF wave functions, although the effect is much less pronounced. Nevertheless, the overestimation may be severe enough to alter the qualitative shape of energy surfaces, for example turning transition structures into minima, as illustrated in Figure 4.10. MCSCF methods are therefore not “black box” methods such as for example HF and MP (Section 4.8.1); selecting a proper number of configurations, and the correct orbitals, to give a balanced description of the problem at hand requires some experimentation and insight.

4.7 Multi-Reference Configuration Interaction

The CI methods described so far consider only CSFs generated by exciting electrons from a single determinant. This corresponds to having an HF type wave function as the reference. However, an MCSCF wave function may also be chosen as the

reference. In that case, a CISD involves excitations of one or two electrons out of *all* the determinants that enter the MCSCF, defining the *Multi-Reference Configuration Interaction* (MRCI) method. Compared with the single-reference CISD, the number of configurations is increased by a factor roughly equal to the number of configurations included in the MCSCF. Large-scale MRCI wave functions (many configurations in the MCSCF) can generate very accurate wave functions, but are also computationally very intensive. Since MRCI methods truncate the CI expansion, they are not size extensive.

Even truncating the (MR) CI expansion at the singles and doubles level frequently generates more configurations than can be handled readily. A further truncation is sometimes performed by selecting only those configurations that have an “interaction” with the reference configuration(s) above a selected threshold, where the “interaction” is evaluated by second-order perturbation theory (Section 4.8). Such *state-selected* CI (or MCSCF) methods all involve a preset cutoff below which configurations are neglected. This may cause problems for comparing energies of different geometries, since the potential energy surface may become discontinuous, i.e. at some point the importance of a given configuration drops below the threshold, and the contribution suddenly disappears.

4.8 Many-Body Perturbation Theory

The idea in perturbation methods is that the problem at hand only differs slightly from a problem that has already been solved (exactly or approximately). The solution to the given problem should therefore in some sense be close to the solution to the already known system. This is described mathematically by defining a Hamiltonian operator that consists of two parts, a reference (\mathbf{H}_0) and a perturbation (\mathbf{H}'). The premise of perturbation methods is that the \mathbf{H}' operator in some sense is “small” compared with \mathbf{H}_0 . Perturbation methods can be used in quantum mechanics for adding corrections to solutions that employ an independent-particle approximation, and the theoretical framework is then called *Many-Body Perturbation Theory* (MBPT).

Let us assume that the Schrödinger equation for the reference Hamiltonian operator is solved.

$$\begin{aligned}\mathbf{H} &= \mathbf{H}_0 + \lambda \mathbf{H}' \\ \mathbf{H}_0 \Phi_i &= E_i \Phi_i \quad i = 0, 1, 2, \dots, \infty\end{aligned}\tag{4.28}$$

The solutions for the unperturbed Hamiltonian operator form a complete set (since \mathbf{H}_0 is Hermitian) which can be chosen to be orthonormal, and λ is a (variable) parameter determining the strength of the perturbation. At present, we will only consider cases where the perturbation is time-independent, and the reference wave function is non-degenerate. To keep the notation simple, we will furthermore only consider the lowest energy state. The perturbed Schrödinger equation is given by eq. (4.29).

$$\mathbf{H}\Psi = W\Psi\tag{4.29}$$

If $\lambda = 0$, then $\mathbf{H} = \mathbf{H}_0$, $\Psi = \Phi_0$ and $W = E_0$. As the perturbation is increased from zero to a finite value, the new energy and wave function must also change continuously, and they can be written as a Taylor expansion in powers of the perturbation parameter λ .

$$\begin{aligned}
W &= \lambda^0 W_0 + \lambda^1 W_1 + \lambda^2 W_2 + \lambda^3 W_3 + \cdots \\
\Psi &= \lambda^0 \Psi_0 + \lambda^1 \Psi_1 + \lambda^2 \Psi_2 + \lambda^3 \Psi_3 + \cdots
\end{aligned}
\tag{4.30}$$

For $\lambda = 0$, it is seen that $\Psi_0 = \Phi_0$ and $W_0 = E_0$, and this is the *unperturbed*, or *zeroth-order* wave function and energy. The Ψ_1, Ψ_2, \dots and W_1, W_2, \dots are the *first-order*, *second-order*, etc., *corrections*. The λ parameter will eventually be set equal to 1, and the n th-order energy or wave function becomes a sum of all terms up to order n .

It is convenient to choose the perturbed wave function to be *intermediately normalized*, i.e. the overlap with the unperturbed wave function should be 1. This has the consequence that all correction terms are orthogonal to the reference wave function.

$$\begin{aligned}
\langle \Psi | \Phi_0 \rangle &= 1 \\
\langle \Psi_0 + \lambda \Psi_1 + \lambda^2 \Psi_2 + \cdots | \Phi_0 \rangle &= 1 \\
\langle \Phi_0 | \Phi_0 \rangle + \lambda \langle \Psi_1 | \Phi_0 \rangle + \lambda^2 \langle \Psi_2 | \Phi_0 \rangle + \cdots &= 1 \\
\langle \Psi_{i \neq 0} | \Phi_0 \rangle &= 0
\end{aligned}
\tag{4.31}$$

Once all the correction terms have been calculated, it is trivial to normalize the total wave function.

With the expansions (eq. (4.30)), the Schrödinger equation (eq. (4.29)) becomes eq. (4.32).

$$\begin{aligned}
(\mathbf{H}_0 + \lambda \mathbf{H}')(\lambda^0 \Psi_0 + \lambda^1 \Psi_1 + \lambda^2 \Psi_2 + \cdots) \\
= (\lambda^0 W_0 + \lambda^1 W_1 + \lambda^2 W_2 + \cdots)(\lambda^0 \Psi_0 + \lambda^1 \Psi_1 + \lambda^2 \Psi_2 + \cdots)
\end{aligned}
\tag{4.32}$$

Since this holds for any value of λ , we can collect terms with the same power of λ to give eq. (4.33).

$$\begin{aligned}
\lambda^0: \mathbf{H}_0 \Psi_0 &= W_0 \Psi_0 \\
\lambda^1: \mathbf{H}_0 \Psi_1 + \mathbf{H}' \Psi_0 &= W_0 \Psi_1 + W_1 \Psi_0 \\
\lambda^2: \mathbf{H}_0 \Psi_2 + \mathbf{H}' \Psi_1 &= W_0 \Psi_2 + W_1 \Psi_1 + W_2 \Psi_0 \\
\lambda^n: \mathbf{H}_0 \Psi_n + \mathbf{H}' \Psi_{n-1} &= \sum_{i=0}^n W_i \Psi_{n-i}
\end{aligned}
\tag{4.33}$$

These are the zero-, first-, second-, n th-order perturbation equations. The zeroth-order equation is just the Schrödinger equation for the unperturbed problem. The first-order equation contains two unknowns, the first-order correction to the energy, W_1 , and the first-order correction to the wave function, Ψ_1 . The n th-order energy correction can be calculated by multiplying from the left by Φ_0 and integrating, and using the “*turnover rule*” $\langle \Phi_0 | \mathbf{H}_0 | \Psi_i \rangle = \langle \Psi_i | \mathbf{H}_0 | \Phi_0 \rangle^*$.

$$\begin{aligned}
\langle \Phi_0 | \mathbf{H}_0 | \Psi_n \rangle + \langle \Phi_0 | \mathbf{H}' | \Psi_{n-1} \rangle &= \sum_{i=0}^{n-1} W_i \langle \Phi_0 | \Psi_{n-i} \rangle + W_n \langle \Phi_0 | \Psi_0 \rangle \\
E_0 \langle \Psi_n | \Phi_0 \rangle + \langle \Phi_0 | \mathbf{H}' | \Psi_{n-1} \rangle &= W_n \langle \Phi_0 | \Psi_0 \rangle \\
W_n &= \langle \Phi_0 | \mathbf{H}' | \Psi_{n-1} \rangle
\end{aligned}
\tag{4.34}$$

From this it would appear that the $(n - 1)$ th-order wave function is required for calculating the n th-order energy. However, by using the turnover rule and the n th- and lower order perturbation equations (4.33), it can be shown that knowledge of the n th-order wave function actually allows a calculation of the $(2n + 1)$ th-order energy.

$$W_{2n+1} = \langle \Psi_n | \mathbf{H}' | \Psi_n \rangle - \sum_{k,l=0}^n W_{2n+1-k-l} \langle \Psi_k | \Psi_l \rangle \quad (4.35)$$

Up to this point, we are still dealing with undetermined quantities, energy and wave function corrections at each order. The first-order equation is one equation with two unknowns. Since the solutions to the unperturbed Schrödinger equation generate a complete set of functions, the unknown first-order correction to the wave function can be expanded in these functions. This is known as *Rayleigh-Schrödinger* perturbation theory, and the λ^1 equation in eq. (4.33) becomes eq. (4.36).

$$\begin{aligned} \Psi_1 &= \sum_i c_i \Phi_i \\ (\mathbf{H}_0 - W_0) \left(\sum_i c_i \Phi_i \right) + (\mathbf{H}' - W_1) \Phi_0 &= 0 \end{aligned} \quad (4.36)$$

Multiplying from the left by Φ_0^* and integrating yields eq. (4.37), where the orthonormality of the Φ_i s is used (this also follows directly from eq. (4.35)).

$$\begin{aligned} \sum_i c_i \langle \Phi_0 | \mathbf{H}_0 | \Phi_i \rangle - W_0 \sum_i c_i \langle \Phi_0 | \Phi_i \rangle + \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle - W_1 \langle \Phi_0 | \Phi_0 \rangle &= 0 \\ \sum_i c_i E_i \langle \Phi_0 | \Phi_i \rangle - c_0 E_0 + \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle - W_1 &= 0 \\ c_0 E_0 - c_0 E_0 + \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle - W_1 &= 0 \\ W_1 &= \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle \end{aligned} \quad (4.37)$$

The last equation shows that the first-order correction to the energy is an average of the perturbation operator over the unperturbed wave function.

The first-order correction to the wave function can be obtained by multiplying eq. (4.33) from the left by a function other than Φ_0 (Φ_j) and integrating to give eq. (4.38).

$$\begin{aligned} \sum_i c_i \langle \Phi_j | \mathbf{H}_0 | \Phi_i \rangle - W_0 \sum_i c_i \langle \Phi_j | \Phi_i \rangle + \langle \Phi_j | \mathbf{H}' | \Phi_0 \rangle - W_1 \langle \Phi_j | \Phi_0 \rangle &= 0 \\ \sum_i c_i E_i \langle \Phi_j | \Phi_i \rangle - c_j E_0 + \langle \Phi_j | \mathbf{H}' | \Phi_0 \rangle &= 0 \\ c_j E_j - c_j E_0 + \langle \Phi_j | \mathbf{H}' | \Phi_0 \rangle &= 0 \\ c_j &= \frac{\langle \Phi_j | \mathbf{H}' | \Phi_0 \rangle}{E_0 - E_j} \end{aligned} \quad (4.38)$$

The expansion coefficients determine the first-order correction to the perturbed wave function (eq. (4.36)), and they can be calculated from the known unperturbed wave functions and energies. The coefficient in front of Φ_0 for Ψ_1 cannot be determined from

the above formula, but the assumption of intermediate normalization (eq. (4.31)) makes $c_0 = 0$. Starting from the second-order perturbation equation (4.33), analogous formulas can be generated for the second-order corrections. Using intermediate normalization ($c_0 = d_0 = 0$), the second-order energy correction is given by eq. (4.39).

$$\begin{aligned}
 \Psi_2 &= \sum_i d_i \Phi_i \\
 (\mathbf{H}_0 - W_0) \left(\sum_i d_i \Phi_i \right) + (\mathbf{H}' - W_1) \left(\sum_i c_i \Phi_i \right) - W_2 \Phi_0 &= 0 \\
 \sum_i d_i \langle \Phi_0 | \mathbf{H}_0 | \Phi_i \rangle - W_0 \sum_i d_i \langle \Phi_0 | \Phi_i \rangle + \\
 \sum_i c_i \langle \Phi_0 | \mathbf{H}' | \Phi_i \rangle - W_1 \sum_i c_i \langle \Phi_0 | \Phi_i \rangle - W_2 \langle \Phi_0 | \Phi_0 \rangle &= 0 \\
 \sum_i d_i E_i \langle \Phi_0 | \Phi_i \rangle - d_0 E_0 + \sum_i c_i \langle \Phi_0 | \mathbf{H}' | \Phi_i \rangle - c_0 W_1 - W_2 &= 0 \\
 d_0 E_0 - d_0 E_0 + \sum_i c_i \langle \Phi_0 | \mathbf{H}' | \Phi_i \rangle - W_2 &= 0 \\
 W_2 = \sum_i c_i \langle \Phi_0 | \mathbf{H}' | \Phi_i \rangle = \sum_{i \neq 0} \frac{\langle \Phi_0 | \mathbf{H}' | \Phi_i \rangle \langle \Phi_i | \mathbf{H}' | \Phi_0 \rangle}{E_0 - E_i}
 \end{aligned} \tag{4.39}$$

The last equation shows that the second-order energy correction may be written in terms of the first-order wave function (c_i) and matrix elements over unperturbed states. The second-order wave function correction is given by eq. (4.40).

$$\begin{aligned}
 \sum_i d_i \langle \Phi_j | \mathbf{H}_0 | \Phi_i \rangle - W_0 \sum_i d_i \langle \Phi_j | \Phi_i \rangle + \\
 \sum_i c_i \langle \Phi_j | \mathbf{H}' | \Phi_i \rangle - W_1 \sum_i c_i \langle \Phi_j | \Phi_i \rangle - W_2 \langle \Phi_j | \Phi_0 \rangle &= 0 \\
 \sum_i d_i E_i \langle \Phi_j | \Phi_i \rangle - d_j E_0 + \sum_i c_i \langle \Phi_j | \mathbf{H}' | \Phi_i \rangle - c_j W_1 &= 0 \\
 d_j E_j - d_j E_0 + \sum_i c_i \langle \Phi_j | \mathbf{H}' | \Phi_i \rangle - c_j \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle &= 0 \\
 d_j = \sum_{i \neq 0} \frac{\langle \Phi_j | \mathbf{H}' | \Phi_i \rangle \langle \Phi_i | \mathbf{H}' | \Phi_0 \rangle}{(E_0 - E_j)(E_0 - E_i)} - \frac{\langle \Phi_j | \mathbf{H}' | \Phi_0 \rangle \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle}{(E_0 - E_j)^2}
 \end{aligned} \tag{4.40}$$

The formulas for higher order corrections become increasingly complex; the third-order energy correction for example is given in eq. (4.41).

$$W_3 = \sum_{i,j \neq 0} \frac{\langle \Phi_0 | \mathbf{H}' | \Phi_i \rangle [\langle \Phi_i | \mathbf{H}' | \Phi_j \rangle - \delta_{ij} \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle \langle \Phi_j | \mathbf{H}' | \Phi_0 \rangle]}{(E_0 - E_i)(E_0 - E_j)} \tag{4.41}$$

The main point, however, is that all corrections can be expressed in terms of matrix elements of the perturbation operator over unperturbed wave functions, and the unperturbed energies.

4.8.1 Møller–Plesset perturbation theory

So far, the theory has been completely general. In order to apply perturbation theory to the calculation of correlation energy, the unperturbed Hamiltonian operator must

be selected. The most common choice is to take this as a sum over Fock operators, leading to *Møller–Plesset* (MP) perturbation theory.⁹ The sum of Fock operators counts the (average) electron–electron repulsion twice (eq. (3.44)), and the perturbation becomes the exact \mathbf{V}_{ee} operator minus twice the $\langle \mathbf{V}_{ee} \rangle$ operator. The operator associated with this difference is often referred to as the *fluctuation potential*. This choice is not really consistent with the basic assumption that the perturbation should be small compared with \mathbf{H}_0 . However, it does fulfil the other requirement that solutions to the unperturbed Schrödinger equation should be known. Furthermore, this is the only choice that leads to a size extensive method, which is a desirable feature.

$$\begin{aligned}
 \mathbf{H}_0 &= \sum_{i=1}^{N_{\text{elec}}} \mathbf{F}_i = \sum_{i=1}^{N_{\text{elec}}} \left(\mathbf{h}_i + \sum_{j=1}^{N_{\text{elec}}} (\mathbf{J}_j - \mathbf{K}_j) \right) \\
 &= \sum_{i=1}^{N_{\text{elec}}} \mathbf{h}_i + \sum_{i=1}^{N_{\text{elec}}} \sum_{j=1}^{N_{\text{elec}}} \langle \mathbf{g}_{ij} \rangle = \sum_{i=1}^{N_{\text{elec}}} \mathbf{h}_i + 2\langle \mathbf{V}_{ee} \rangle \\
 \mathbf{H}' &= \mathbf{H} - \mathbf{H}_0 = \sum_{i=1}^{N_{\text{elec}}} \sum_{j>i}^{N_{\text{elec}}} \mathbf{g}_{ij} - \sum_{i=1}^{N_{\text{elec}}} \sum_{j=1}^{N_{\text{elec}}} \langle \mathbf{g}_{ij} \rangle = \mathbf{V}_{ee} - 2\langle \mathbf{V}_{ee} \rangle
 \end{aligned} \tag{4.42}$$

The zeroth-order wave function is the HF determinant, and the zeroth-order energy is just a sum of MO energies.

$$W_0 = \langle \Phi_0 | \mathbf{H}_0 | \Phi_0 \rangle = \left\langle \Phi_0 \left| \sum_{i=1}^{N_{\text{elec}}} \mathbf{F}_i \right| \Phi_0 \right\rangle = \sum_{i=1}^{N_{\text{elec}}} \epsilon_i \tag{4.43}$$

Recall that the orbital energy is the energy of an electron in the field of all the nuclei and includes the repulsion to all other electrons, eq. (3.44), and therefore counts the electron–electron repulsion twice. The first-order energy correction is the average of the perturbation operator over the zeroth-order wave function (eq. (4.37)).

$$W_1 = \langle \Phi_0 | \mathbf{H}' | \Phi_0 \rangle = \langle \mathbf{V}_{ee} \rangle - 2\langle \mathbf{V}_{ee} \rangle = -\langle \mathbf{V}_{ee} \rangle \tag{4.44}$$

This yields a correction for the overcounting of the electron–electron repulsion at zeroth order. Comparing eq. (4.44) with the expression for the total energy in eq. (3.32), it is seen that the first-order energy (sum of W_0 and W_1) is exactly the HF energy. Using the notation $E(\text{MP}n)$ to indicate the correction at order n , and $\text{MP}n$ to indicate the total energy up to order n , we have eq. (4.45).

$$\begin{aligned}
 \text{MP}0 &= E(\text{MP}0) = \sum_{i=1}^{N_{\text{elec}}} \epsilon_i \\
 \text{MP}1 &= E(\text{MP}0) + E(\text{MP}1) = E(\text{HF})
 \end{aligned} \tag{4.45}$$

Electron correlation energy thus starts at order two with this choice of \mathbf{H}_0 .

In developing perturbation theory, it was assumed that the solutions to the unperturbed problem formed a complete set. This in general means that there must be an infinite number of functions, which is impossible in actual calculations. The lowest energy solution to the unperturbed problem is the HF wave function, additional higher energy solutions are excited Slater determinants, analogous to the CI method. When a finite basis set is employed, it is only possible to generate a finite number of excited determinants. The expansion of the many-electron wave function is therefore truncated.

Let us look at the expression for the second-order energy correction, eq. (4.39). This involves matrix elements of the perturbation operator between the HF reference and all possible excited states. Since the perturbation is a two-electron operator, all matrix elements involving triple, quadruple, etc., excitations are zero. When canonical HF orbitals are used, matrix elements with singly excited states are also zero, as indicated in eq. (4.46).

$$\begin{aligned}
 \langle \Phi_0 | \mathbf{H}' | \Phi_i^a \rangle &= \left\langle \Phi_0 \left| \mathbf{H}_0 - \sum_j^{N_{\text{elec}}} \mathbf{F}_j \right| \Phi_i^a \right\rangle \\
 &= \langle \Phi_0 | \mathbf{H}_0 | \Phi_i^a \rangle - \left\langle \Phi_0 \left| \sum_j^{N_{\text{elec}}} \mathbf{F}_j \right| \Phi_i^a \right\rangle \\
 &= \langle \Phi_0 | \mathbf{H}_0 | \Phi_i^a \rangle - \left(\sum_j^{N_{\text{elec}}} \epsilon_j \right) \langle \Phi_0 | \Phi_i^a \rangle = 0
 \end{aligned} \tag{4.46}$$

The first bracket is zero owing to Brillouin's theorem (Section 4.2.1), and the second set of brackets is zero owing to the orbitals being eigenfunctions of the Fock operators and orthogonal to each other. The second-order correction to the energy, which is the first contribution to the correlation energy, thus only involves a sum over doubly excited determinants. These can be generated by promoting two electrons from occupied orbitals i and j to virtual orbitals a and b . The summation must be restricted such that each excited state is only counted once.

$$W_2 = \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{vir}} \frac{\langle \Phi_0 | \mathbf{H}' | \Phi_{ij}^{ab} \rangle \langle \Phi_{ij}^{ab} | \mathbf{H}' | \Phi_0 \rangle}{E_0 - E_{ij}^{ab}} \tag{4.47}$$

The matrix elements between the HF and a doubly excited state are given by two-electron integrals over MOs (eq. (4.8)). The difference in total energy between two Slater determinants becomes a difference in MO energies (essentially Koopmans' theorem), and the explicit formula for the second-order Møller–Plesset correction is given in eq. (4.48).

$$E(\text{MP2}) = \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{vir}} \frac{(\langle \phi_i \phi_j | \phi_a \phi_b \rangle - \langle \phi_i \phi_j | \phi_b \phi_a \rangle)}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \tag{4.48}$$

Once the two-electron integrals over MOs are available, the second-order energy correction can be calculated as a sum over such integrals. There are of the order of M_{basis}^4 integrals, thus the calculation of the energy (only) increases as M_{basis}^4 with the system size. However, the transformation of the integrals from the AO to the MO basis grows as M_{basis}^5 (Section 4.2.1). MP2 is an M_{basis}^5 method, but fairly inexpensive as not all two-electron integrals over MOs are required. Only those corresponding to the combination of two occupied and two virtual MOs are needed. In practical calculations, this means that the MP2 energy for systems with a few hundred basis functions can be calculated at a cost similar to or less than what is required for calculating the HF energy. MP2 typically accounts for 80–90% of the correlation energy, and it is the most economical method for including electron correlation.

The formula for the first-order correction to the wave function (eq. (4.38)) similarly only contains contributions from doubly excited determinants. Since knowledge of the

first-order wave function allows calculation of the energy up to third order ($2n + 1 = 3$, eq. (4.35)), it is immediately clear that the third-order energy also only contains contributions from doubly excited determinants. Qualitative speaking, the MP2 contribution describes the correlation between pairs of electrons while MP3 describes the interaction between pairs. The formula for calculating this contribution is given in eq. (4.41) and involves a computational effort that formally increases as M_{basis}^6 . The third-order energy typically accounts for 90–95% of the correlation energy.

The formula for the second-order correction to the wave function (eq. (4.40)) contains products of the type $\langle \Phi_j | \mathbf{H}' | \Phi_i \rangle \langle \Phi_i | \mathbf{H}' | \Phi_0 \rangle$. The Φ_0 is the HF determinant and the *last* bracket can only be non-zero if Φ_i is a doubly excited determinant. This means that the *first* bracket only can be non-zero if Φ_j is either a singly, doubly, triply or quadruply excited determinant (since \mathbf{H}' is a two-electron operator). The second-order wave function allows calculation of the fourth- and fifth-order energies, and these terms therefore have contributions from determinants that are singly, doubly, triply or quadruply excited. The computational cost of the fourth-order energy without the contribution from the triply excited determinants, MP4(SDQ), increases as M_{basis}^6 , while the triples contribution increases as M_{basis}^7 . MP4 is still a computationally feasible model for many molecular systems, requiring a time similar to CISD. In typical calculations, the T contribution to MP4 will take roughly the same amount of time as the SDQ contributions, but the triples are often the most important at fourth order. The full fourth-order energy typically accounts for 95–98% of the correlation energy.

The fifth-order correction to the energy also involves S, D, T and Q contributions, and the sixth-order term introduces quintuple and sextuple excitations. The working formulas for the MP5 and MP6 contributions are so complex that actual calculations are only possible for small systems. The computational effort for MP5 increases as M_{basis}^8 and for MP6 as M_{basis}^9 . There is very little experience with the performance of $\text{MP}n$ beyond MP4.

As shown in Table 4.2, the most important contribution to the energy in a CI procedure comes from doubly excited determinants. This is also shown by the perturbation expansion, the second- and third-order energy corrections only involve doubles. At fourth order the singles, triples and quadruples enter the expansion for the first time. This is again consistent with Table 4.2, which shows that these types of excitations are of similar importance.

CI methods determine the energy by a variational procedure, and the energy is consequently an upper bound to the exact energy. There is no such guarantee for perturbation methods, and it is possible that the energy will be lower than the exact energy. This is rarely a problem and may in fact be advantageous. Limitations in the basis set often mean that the error in total energy is several au (thousands of kJ/mol) anyway. In the large majority of cases, the interest is not in total energies but in energy differences. Having a variational upper bound for two energies does not give any bound for the difference between these two numbers. The main interest is therefore that the error remains relatively constant for different systems, and the absence of a variational bound can allow for error cancellations. The lack of size extensivity of CI methods, on the other hand, is disadvantageous in this respect. The MP perturbation method *is* size extensive, but other forms of MBPT are not. It is now generally recognized that size extensivity is an important property, and the MP form of MBPT is used almost exclusively.

The main limitation of perturbation methods is the assumption that the zeroth-order wave function is a reasonable approximation to the real wave function, i.e. the perturbation operator is sufficiently “small”. The more poorly the HF wave function describes the system, the larger are the correction terms, and the more terms must be included to achieve a given level of accuracy. If the reference state is a poor description of the system, the convergence may be so slow or erratic that perturbation methods cannot be used. Actually, it is difficult to assess whether the perturbation expansion is convergent or not, although the first few terms for many systems show a behaviour that suggests that it is the case. This may to some extent be deceptive, as it has been demonstrated that the convergence properties depend on the size of the basis set,¹⁰ and the majority of studies have employed small- or medium-sized basis sets. A convergent series in for example a DZP type basis may become divergent or oscillating in a larger basis, especially if diffuse functions are present.

The convergence properties for the perturbation series can be analyzed by considering the partitioned Hamiltonian in eq. (4.28), with λ being a parameter connecting the reference system ($\lambda = 0$) with the real physical system ($\lambda = 1$).¹¹ For analyzing the convergence behaviour, we must allow λ also to have complex values.

$$\mathbf{H}(\lambda) = \mathbf{H}_0 + \lambda \mathbf{H}' \quad (4.49)$$

For a given λ value, the (exact) energy of the ground state can be written as an infinite summation of all perturbation terms.

$$E(\lambda) = \sum_{i=0}^{\infty} W_i \lambda^i \quad (4.50)$$

The mathematical theory of infinite series states that this summation is only convergent within a given radius R , i.e. the infinite series in eq. (4.50) only has a well-defined value if $|\lambda| < R$. Since we are interested in the situation where $\lambda = 1$, this translates into the condition $R > 1$. The convergence radius is determined by the smallest value of λ where another state becomes degenerate with the ground state, i.e. the MP perturbation series is only convergent if there are no excited states that become degenerate with the ground state within the circle in the complex plane corresponding to $|\lambda| = 1$. This includes non-physical situations where λ is negative, i.e. where the perturbation corresponds to the electron–electron interaction being *attractive*. In MP theory the zeroth-order energy is the sum of orbital energies, which includes the average electron–electron interaction twice, and the first-order energy correction W_1 is the negative of the average electron–electron interaction (eq. 4.44). Since W_1 is significantly smaller for excited states than for the ground state (more diffuse orbitals), this means that a negative λ value will raise the ground state more in energy than an excited state, and this may be sufficient to overcome the energy separation at $\lambda = 0$. This is especially true when diffuse basis functions are included, since they preferentially improve the description of excited states, and such *intruder states* are the reason for the non-convergent behaviour of the MP perturbation series. A complete search for intruder states within the complex plane corresponding to $|\lambda| = 1$ is difficult even for simple systems, and a less rigorous search for avoided crossings along the real axis is also demanding. Establishing the convergence or divergence of the MP expansion

on a case-by-case basis is unmanageable, and one is therefore limited to observing the behaviour for the first few terms.

In the ideal case, the HF, MP2, MP3 and MP4 results show a monotonic convergence towards a limiting value, with the corrections being of the same sign and numerically smaller as the order of perturbation increases. Unfortunately, this is not the typical behaviour. Even in systems where the reference is well described by a single determinant, oscillations in a given property as a function of perturbation order are often observed. An analysis by Cremer and He indicates that a smooth convergence (of the total energy) is only expected for systems containing well-separated electron pairs, and that oscillations occur when this is not that case.¹² The latter encompass system containing lone pairs and/or multiple bonds, covering the large majority of molecules. It should be noted that one cannot conclude anything about the convergence properties of the whole perturbation series from either the monotonic or oscillating behaviour of the first few terms.

In practice, only low orders of perturbation theory can be carried out, and it is often observed that the HF and MP2 results differ considerably, the MP3 result moves back towards the HF, and MP4 away again. For “well-behaved” systems the correct answer is often somewhere between the MP3 and MP4 results. MP2 typically overshoots the correlation effect, but often gives a better answer than MP3, at least if medium-sized basis sets are used. Just as the first term involving doubles (MP2) tends to overestimate the correlation effect, it is often observed that MP4 overestimates the effect of the singles and triples contributions, since they enter the series for the first time at fourth order.




Figure 4.13 Typical oscillating behaviour of results obtained with the MP method

When the reference wave function contains substantial multi-reference character, a perturbation expansion based on a single determinant will display poor convergence. If the reference wave function suffers from symmetry breaking (Section 3.8.3), the MP method is almost guaranteed to give absurd results. The questionable convergence of the MP method has caused it to be significantly less popular in recent years, although

MP2 continues to be a computationally cheap way of including the majority of the electron correlation effect.

4.8.2 *Unrestricted and projected Møller–Plesset methods*

When the reference is an RHF type wave function the dissociation limit will normally be incorrect. As a bond is stretched, RHF gives an increasingly poorer description of the wave function, and consequently causes the perturbation series to break down. The use of a UHF wave function allows a correct dissociation limit in terms of energy but at the cost of introducing spin contamination (Sections 4.3 and 4.4). It is straightforward to derive an MP method based on a UHF reference wave function (UMP): in this case the unperturbed Hamiltonian operator is a sum of the α and β Fock operators. The addition of electron correlation decreases the spin contamination of the wave function (in the full CI limit the spin contamination is zero) but the improvement is usually small at low orders (2–4) of perturbation theory. As illustrated in Section 4.4, the UHF energy is lower than that of RHF owing to the inclusion of some electron correlation (mainly static), but it also contains some amounts of higher energy spin states. Since MP methods recover a large part of the electron correlation (both static and dynamical), the net effect at the UMP level is an increase in energy due to spin contamination. In the dissociation limit, this has no consequence, as the different spin states have equal energies. In the intermediate region, where the bond is not completely broken, it is usually observed that the RMP n energy is *lower* than the UMP n energy, although the RHF energy is *higher* than the UHF (see also Section 11.5.2). The spin contamination in UHF wave functions causes an UMP n expansion to converge more slowly than RMP n .¹³ For open-shell systems, where RHF cannot be used, this would suggest that the reference wave function should be of the ROHF type, instead of UHF. Formulation of ROHF-based perturbation methods, however, is somewhat more difficult than for the UHF case. The reason is that for an ROHF wave function it is not possible to choose a set of MOs that makes the matrix of Lagrange multipliers diagonal (eqs (3.40) and (3.41)). There is thus not a unique set of canonical MOs to be used in the perturbation expansion, which again has the consequence that several choices of the unperturbed Hamiltonian operator are possible.¹⁴ Different ROMP methods therefore give different energies, and there are no firm theoretical grounds for choosing one over the other. In practice, however, different choices of the unperturbed Hamiltonian operator lead to similar results, and perturbation calculations based on ROHF type wave functions are now routine.

While projection methods for removing spin contamination are not recommended at the HF level, they work quite well at the UMP level. Formulas have been derived for removing all contaminants at the UMP2 level, and also the first few states at the UMP3 and UMP4 levels.¹⁵ The associated acronyms are PUMP and PMP, denoting slightly different methods, although in practice they give similar results. For singlet wave functions with bond lengths only slightly longer than the RHF/UHF instability point, such PUMP methods tend to give results very similar to those based on an RHF wave function. At longer bond lengths the RMP perturbation series eventually breaks down, while the PUMP methods approach the correct dissociation limit. It would therefore appear that PUMP methods should always be preferred. There are, however, also some computational factors to consider. First, UMP methods are by nature a factor

of ~ 2 more expensive since there are twice as many MO coefficients. Second, the projection itself also uses CPU time. This is especially true if many of the higher spin states need to be removed, or for projection at the MP4 level. Third, it is difficult to formulate derivatives of projected wave functions, which limits PUMP methods to the calculation of energies. A rule of thumb says that for uncomplicated systems the RMP4 treatment gives acceptable accuracy (relative errors of the order of a ~ 10 kJ/mol) up to bond lengths ~ 1.5 times the equilibrium length. Longer bonds are better treated by PUMP methods (see also Section 11.5.2). Most transition structures have bond lengths shorter than ~ 1.5 times the equilibrium length and RMP4 often gives quite accurate activation energies.

Just as single-reference CI can be extended to MRCI, it is also possible to use perturbation methods with a multi-determinant reference wave function. A formulation of MR-MBPT methods, however, is not straightforward. The main problem here is similar to that with ROMP methods: the choice of the unperturbed Hamiltonian operator. Several different choices are possible, which will give different answers when the theory is carried out only to low order. Nevertheless, there are now several different implementations of MP2 type expansions based on a CASSCF reference, denoted CASMP2 or CASPT2.¹⁶ Experience of their performance is still somewhat limited.

4.9 Coupled Cluster

Perturbation methods add *all types of corrections* (S, D, T, Q, etc.) to the reference wave function *to a given order* (2, 3, 4, etc.). The idea in *Coupled Cluster* (CC) methods is to include *all corrections of a given type to infinite order*.¹⁷ Let us start by defining an excitation operator \mathbf{T} as in eq. (4.51).

$$\mathbf{T} = \mathbf{T}_1 + \mathbf{T}_2 + \mathbf{T}_3 + \cdots + \mathbf{T}_{N_{\text{elec}}} \quad (4.51)$$

The \mathbf{T}_i operator acting on an HF reference wave function Φ_0 generates all i th excited Slater determinants.

$$\begin{aligned} \mathbf{T}_1 \Phi_0 &= \sum_i^{\text{occ}} \sum_a^{\text{vir}} t_i^a \Phi_i^a \\ \mathbf{T}_2 \Phi_0 &= \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{vir}} t_{ij}^{ab} \Phi_{ij}^{ab} \end{aligned} \quad (4.52)$$

In coupled cluster theory it is customary to use the term *amplitudes* for the expansion coefficients t , which are equivalent to the a_i coefficients in eq. (4.1).

Using intermediate normalization, a CI wave function can be generated by allowing the excitation operator to work on an HF wave function.

$$\Psi_{\text{CI}} = (\mathbf{1} + \mathbf{T})\Phi_0 = (\mathbf{1} + \mathbf{T}_1 + \mathbf{T}_2 + \mathbf{T}_3 + \mathbf{T}_4 + \cdots)\Phi_0 \quad (4.53)$$

The corresponding coupled cluster wave function, on the other hand, is defined in eq. (4.54).

$$\begin{aligned} \Psi_{\text{CC}} &= e^{\mathbf{T}} \Phi_0 \\ e^{\mathbf{T}} &= \mathbf{1} + \mathbf{T} + \frac{1}{2} \mathbf{T}^2 + \frac{1}{6} \mathbf{T}^3 + \cdots = \sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{T}^k \end{aligned} \quad (4.54)$$

From eqs (4.51) and (4.54) the exponential operator may be written as in eq. (4.55).

$$\begin{aligned} e^{\mathbf{T}} = & 1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2}\mathbf{T}_1^2) + (\mathbf{T}_3 + \mathbf{T}_2\mathbf{T}_1 + \frac{1}{6}\mathbf{T}_1^3) + \\ & (\mathbf{T}_4 + \mathbf{T}_3\mathbf{T}_1 + \frac{1}{2}\mathbf{T}_2^2 + \frac{1}{2}\mathbf{T}_2\mathbf{T}_1^2 + \frac{1}{24}\mathbf{T}_1^4) + \dots \end{aligned} \quad (4.55)$$

The first term generates the reference HF and the second all singly excited states. The first parenthesis generates all doubly excited states, which may be considered as *connected* (\mathbf{T}_2) or *disconnected* (\mathbf{T}_1^2). The second parenthesis generates all triply excited states, which again may be either “true” (\mathbf{T}_3) or “product” triples ($\mathbf{T}_2\mathbf{T}_1$, \mathbf{T}_1^3). The quadruply excited states can similarly be viewed as composed of five terms, a true quadruple and four product terms. Physically, a connected type such as \mathbf{T}_4 corresponds to four electrons interacting simultaneous, while a disconnected term such as \mathbf{T}_2^2 corresponds to two non-interacting pairs of interacting electrons. By comparison with the CI wave function in eq. (4.53), it is seen that the CC wave function at each excitation level contains additional terms arising from products of excitations.

With the coupled cluster wave function in eq. (4.54) the Schrödinger equation becomes eq. (4.56).

$$\mathbf{H}e^{\mathbf{T}}\Phi_0 = Ee^{\mathbf{T}}\Phi_0 \quad (4.56)$$

At this point, one could proceed analogously to CI and evaluate the energy as an expectation value of the CC wave function, and use the variational principle to determine the amplitudes.

$$\begin{aligned} E_{CC}^{\text{var}} &= \frac{\langle \Psi_{CC} | \mathbf{H} | \Psi_{CC} \rangle}{\langle \Psi_{CC} | \Psi_{CC} \rangle} = \frac{\langle e^{\mathbf{T}}\Phi_0 | \mathbf{H} | e^{\mathbf{T}}\Phi_0 \rangle}{\langle e^{\mathbf{T}}\Phi_0 | e^{\mathbf{T}}\Phi_0 \rangle} \\ E_{CC}^{\text{var}} &= \frac{\langle (1 + \mathbf{T} + \frac{1}{2}\mathbf{T}^2 \dots \frac{1}{N!}\mathbf{T}^N)\Phi_0 | \mathbf{H} | (1 + \mathbf{T} + \frac{1}{2}\mathbf{T}^2 \dots \frac{1}{N!}\mathbf{T}^N)\Phi_0 \rangle}{\langle (1 + \mathbf{T} + \frac{1}{2}\mathbf{T}^2 \dots \frac{1}{N!}\mathbf{T}^N)\Phi_0 | (1 + \mathbf{T} + \frac{1}{2}\mathbf{T}^2 \dots \frac{1}{N!}\mathbf{T}^N)\Phi_0 \rangle} \end{aligned} \quad (4.57)$$

Expansion of the numerator and denominator according to eq. (4.54) unfortunately leads to a series of non-vanishing terms all the way up to order N_{elec} , which makes a variational coupled cluster approach unmanageable for all but the smallest systems.¹⁸

The standard formulation of coupled cluster theory instead proceeds by *projecting* the coupled cluster Schrödinger equation (4.56) onto the reference wave function. Multiplying from the left by Φ_0^* and integrating gives eq. (4.58).

$$\begin{aligned} \langle \Phi_0 | \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= E_{CC} \langle \Phi_0 | e^{\mathbf{T}} \Phi_0 \rangle \\ \langle \Phi_0 | \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= E_{CC} \langle \Phi_0 | (1 + \mathbf{T}_1 + \mathbf{T}_2 + \dots) \Phi_0 \rangle \\ E_{CC} &= \langle \Phi_0 | \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle \end{aligned} \quad (4.58)$$

Expanding out the exponential in eq. (4.54) and using the fact that the Hamiltonian operator contains only one- and two-electron operators (eq. (3.24)) we get eq. (4.59).

$$\begin{aligned} E_{CC} &= \langle \Phi_0 | \mathbf{H} (1 + \mathbf{T}_1 + \mathbf{T}_2 + \frac{1}{2}\mathbf{T}_1^2) | \Phi_0 \rangle \\ E_{CC} &= \langle \Phi_0 | \mathbf{H} | \Phi_0 \rangle + \langle \Phi_0 | \mathbf{H} | \mathbf{T}_1 \Phi_0 \rangle + \langle \Phi_0 | \mathbf{H} | \mathbf{T}_2 \Phi_0 \rangle + \frac{1}{2} \langle \Phi_0 | \mathbf{H} | \mathbf{T}_1^2 \Phi_0 \rangle \\ E_{CC} &= E_0 + \sum_i^{\text{occ}} \sum_a^{\text{vir}} t_i^a \langle \Phi_0 | \mathbf{H} | \Phi_i^a \rangle + \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{vir}} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) \langle \Phi_0 | \mathbf{H} | \Phi_{ij}^{ab} \rangle \end{aligned} \quad (4.59)$$

Note that the infinite expansion of the exponential operator in eqs (4.58) and (4.59) terminates at the 1 and \mathbf{T}_2 levels, in contrast to \mathbf{T}_N in eq. (4.57). Furthermore, when using HF orbitals for constructing the Slater determinants, the first matrix elements in eq. (4.59) are zero (Brillouin's theorem) and the second matrix elements are just two-electron integrals over MOs (eq. (4.8)).

$$E_{CC} = E_0 + \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{vir}} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) (\langle \phi_i \phi_j | \phi_a \phi_b \rangle - \langle \phi_i \phi_j | \phi_b \phi_a \rangle) \quad (4.60)$$

The coupled cluster correlation energy is therefore determined completely by the singles and doubles amplitudes and the two-electron MO integrals.

Equations for the amplitudes can be obtained by projecting the Schrödinger equation (4.56) onto the space of singly, doubly, triply, etc., excited determinants. While this can be done analogously to eq. (4.58), a more elegant formulation is possible by using a *similarity transformation* of the Hamiltonian operator. Consider eq. (4.56) where we multiply from the left by $e^{-\mathbf{T}}$.

$$e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} \Phi_0 = E_{CC} \Phi_0 \quad (4.61)$$

Just as $e^{\mathbf{T}}$ is an *excitation* operator working on the function to the *right*, $e^{-\mathbf{T}}$ is a *deexcitation* operator working on the function to the *left*. Multiplying with Φ_0^* from the left and integrating leads directly to the energy equation.

$$E_{CC} = \langle \Phi_0 | e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle \quad (4.62)$$

Note that equation (4.62) can be considered as the expectation value of a similarity transformed (non-Hermitian) Hamiltonian. Since $e^{-\mathbf{T}}$ tries to generate deexcitations from the reference Φ_0^* , which is impossible, eq. (4.62) is identical to eq. (4.58).

Equations for the amplitudes are obtained by multiplying with an excited state.

$$\begin{aligned} \langle \Phi_m^e | e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= 0 \\ \langle \Phi_{mn}^{ef} | e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= 0 \\ \langle \Phi_{mnl}^{efg} | e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= 0 \\ &\vdots \end{aligned} \quad (4.63)$$

The deexcitation operator $e^{-\mathbf{T}}$ working on $\langle \Phi_m^e |$ now generates the reference wave function in addition to the singly excited state.

$$\begin{aligned} \langle \Phi_m^e | e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= 0 \\ \langle \Phi_m^e (1 - \mathbf{T}_1) | \mathbf{H} | (1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_3 + \mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3)) | \Phi_0 \rangle &= 0 \end{aligned} \quad (4.64)$$

Only the indicated terms survive in the expansion when the orthogonality of the Slater determinants and the nature of the Hamiltonian operator (only one- and two-electron terms) are considered. The terms involving singly excited states and the reference wave function are again zero owing to Brillouin's theorem, and the remaining terms form a coupled set of equations with single, double and triple amplitudes as the variables.

Similarly, $e^{-\mathbf{T}}$ working on $\langle \Phi_{mn}^{ef} |$ generates both the reference and singly excited states, in addition to the doubly excited states.

$$\begin{aligned} \langle \Phi_{mn}^{ef} | e^{-\mathbf{T}} \mathbf{H} e^{\mathbf{T}} | \Phi_0 \rangle &= 0 \\ \left\langle \Phi_{mn}^{ef} (1 - \mathbf{T}_1 - \mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) | \mathbf{H} \left((1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_3 + \mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3) + \right. \right. \right. \\ &\quad \left. \left. \left. (\mathbf{T}_4 + \mathbf{T}_3 \mathbf{T}_1 + \frac{1}{2} \mathbf{T}_2^2 + \frac{1}{2} \mathbf{T}_2 \mathbf{T}_1^2 + \frac{1}{24} \mathbf{T}_1^4) \right) \right| \Phi_0 \right\rangle = 0 \end{aligned} \quad (4.65)$$

Equation (4.65) now in addition has quadruple amplitudes, and additional terms coupling the lower order amplitudes. More equations connecting amplitudes may be obtained by projection against a triple, quadruple, etc., excited determinant.

4.9.1 Truncated coupled cluster methods

So far, everything has been exact. If all cluster operators up to \mathbf{T}_N are included in \mathbf{T} , all possible excited determinants are generated and the coupled cluster wave function is equivalent to full CI. This is, as already stated, impossible for all but the smallest systems. The cluster operator must therefore be truncated at some excitation level. When the \mathbf{T} operator is truncated, some of the terms in the amplitude equations will become zero, and the amplitudes derived from these approximate equations will no longer be exact. The energy calculated from these approximate singles and doubles amplitudes (eq. (4.60)) will therefore also be approximate. How severe the approximation is depends on how many terms are included in \mathbf{T} . Including only the \mathbf{T}_1 operator does not give any improvement over HF, as matrix elements between the HF and singly excited states are zero. The lowest level of approximation is therefore $\mathbf{T} = \mathbf{T}_2$, referred to as *Coupled Cluster Doubles* (CCD). Compared with the number of doubles, there are relatively few singly excited states. Using $\mathbf{T} = \mathbf{T}_1 + \mathbf{T}_2$ gives the CCSD model, which is only slightly more demanding than CCD, and yields a more complete model. Both CCD and CCSD involve a computational effort that scales as M_{basis}^6 in the limit of a large basis set. The next higher level has $\mathbf{T} = \mathbf{T}_1 + \mathbf{T}_2 + \mathbf{T}_3$, giving the CCSDT model.¹⁹ This involves a computational effort that scales as M_{basis}^8 and is more demanding than CISDT. It (and higher order methods such as CCSDTQ) can consequently only be used for small systems, and CCSD is the only generally applicable coupled cluster method.

Let us look in a bit more detail at the CCSD method. In this case, we have from eq. (4.55).

$$e^{\mathbf{T}_1 + \mathbf{T}_2} = 1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3) + (\frac{1}{2} \mathbf{T}_2^2 + \frac{1}{2} \mathbf{T}_2 \mathbf{T}_1^2 + \frac{1}{24} \mathbf{T}_1^4) + \dots \quad (4.66)$$

The CCSD energy is given by the general CC equation (4.60), and amplitude equations are derived from (4.64).

$$\begin{aligned} \langle \Phi_m^e | (1 - \mathbf{T}_1) | \mathbf{H} | (1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3)) \Phi_0 \rangle &= 0 \\ \langle \Phi_m^e | \mathbf{H} | (\mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3)) \Phi_0 \rangle - t_m^e \langle \Phi_0 | \mathbf{H} | (1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2)) \Phi_0 \rangle &= 0 \\ \sum_{ia} t_i^a \langle \Phi_m^e | \mathbf{H} | \Phi_i^a \rangle + \sum_{ijab} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) \langle \Phi_m^e | \mathbf{H} | \Phi_{ij}^{ab} \rangle + \\ \sum_{ijkabc} (t_{ij}^{ab} t_k^c + \dots + t_i^a t_j^b t_k^c + \dots) \langle \Phi_m^e | \mathbf{H} | \Phi_{ijk}^{abc} \rangle - \\ t_m^e E_0 - t_m^e \sum_{ijab} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) \langle \Phi_0 | \mathbf{H} | \Phi_{ij}^{ab} \rangle &= 0 \end{aligned} \quad (4.67)$$

The notation $(t_i^a t_j^b t_k^c + \dots)$ indicates that terms involving permutations of the indices are omitted. From eq. (4.65) we obtain eq. (4.68).

$$\begin{aligned}
& \left\langle \Phi_{mn}^{ef} (1 - \mathbf{T}_1 - \mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) | \mathbf{H} \left((1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3)) \right. \right. \\
& \quad \left. \left. + (\frac{1}{2} \mathbf{T}_2^2 + \frac{1}{2} \mathbf{T}_2 \mathbf{T}_1^2 + \frac{1}{24} \mathbf{T}_1^4) \right) \Phi_0 \right\rangle = 0 \\
& \langle \Phi_{mn}^{ef} | \mathbf{H} | (1 + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3) + (\frac{1}{2} \mathbf{T}_2^2 + \frac{1}{2} \mathbf{T}_2 \mathbf{T}_1^2 + \frac{1}{24} \mathbf{T}_1^4)) \Phi_0 \rangle - \\
& \quad t_n^f \langle \Phi_m^e | \mathbf{H} | (\mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2) + (\mathbf{T}_2 \mathbf{T}_1 + \frac{1}{6} \mathbf{T}_1^3)) \Phi_0 \rangle + \\
& \quad (-t_{mn}^{ef} + t_m^e t_n^f - t_m^f t_n^e) \langle \Phi_0 | \mathbf{H} | (1 + (\mathbf{T}_2 + \frac{1}{2} \mathbf{T}_1^2)) \Phi_0 \rangle = 0 \\
& \langle \Phi_{mn}^{ef} | \mathbf{H} | \Phi_0 \rangle + \sum_{ia} t_i^a \langle \Phi_{mn}^{ef} | \mathbf{H} | \Phi_i^a \rangle + \sum_{ijab} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) \langle \Phi_{mn}^{ef} | \mathbf{H} | \Phi_{ij}^{ab} \rangle + \\
& \quad \sum_{ijkabc} (t_{ij}^{ab} t_k^c + \dots + t_i^a t_j^b t_k^c + \dots) \langle \Phi_{mn}^{ef} | \mathbf{H} | \Phi_{ijk}^{abc} \rangle + \\
& \quad \sum_{ijklabcd} (t_{ij}^{ab} t_{kl}^{cd} + \dots + t_{ij}^{ab} t_k^c t_l^d + \dots + t_i^a t_j^b t_k^c t_l^d + \dots) \langle \Phi_{mn}^{ef} | \mathbf{H} | \Phi_{ijkl}^{abcd} \rangle - \\
& \quad t_n^f \sum_{ia} t_i^a \langle \Phi_m^e | \mathbf{H} | \Phi_i^a \rangle - t_n^f \sum_{ijab} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) \langle \Phi_m^e | \mathbf{H} | \Phi_{ij}^{ab} \rangle - \\
& \quad t_n^f \sum_{ijkabc} (t_{ij}^{ab} t_k^c + \dots + t_i^a t_j^b t_k^c + \dots) \langle \Phi_m^e | \mathbf{H} | \Phi_{ijk}^{abc} \rangle + \\
& \quad (-t_{mn}^{ef} + t_m^e t_n^f - t_m^f t_n^e) \left\{ E_0 + \sum_{ijab} (t_{ij}^{ab} + t_i^a t_j^b - t_i^b t_j^a) \langle \Phi_0 | \mathbf{H} | \Phi_{ij}^{ab} \rangle \right\} = 0
\end{aligned} \tag{4.68}$$

Equations (4.67) and (4.68) involve matrix elements between singles and triples, and between doubles and quadruples. However, since the Hamiltonian operator only contains one- and two-electron operators, these are actually identical to matrix elements between the reference and a doubly excited state. Consider for example $\langle \Phi_m^e | \mathbf{H} | \Phi_{ijk}^{abc} \rangle$. Unless m equals either i, j or k , and e equals either a, b or c , there will be one overlap integral between different MOs which makes the matrix element zero. If for example $m = k$ and $e = c$, then the MO integral over these indices factor out as 1, and the rest is equal to a matrix element $\langle \Phi_0 | \mathbf{H} | \Phi_{ij}^{ab} \rangle$. Similarly, the matrix element $\langle \Phi_{mn}^{ef} | \mathbf{H} | \Phi_{ijkl}^{abcd} \rangle$, between a doubly and a quadruply excited determinant, is only non-zero if mn matches up with two of the $ijkl$ indices, and ef matches up with $abcd$. Again, such non-zero matrix elements are equal to matrix elements between the reference and a doubly excited determinant, eq. (4.8).

All the matrix elements can be evaluated in terms of MO integrals, and the expressions in eqs (4.67) and (4.68) form coupled non-linear equations for the singles and doubles amplitudes. The equations contain terms up to quartic in the amplitudes, e.g. $(t_i^e)^4$ (since \mathbf{H} contains one- and two-electron operators), and must be solved by iterative techniques. Once the amplitudes are known, the energy and wave function can be calculated. The important aspect in coupled cluster methods is that excitations of higher order than the truncation of the \mathbf{T} operator enter the amplitude equation. Quadruply excited states, for example, are generated by the \mathbf{T}_2^2 operator in CCSD, and they enter the amplitude equations with a weight given as a *product* of doubles amplitudes. Quadruply excited states influence the doubles amplitudes, and thereby also the CCSD energy. It is the inclusion of these products of excitations that makes coupled cluster theory size extensive. For the case of a single H_2 molecule, a CISD calculation is equivalent to CCSD, and is also equivalent to a full CI calculation. For two H_2 molecules separated by 100 Å, however, a CISD is not equivalent

to a full CI (it is missing the T and Q excitations), but a CCSD calculation is still equivalent to a full CI.

4.10 Connections between Coupled Cluster, Configuration Interaction and Perturbation Theory

The general cluster operator is given by eq. (4.69), where terms have been collected according to the excitation they generate.

$$e^T = \mathbf{1} + \mathbf{T}_1 + (\mathbf{T}_2 + \frac{1}{2}\mathbf{T}_1^2) + (\mathbf{T}_3 + \mathbf{T}_2\mathbf{T}_1 + \frac{1}{6}\mathbf{T}_1^3) + (\mathbf{T}_4 + \mathbf{T}_3\mathbf{T}_1 + \frac{1}{2}\mathbf{T}_2^2 + \frac{1}{2}\mathbf{T}_2\mathbf{T}_1^2 + \frac{1}{24}\mathbf{T}_1^4) + \dots \quad (4.69)$$

Each of the operators in a given parenthesis generates *all* the excited determinants of the given type. Both \mathbf{T}_2 and \mathbf{T}_1^2 generate all doubly excited determinants, and the terms in eq. (4.69) generate all determinants that are included in a CISDTQ calculation. The cluster expansion can be viewed as a method of dividing up the contributions from each excitation type. The total contribution from double excitations is the sum of two terms, one that is the square of the singles contributions and the remaining is (by definition) the connected doubles. Similarly, the total contribution from triple excitations is a sum of three terms, the cube of the singles contributions, the product of the singles and doubles contribution, and the remaining is the connected triples.

The \mathbf{T}_1 effect is small when canonical HF orbitals are used, although not zero since singles enter indirectly via the doubly excited states (note that if non-canonical orbitals are used, the \mathbf{T}_1 term can be large). From CI we know that the effect of doubles is the most important (Section 4.2.3). In coupled cluster theory the doubles contribution is divided into \mathbf{T}_1^2 and \mathbf{T}_2 . If \mathbf{T}_1 is small, then \mathbf{T}_1^2 must also be small, and the most important term is \mathbf{T}_2 . For the triple excitations, \mathbf{T}_1^3 must be negligible, and $\mathbf{T}_1\mathbf{T}_2$ is small owing to \mathbf{T}_1 . The most important contribution is therefore from connected triples \mathbf{T}_3 . For the quadruple excitations, all the terms involving \mathbf{T}_1 must again be small, and since \mathbf{T}_2 is large, we expect the disconnected quadruples \mathbf{T}_2^2 to be the dominant term. This again suggests that the connected quadruples term \mathbf{T}_4 is small, which is reasonable since it correspond to a simultaneous correlation of four electrons. Higher order excitations will always contain terms appearing as powers and/or products of \mathbf{T}_2 and \mathbf{T}_3 , which will normally dominate. Higher order connected terms, \mathbf{T}_n with $n > 4$, are therefore expected to have small effects. This is consistent with the physical picture that connected \mathbf{T}_n operators correspond to n electrons interacting simultaneously. As n becomes large, this is increasingly improbable. It should be noted, however, that the higher order cluster operators ($\mathbf{T}_4, \mathbf{T}_5, \dots$) are expected to become more and more important as the number of electrons increases.

The principal deficiency of CISD is the lack of the \mathbf{T}_2^2 term, which is the main reason for CISD not being size extensive. Furthermore, this term becomes more and more important as the number of electrons increases, and CISD therefore recovers a smaller and smaller percentage of the correlation energy as the system increases. There are various approximate corrections for this lack of size extensivity that can be added to standard CISD. The most widely known of these is the *Davidson correction*, sometimes denoted CISD+Q(Davidson), where the quadruples contribution is approximated as in eq. (4.70), with a_0 being the coefficient for the HF reference wave function.

$$\Delta E_0 = (1 - a_0^2) \Delta E_{\text{CISD}} \quad (4.70)$$

If the renormalization of the wave function is also taken into account, the $(1 - a_0^2)$ quantity is divided by a_0^2 , and the corresponding correction is called the *renormalized Davidson correction*. The effect of higher order excitations is thus estimated from the correlation energy obtained at the CISD level times a factor that measures how important the single-determinant reference is at the CISD level. The Davidson correction does not yield zero for two-electron systems, where CISD is equivalent to full CI, and it is likely that it overestimates the higher order corrections for systems with few electrons. More complicated correction schemes have also been proposed,²⁰ but are rarely used.

Coupled cluster is closely connected with Møller–Plesset perturbation theory, as mentioned at the start of this section. The infinite Taylor expansion of the exponential operator (eq. (4.54)) ensures that the contributions from a given excitation level are included to infinite order. Perturbation theory indicates that doubles are the most important, since they are the only contributors to MP2 and MP3. At fourth order, there are contributions from singles, doubles, triples and quadruples. The MP4 quadruples contribution is actually the disconnected \mathbf{T}_2^2 term in the coupled cluster language, and the triples contribution corresponds to \mathbf{T}_3 . This is consistent with the above analysis, the most important is \mathbf{T}_2 (and products thereof) followed by \mathbf{T}_3 . The CCD energy is equivalent to $\text{MP}\infty(\text{D})$ where all disconnected contributions of products of doubles are included. If the perturbation series is reasonably converged at fourth order, we expect that CCD will be comparable to $\text{MP4}(\text{DQ})$, and CCSD will be comparable to $\text{MP4}(\text{SDQ})$. The MP2, MP3 and $\text{MP4}(\text{SDQ})$ results may be obtained in the first iteration for the CCSD amplitudes, allowing a direct test of the convergence of the MP series. This also points out the principal limitation of the CCSD method: the neglect of the connected triples. Including \mathbf{T}_3 in the \mathbf{T} operator leads to the CCSDT method which, as mentioned above, is too demanding computationally for all but the smallest systems. Alternatively, the triples contribution may be evaluated by perturbation theory and *added* to the CCSD results. Several such hybrid methods have been proposed, but only the method with the acronym CCSD(T) is commonly used.²¹ In this case, the triples contribution is calculated from the formula given by MP4, but using the CCSD amplitudes instead of the perturbation coefficients for the wave function corrections and adding a term arising from fifth-order perturbation theory, describing the coupling between singles and triples. Higher order hybrid methods such as CCSD(TQ), where the connected quadruples contribution is estimated by fifth-order perturbation theory, are also possible, but they are again so demanding that they can only be used for small systems.²²

As mentioned, the singles make a fairly small contribution to the correlation energy when canonical HF orbitals are used. *Brueckner theory* is a variation of coupled cluster where the orbitals used for constructing the Slater determinants are optimized such that the contribution from singles is exactly zero, i.e. $t_i^a = 0$.²³ The lowest level of Brueckner theory includes only doubles, giving the acronym BD. Although BD in theory should be slightly better than CCSD, since it includes orbital relaxation, they give in practice essentially identical results (differences between BD and CCSD are of fifth-order or higher in terms of perturbation theory). This is presumably rooted in the fact that the singles in CCSD introduce orbital relaxation.²⁴ The computational cost is

also very similar for CCSD and BD.²⁵ Similarly, BD(T) is essentially equivalent to CCSD(T),²⁶ and BD(TQ) to CCSD(TQ).

Since the singly excited determinants effectively relax the orbitals in a CCSD calculation, non-canonical HF orbitals can also be used in coupled cluster methods. This allows for example the use of open-shell singlet states (which require two Slater determinants) as reference for a coupled cluster calculation.²⁷

Another commonly used method is *Quadratic Cisd* (QCISD). It was originally derived from Cisd by including enough higher order term to make it size extensive.²⁸ It has since been shown that the resulting equations are identical to CCSD where some of the terms have been omitted.²⁹ The omitted terms are computationally inexpensive, and there appears to be no reason for using the less complete QCISD over CCSD (or QCISD(T) in place of CCSD(T)), although in practice they normally give very similar results.³⁰ There are a few other methods that may be considered either as Cisd with the addition of extra terms to make them approximately size extensive, or as approximate versions of CCSD. Some of the methods falling into this category are *Averaged Coupled-Pair Functional* (ACPF) and *Coupled Electron Pair Approximation* (CEPA). The simplest form of CEPA, CEPA-0, is also known as *Linear Coupled Cluster Doubles* (LCCD).

More recently two new intermediate coupled cluster methods have been defined, known as CC2 and CC3.³¹ As already mentioned, the single excitations allow the MOs to relax from their HF form but do not give any direct contribution to the energy due to Brillouin's theorem. For studying properties that measure the response of the energy to a perturbation, the HF orbitals are no longer optimum, and the singles are at least as important as the doubles. The CC2 method is derived from CCSD by only including the doubles contribution arising from the lowest (non-zero) order in perturbation theory, where the perturbation is defined as in MP theory (i.e. as the true electron-electron potential minus twice the average repulsion). The amplitude equations corresponding to multiplication of a doubly excited determinant in the CCSD equations (eq. (4.68)) thereby reduce to an MP2-like expression, and the t_2 amplitudes may be expressed directly in terms of the t_1 amplitudes and MO integrals. The iterative procedure therefore only involves the t_1 amplitudes. CC2 may loosely be defined as MP2 with the added feature of orbital relaxation arising from the singles. Similarly, CC3 is an approximation to the full CCSDT model, where the triples contribution is approximated by the expression arising from the lowest non-vanishing order in perturbation theory. The triples amplitudes can then be expressed directly in terms of the singles and doubles amplitudes, and MO integrals. Both in terms of computational cost and accuracy, the following progression is expected, although the CC2 and CC3 models are so new that there are few data for comparison.

$$\text{HF} \ll \text{CC2} < \text{CCSD} < \text{CC3} < \text{CCSDT} \quad (4.71)$$

Analogously to MP methods, coupled cluster theory may also be based on a UHF reference wave function. The resulting UCC methods again suffer from spin contamination of the underlying UHF, but the infinite nature of coupled cluster methods is substantially better at reducing spin contamination relative to UMP.³² Projection methods analogous to the PUMP case have been considered but are not commonly used. ROHF-based coupled cluster methods have also been proposed but appear to give results very similar to UCC, especially at the CCSD(T) level.³³

Standard coupled cluster theory is based on a single-determinant reference wave function. It suffers from the same problem as MP, in that it works best if the zeroth-order wave function is sufficiently “good”. Owing to the summation of contributions to infinite order, however, coupled cluster is somewhat more tolerant to a poor reference wave function than MP methods. Since the singly excited determinants allow the MOs to relax in order to describe the multi-reference character in the wave function, the magnitude of the singles amplitude is an indication of how good the HF single determinant is as the reference. The T_1 -diagnostic defined as the norm of the singles amplitude vector divided by the square root of the number of electrons has been suggested as an internal evaluation of the quality of a CCSD wave function.³⁴

$$T_1 = \frac{1}{\sqrt{N_{\text{elec}}}} |\mathbf{t}_1| \quad (4.72)$$

Specifically, if $T_1 < 0.02$, the CCSD(T) method is expected to give results close the full CI limit for the given basis set. If T_1 is larger than 0.02, it indicates that the reference wave function has significant multi-determinant character, and multi-reference coupled cluster should preferentially be employed. Such methods are being developed,³⁵ but have not yet seen any extensive use. The T_1 -diagnostic in eq. (4.70) is not completely independent of the system size, and other diagnostics have also been proposed.³⁶

4.10.1 Illustrating correlation methods for the beryllium atom

The beryllium atom has four electrons ($1s^2 2s^2$ electron configuration) and the ground state wave function contains significant multi-reference character owing to the presence of the low-lying 2p-orbital. The correlation energies calculated with MP, CI and CC methods in a 4s2p basis set (cc-pVDZ basis augmented with one set of tight s - and p -functions for correlating the 1s-electrons) are given in Table 4.5.³⁷

Table 4.5 Correlation energies for the beryllium atom in a 4s2p basis set

Level	ΔE_{corr} (au)	%	Level	ΔE_{corr} (au)	%	Level	ΔE_{corr} (au)	%
MP2	0.053174	67.85						
MP3	0.067949	86.70	CISD	0.075277	96.05	CCSD	0.078176	99.75
MP4	0.074121	94.58				CCSD(T)	0.078361	99.99
MP5	0.076918	98.15	CISDT	0.075465	96.29	CCSDT	0.078364	99.99
MP6	0.078090	99.64						
MP7	0.078493	100.15	CISDTQ	0.078372	100	CCSDTQ	0.078372	100

Since beryllium only has four electrons, CISDTQ is a full CI treatment and completely equivalent to a CCSDTQ calculation. The multi-reference character displays itself as a relatively slow convergence of the perturbation series, with millihartree accuracy being attained at the MP6 level and inclusion of terms up to MP20 is required in order to converge the energy to within 10^{-6} au of the exact answer. Note also that the correlation energy is overestimated at order seven, i.e. the perturbation series oscillates at higher orders. The contribution from triply excited states is minute, as expected

Table 4.6 Coefficients for dominating excited states

Excitation	Type	a_{CISDTQ}	t_{CCSDTQ}	$t_{\text{CCSDTQ}} \cdot t_{\text{CCSDTQ}}$
$2s^2 \rightarrow 2p^2$	D	-0.18612	-0.18523	
$2s^2 \rightarrow 2s'^2$	D	-0.04341	-0.04376	
$1s^2 \rightarrow 1s'^2$	D	-0.02171	-0.02170	
$1s^2 2s^2 \rightarrow 1s'^2 2p^2$	Q	0.00407	$4 \cdot 10^{-7}$	0.00402
$1s^2 2s^2 \rightarrow 1s'^2 2s'^2$	Q	0.00092		0.00095

for a system with two well-separated electron pairs, i.e. CISDT is only a marginal improvement over CISD.

The coefficients in the CISDTQ and CCSDTQ wave functions (using intermediate normalization) for the dominating excitations are given in Table 4.6.

There is little difference between the full CI and CCSD coefficients for the three most important doubly excited states. The quadruply excited states enter the CI wave function with non-negligible weights, contributing ~4% of the correlation energy (Table 4.5), but as shown in the last column of Table 4.6, these contributions are estimated very well by the product terms in the CC wave function. The CCSD energy in Table 4.5 and the t_4 amplitude in Table 4.6 show that the quadruply excited states in the CI wave function are mainly of the product type, and not a true quadruply excited state. It is this feature that makes CC superior to CI-based methods.

4.11 Methods Involving the Interelectronic Distance

The necessity of going beyond the HF approximation is due to the fact that electrons are further apart than described by the product of their orbital densities, i.e. their motions are correlated. This arises from the electron–electron repulsion operator, which is a sum of terms of the type shown in eq. (4.73).

$$\frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} = \frac{1}{r_{12}} \quad (4.73)$$

Without these terms, the Schrödinger equation can be solved exactly, with the solution being a Slater determinant composed of orbitals.

The electron–electron repulsion operator has a *singularity* for $r_{12} = 0$ which results in the exact wave function having a *cusp* (discontinuous derivative),³⁸ since the kinetic energy must cancel the infinity of the potential energy to give a finite result.

$$\left(\frac{\partial \Psi}{\partial r_{12}} \right)_{r_{12}=0} = \frac{1}{2} \Psi(r_{12}) \quad (4.74)$$

The cusp condition implies that the exact wave function must be linear in the interelectronic distance for small values of r_{12} .

$$\Psi_{\text{exact}}(r_{12}) = \text{constant} + \frac{1}{2} r_{12} + \dots \quad (4.75)$$

It would therefore seem natural that the interelectronic distance should be a necessary variable for describing electron correlation. For two-electron systems, extremely

accurate wave functions may be generated by taking a trial wave function consisting of an orbital product times an expansion in electron coordinates, as given in eq. (4.76), and variationally optimizing the α_i and C_{klm} parameters.

$$\Psi(r_1, r_2) = e^{-\alpha_1 r_1} e^{-\alpha_2 r_2} \sum_{klm} C_{klm} (r_1 + r_2)^k (r_1 - r_2)^l r_{12}^m \quad (4.76)$$

Expansions such as eq. (4.76) are known as *Hylleraas* type wave functions.³⁹ For the hydrogen molecule, it is possible to converge the total energy to $\sim 10^{-9}$ au, which is more accurate than what can be determined experimentally. In fact, the prediction that the experimental dissociation energy for H_2 was wrong, based on calculations, was one of the first hallmarks of quantum chemistry.⁴⁰ Such wave functions unfortunately become impractical for more than 3–4 electrons.

All electron correlation methods based on expanding the N -electron wave function in terms of Slater determinants built from orbitals (one-electron functions) suffer from an agonizingly slow convergence. Literally millions or billions of determinants are required for obtaining results that in an absolute sense are close to the exact results. This is due to the fact that products of one-electron functions are poor at describing the cusp behaviour of the wave function when two electrons are close together. At the second-order perturbation level (i.e. MP2) it may be shown that the *error* in the correlation energy behaves asymptotically as $(l + 1/2)^{-4}$, where l is the highest angular momentum in the basis set. For a general wave function the convergence is $(l + 1/2)^{-4} + (l + 1/2)^{-5} + (l + 1/2)^{-6} + \dots$. This means that the total energy will converge as $(L + 1)^{-3} + (L + 1)^{-4} + (L + 1)^{-5} + \dots$, if the basis set is saturated up to angular momentum L .⁴¹ For sufficiently large values of L the convergence is thus $\sim (L + 1)^{-3}$, which is quite slow.

In order to achieve a high accuracy, it would seem desirable to explicitly include terms in the wave functions that are linear in the interelectronic distance. This is the idea in the R12 methods developed by Kutzelnigg and coworkers.⁴² The first-order correction to the HF wave function only involves doubly excited determinants (eq. (4.38)). In R12 methods additional terms are included, which is essentially the HF determinant multiplied with r_{ij} factors.

$$\Psi_{R12} = \Phi_{HF} + \sum_{ijab} a_{ijab} \Phi_{ij}^{ab} + \sum_{ij} b_{ij} r_{ij} \Phi_{HF} \quad (4.77)$$

The exact definition is slightly more complicated, since the wave function has to be properly antisymmetrized and projected onto the actual basis but, for illustration, the above form is sufficient. Such R12 wave functions may then be used in connection with the CI, MBPT or CC methods described above. Consider for example a CI calculation with an R12 type wave function. The energy is given by eq. (4.78), where the a_{ijab} and b_{ij} parameters in (4.77) are optimized variationally.

$$E = \langle \Psi_{R12} | \mathbf{H} | \Psi_{R12} \rangle \quad (4.78)$$

The overwhelming problem is that matrix elements from eq. (4.78) now involve integrals depending on three and four electron coordinates. Consider for example the following terms arising from the r_{ij} operator written out in terms of the one- and two-electron operators (\mathbf{h} and \mathbf{g} , eq. (3.24)).

$$\begin{aligned}\langle \Phi_{\text{HF}} | \mathbf{H} | r_{ij} \Phi_{\text{HF}} \rangle &= \langle \Phi_{\text{HF}} | \mathbf{h} | r_{ij} \Phi_{\text{HF}} \rangle + \langle \Phi_{\text{HF}} | \mathbf{g} | r_{ij} \Phi_{\text{HF}} \rangle \\ \langle r_{ij} \Phi_{\text{HF}} | \mathbf{H} | r_{ij} \Phi_{\text{HF}} \rangle &= \langle r_{ij} \Phi_{\text{HF}} | \mathbf{h} | r_{ij} \Phi_{\text{HF}} \rangle + \langle r_{ij} \Phi_{\text{HF}} | \mathbf{g} | r_{ij} \Phi_{\text{HF}} \rangle\end{aligned}\quad (4.79)$$

The \mathbf{g} operator leads to integrals over molecular orbitals of the type shown in eq. (4.80).

$$\begin{aligned}& \left\langle \phi_i(1)\phi_j(2)\phi_k(3) \left| \frac{r_{12}}{r_{13}} \right| \phi_i(1)\phi_j(2)\phi_k(3) \right\rangle \\ & \left\langle \phi_i(1)\phi_j(2)\phi_k(3) \left| \frac{r_{12}r_{23}}{r_{13}} \right| \phi_i(1)\phi_j(2)\phi_k(3) \right\rangle \\ & \left\langle \phi_i(1)\phi_j(2)\phi_k(3)\phi_l(4) \left| \frac{r_{12}r_{34}}{r_{23}} \right| \phi_i(1)\phi_j(2)\phi_k(3)\phi_l(4) \right\rangle\end{aligned}\quad (4.80)$$

Not only are such integrals difficult to calculate but, when the MOs are expanded in a basis set consisting of M_{basis} AOs, there will be of the order of M_{basis}^6 three-electron integrals and of the order of M_{basis}^8 four-electron integrals. Such methods are therefore inherently more expensive than for example the full CCSDT model.

The trick for turning the R12 method into a viable computational tool is to avoid calculating the three- and four-electron integrals, without jeopardizing the accuracy. In a complete basis, a three-electron integral may be written in terms of products of two-electron integrals by inserting a “resolution of the identity” between the two operators.

$$\begin{aligned}\mathbf{1} &= \sum_p |\phi_p\rangle \langle \phi_p| = \sum_p |\phi_p \phi_q \phi_r\rangle \langle \phi_p \phi_q \phi_r| \\ \left\langle \phi_i \phi_j \phi_k \left| \frac{r_{12}}{r_{13}} \right| \phi_i \phi_j \phi_k \right\rangle &= \sum_{pqr} \langle \phi_i \phi_j \phi_k | r_{12} | \phi_p \phi_q \phi_r \rangle \left\langle \phi_p \phi_q \phi_r \left| \frac{1}{r_{13}} \right| \phi_i \phi_j \phi_k \right\rangle \\ &= \sum_{pqr} (\delta_{kr} \langle \phi_i \phi_j | r_{12} | \phi_p \phi_q \rangle) \left(\delta_{qj} \left\langle \phi_p \phi_r \left| \frac{1}{r_{13}} \right| \phi_i \phi_k \right\rangle \right) \\ &= \sum_p \langle \phi_i \phi_j | r_{12} | \phi_p \phi_j \rangle \left\langle \phi_p \phi_k \left| \frac{1}{r_{13}} \right| \phi_i \phi_k \right\rangle\end{aligned}\quad (4.81)$$

The first reduction occurs since the r_{12} and r_{13}^{-1} operators only involve two electron coordinates, the second reduction is due to the two delta functions. Three- and four-electron integrals can therefore be written as a sum over products of integrals involving only two electron coordinates. In a finite basis set, the resolution is not exact, and the identities in eq. (4.81) become approximations. The beauty of the R12 methods is that this error can be controlled, albeit at the price of calculating and handling a significantly larger number (and different types) of two-electron integrals. In the original method, the basis set for the resolution of the identity was the same as for expanding the orbitals and therefore needed to be large for the identity resolution to be reasonably fulfilled.⁴³ In more recent work an auxiliary basis set was used for the identity resolution, which significantly improved the computational efficiency.⁴⁴

The significance of R12 methods is that the energy error in terms of angular momentum of the basis set now behaves approximately as $(L + 1)^{-7}$, which is a significant improvement over standard methods. It should be noted that in the limit of a complete basis set the MP2-R12 (for example) will give the same result as a traditional

MP2 calculation, i.e. the R12 approach speeds up the basis set convergence, but does not change the fundamental characteristics of the MP2 method. The drawback is that R12 methods work inefficiently with small basis sets. While the convergence changes from $(L + 1)^{-3}$ to $(L + 1)^{-7}$, the effective improvement only becomes significant for quite high L values. Recent work has investigated whether other correlation factors may be incorporated in order to speed up the basis set convergence.⁴⁵

4.12 Direct Methods

Conventional HF methods rely on storing the two-electron integrals over atomic orbitals on disk, and reading them in each SCF iteration, while direct methods generate the integrals as they are needed (Section 3.8.5). This is an easy change in algorithm since the HF energy is expressed directly in terms of AO integrals. Methods involving electron correlation, however, require matrix elements between Slater determinants, which can be expressed in terms of integrals over MOs (eq. (4.8)). Conventional methods for the integral transformation (Section 4.2.1) read the AOs, perform the multiplications with the MO coefficients (eq. (4.12)), and write the MO integrals to disk. These can then be read in and used in the correlation treatment. Although the number of MO integrals typically is somewhat smaller than the number of AO integrals (for example MO integrals involving four virtual orbitals may not be needed), the disk space requirements are still significant if more than a few hundred basis functions are used. To eliminate the disk space requirements, and remove the relatively inefficient data transfer step for reading/writing to disk, it is desirable also to have direct algorithms for electron correlation methods. Direct in this context means that the integrals are calculated as needed and then discarded. The need for integrals over MOs instead of AOs, however, makes the development of direct methods in electron correlation somewhat more complicated than at the HF level.

Consider for example the MP2 energy expression given in eq. (4.82).⁴⁶

$$E(\text{MP2}) = \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{vir}} \frac{(\langle \phi_i \phi_j | \phi_a \phi_b \rangle - \langle \phi_i \phi_j | \phi_b \phi_a \rangle)}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (4.82)$$

The MO integrals are given in eq. (4.83).

$$\langle \phi_i \phi_j | \phi_k \phi_l \rangle = \sum_{\alpha}^{M_{\text{basis}}} \sum_{\beta}^{M_{\text{basis}}} \sum_{\gamma}^{M_{\text{basis}}} \sum_{\delta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta j} c_{\gamma k} c_{\delta l} \langle \chi_{\alpha} \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle \quad (4.83)$$

Since each MO integral in principle contains contributions from *all* the AO integrals, a straightforward calculation of an MO integral each time it is needed will involve a generation of all the AO integrals. In other words, it would be necessary to recalculate the AO integrals $\sim O^2 V^2$ times (O and V being the number of occupied and virtual orbitals, respectively), compared with the 15–20 times in an SCF calculation. The MP2 method would therefore change from being an M_{basis}^5 to an M_{basis}^8 method, which clearly is an unacceptably large penalty for a direct method.

The M_{basis}^8 dependence is a consequence of performing the four index transformation with all four indices at once. As shown in Section 4.2.1, it is advantageous to perform the transformation one index at a time.

$$\begin{aligned}
\langle \phi_i \phi_j | \phi_a \phi_b \rangle &= \sum_{\delta}^{M_{\text{basis}}} c_{\delta b} \langle \phi_i \phi_j | \phi_a \chi_{\delta} \rangle \\
\langle \phi_i \phi_j | \phi_a \chi_{\delta} \rangle &= \sum_{\gamma}^{M_{\text{basis}}} c_{\gamma a} \langle \phi_i \phi_j | \chi_{\gamma} \chi_{\delta} \rangle \\
\langle \phi_i \phi_j | \chi_{\gamma} \chi_{\delta} \rangle &= \sum_{\beta}^{M_{\text{basis}}} c_{\beta j} \langle \phi_i \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle \\
\langle \phi_i \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle &= \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \langle \chi_{\alpha} \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle
\end{aligned} \tag{4.84}$$

By choosing the right order of the transformation the scaling can be reduced considerably. In eq. (4.84) the indices corresponding to the occupied orbitals may be transformed before the virtuals. There are of the order of M_{basis}^4 of the AO integrals, $\langle \chi_{\alpha} \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle$, but only OM_{basis}^3 of the quarter transformed integrals, $\langle \phi_i \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle$. Instead of storing and reading the AO integrals from the SCF step, they can be recalculated in the transformation step, reducing the storage from M_{basis}^4 to OM_{basis}^3 . The subsequent quarter transformations require less storage, i.e. the next transformation with an occupied index reduces the number of integrals to $O^2M_{\text{basis}}^2$, the third to O^2VM_{basis} , and the last to O^2V^2 . Since the MP2 energy can be written as a sum of contributions from each occupied orbital, the occupied orbitals can be treated one at a time, i.e. first sum all contributions of $\langle \phi_i \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle$ then $\langle \phi_2 \chi_{\beta} | \chi_{\gamma} \chi_{\delta} \rangle$, etc. This reduces the necessary storage to only order M_{basis}^3 . It may be further reduced to $OV M_{\text{basis}}$ by proper scheduling of the evaluation order of the remaining three indices. The $OV M_{\text{basis}}$ number of integrals is much less than the original M_{basis}^4 , and will in many cases fit into memory. The net result is that disk storage is effectively eliminated, or at least greatly reduced. If only one occupied orbital is treated at a time, O integral evaluations are required, however, the more memory that is available, the more of the occupied orbitals can be treated in a single sweep, decreasing the number of integral evaluations.

The above is an example of how direct algorithms may be formulated for methods involving electron correlation. It illustrates that it is not as straightforward to apply direct methods at the correlated level, as at the SCF level. However, the steady increase in CPU performance, and especially the evolution of multi-processor machines, favours direct (and semi-direct where some intermediate results are stored on disk) algorithms.⁴⁷

4.13 Localized Orbital Methods

Ab initio calculations involving electron correlation usually build on a set of canonical HF orbitals and this leads to a computational effort that increases as a rather high power of the system size, i.e. M^5 – M^8 . Considering that the fundamental physical force is only between pairs of particles, this scaling is “non-physical”. One of the reasons for the high scaling is the fact that canonical orbitals are delocalized over the whole molecule, i.e. essentially all orbitals make a (small) contribution to the wave function for a specific part of the molecule. This suggested that a set of localized orbitals may be a better starting point, since a single or only a few orbitals would then contribute the large majority at a given point, and the remaining contributions could simply be neglected. Alternatively, the problem may be formulated directly in the atomic orbital

basis, since the basis functions are naturally localized on a single atom. Such local MP2 and local CC methods have started to appear, but are not yet commonly used.⁴⁸ These methods are somewhat more complicated to formulate as the Fock matrix is only diagonal in the canonical orbitals. Nevertheless, methods based on localized orbitals hold the promise of a near-linear scaling with problem size in the large-scale limit. It is at present not clear exactly how large the systems need to be to reach the “large-scale” limit, but the present methods appear to have cross-over points in terms of computational resources in the few hundred atoms region.

The use of localized orbitals will only lead to a computational saving if it is combined with criteria for neglecting a (large) fraction of the terms, and such criteria all involve one or more cutoff distances. A formal problem with local orbital-based methods is the risk of producing non-continuous energy surfaces, e.g. during a geometry optimization a given atomic distance may increase beyond one of the cutoff values, and some of the correlation contributions suddenly drop to zero. The energy will thus experience a discontinuity upon increasing an interatomic distance by an infinitesimal amount, leading to corresponding discontinuities in the gradient and second derivatives. This in turn may lead to problems in the geometry optimization and/or vibrational frequencies. These errors can of course be controlled by choosing cutoff distances sufficiently large that the discontinuities are well below chemical significance, but large cutoff distances are counterproductive from a computational efficiency point of view, and a compromise is necessary. It is at present unclear how large this problem is in practice.⁴⁹ Another aspect is that the total energy is often only accurate to a few millihartrees (~ 1 kJ/mol), and different implementations may thus give slightly different results. Finally, it should be noted that there are certain systems, typically having stretched bond (TS) or being aromatic, where it is difficult to obtain localized orbitals, or where more than one set of localized orbitals are possible, and a small geometry change may in such cases suddenly switch from one localization to another.

Another method for reducing the computational cost relies on the so-called “resolution of the identity” technique, where the calculation of four-index integrals is replaced by three- and two-index quantities, via the use of an auxiliary basis set. Although the formal scaling is unchanged, the computational prefactor is significantly reduced, leading to an overall efficiency gain of approximately an order of magnitude.⁵⁰ These methods furthermore have the appealing feature that the efficiency gain increases with the size of the basis set, e.g. for basis sets such as cc-pVQZ the gain is over two orders of magnitude. A similar idea is used in the Cholesky decomposition method, where the number of required two-electron integrals for a given accuracy is reduced by decomposition of the full two-integral matrix,⁵¹ and in methods where advanced integral screening protocols are used.⁵² The divide and conquer method has also been used in correlated calculations, where the full system is broken into smaller subsections that are treated separately, and the results subsequently assembled to that of the full system.⁵³

4.14 Summary of Electron Correlation Methods

The only generally applicable methods are CISD, MP2, MP3, MP4, CCSD and CCSD(T). CISD is variational, but not size extensive, while MP and CC methods are non-variational but size extensive. CISD and MP are in principle non-iterative

methods, although the matrix diagonalization involved in CISD is usually so large that it has to be done iteratively. Solution of the coupled cluster equations must be done by an iterative technique since the parameters enter in a non-linear fashion. In terms of the most expensive step in each of the methods they may be classified according to how they formally scale in the large system limit, as shown in Table 4.7.

Table 4.7 Limiting scaling in terms of basis set size M for various methods

Scaling	CI methods	MP methods	CC methods (iterative)
M^5	CIS	MP2	CC2
M^6	CISD	MP3	CCSD
M^7		MP4	CC3, CCSD(T)
M^8	CISDT	MP5	CCSDT
M^9		MP6	
M^{10}	CISDTQ	MP7	CCSDTQ

We have so far been careful to use the wording “formal scaling”. As already discussed, HF is formally an M^4 method but in practice the scaling may be reduced all the way down to M^1 . Similarly, MP2 is formally an M^5 method. However, an MP2 calculation consists of three main parts: the HF calculation, the AO to MO integral transformation, and the MP2 energy calculation. Only the second part has a formal scaling of M^5 , the others are (formal) M^4 steps. In the large system limit, the transformation required for the MP2 procedure *will* become the most expensive step, however, in practice where calculations may be restricted to a few hundred basis functions, it is often observed that the MP2 step takes less time than the HF step. The formal scaling only indicates what the rate limiting step will be in the large system limit. Whether this limit actually is reached in practical calculations is another matter.

The lower value of M^5 scaling for methods involving electron correlation arises from the transformation of the two-electron integrals from the AO to MO basis, but if the transformation is carried out with one of the indices belonging to an occupied MO first, the scaling is actually the number of occupied orbitals (O) times M^4 . If we consider making the system larger by doubling the fundamental unit (for example calculations on a series of increasingly larger water clusters), keeping the basis set per atom constant, O scales linearly with M , and we arrive at the M^5 scaling. This assumption (increasing system size) is the basis for Table 4.7. More often, however, a series of calculations is performed on the same system with increasingly larger basis sets. In this case, the number of electrons (occupied orbitals) is constant and the scaling is M^4 . Many of the commonly employed methods for electron correlation (including for example MP2, MP3, MP4, CISD, CCSD and CCSD(T)) scale in fact as M^4 when the number of occupied orbitals is constant.

In terms of accuracy with a medium-sized basis set the following order is often observed.

$$\text{HF} \ll \text{MP2} < \text{CISD} < \text{MP4(SDQ)} \sim \text{CCSD} < \text{MP4} < \text{CCSD(T)} \quad (4.85)$$

All of these are single-determinant-based methods. Multi-reference methods cannot easily be classified as the quality of the results depends heavily on the size of the

reference. A two-configurational reference is only a slight improvement over HF, but including all configurations generates a full CI. The ordering above is only valid when the HF reference is a “good” zeroth-order description of the system. The more multi-reference character in the wave function, the better the “infinite” order coupled cluster performs relative to perturbation methods.

MP3 has not been included in the above comparison. As already mentioned, MP3 results are often inferior to those at MP2. In fact, MP2 often gives surprisingly good results, especially if large basis sets are used.⁵⁴ Furthermore, it should be kept in mind that the MP perturbation series in many cases may actually be divergent, although corrections carried out to low order (i.e. 2–4) rarely display excessive oscillations.

HF results should by modern standards be considered as model calculations, like semi-empirical methods such as AM1 and PM3. Minimal basis HF calculations often give results that are worse than AM1 or PM3, but at a computational cost of maybe 100 times as much. Medium and large basis set HF calculations usually do not give absolute results that are particularly close to experimental values, but since the errors to a certain degree are systematic (such as all vibrational frequencies being overestimated by ~10%), they can be used with more or less “empirical” corrections to treat systems for which correlated calculations are not possible. The distinct advantage of *ab initio* methods is the ability to treat all systems at an equal level of accuracy, independent of whether experimental data exist or not. A detailed assessment of the level of accuracy that can be expected at a given level of theory is difficult to establish as it is heavily dependent on the quality of the basis set. Given a sufficiently large basis set, however, the CCSD(T) method is able to meet the goal of an accuracy of ~4 kJ/mol (~1 kcal/mol) for most systems. Even with less complete methods (such as MP4) and medium-size basis sets such as DZP or TZP, it is often possible to get accuracies of the order of a few tens of kJ/mol.

The use of CI methods has been declining in recent years at the expense of MP and especially CC methods. It is now recognized that size extensivity is important for obtaining accurate results. Excited states, however, are somewhat difficult to treat by perturbation or coupled cluster methods, and CI- or MCSCF-based methods have been the preferred methods here. More recently linear response methods (Section 10.9) have been developed for coupled cluster wave functions, and which allow calculation of excited state properties.

Finally, a few words on the size of systems that can be treated. The limiting parameters will again be taken as the number of basis functions although, as noted above, a more detailed breakdown in terms of occupied and virtual MOs can be done. Note also that a given limit in terms of basis functions may translate either into a large molecular system with a small basis per atom or a small molecular system with a very large basis set on each atom. The ordering in eq. (4.85) suggests three levels of electron correlation: none (HF), MP2, or extended (MP4 or CCSD(T)). HF methods are in general possible with up to ~5000 basis functions, MP2 is fairly routine up to ~800 basis functions, while the advanced correlation methods are limited to ~300–400 basis functions. With a DZP basis set these values translate into roughly 200, 30 and 10 CH₂ fragments, respectively. The limits hold for just calculating the energy at a single geometry. If more advanced features are desired, such as optimizing the geometry or calculating frequencies, the limits drop to roughly half of the above. Unfortunately, MP2 and higher correlated methods require basis sets larger than DZP to fully exploit the

inherent accuracy in these methods, and this further reduces the size of the systems that can be handled.

With the continuing advances of computer hardware and more efficient algorithms, these limits are gradually being shifted upwards. Owing to the rather steep scaling with system size, however, they will (barring a fundamental breakthrough) give a rough idea of the size of systems that can be handled also in the future. Currently the speed of computer hardware improves by a factor of two in a timespan of about 18 months. In other words, a factor of 10 in terms of performance for the same price is gained roughly every 5 years. Owing to the scaling between 4 and 7 of the various methods, however, a factor of 10 increase in raw speed only translates into an increase of system size of 1.7 (M^4 scaling) or 1.4 (M^7 scaling). Linear scaling methods in Hartree–Fock methods will of course benefit fully from increased computational speed.

4.15 Excited States

The development of HF and correlated methods in the previous chapters has focused on the electronic ground state. In some cases it is also of interest to consider electronically excited states. It is useful to distinguish between two cases, depending on whether the excited state has the *same* or a *different* symmetry than the lower state(s). The different symmetry case is easy to handle, as the lowest energy state of a given symmetry may be handled completely analogously to the ground state. An HF wave function may be obtained by a proper specification of the occupied orbitals, and the resulting wave function can be improved by adding electron correlation by for example CI, MP or CC methods. The only caveat may be that the state is an open shell, which often requires a (small) MCSCF wave function for an adequate zeroth-order description.

Excited states having lower energy solutions of the same symmetry are somewhat more difficult to treat. It is difficult to generate an HF type wave function for such states, as the variational optimization will collapse to the lowest energy solution of the given symmetry. The lack of a proper HF solution means that perturbation and coupled cluster methods are not well suited for calculating excited states, although excited state *properties* (for example excitation energies) may be calculated *directly* with response methods (Section 10.9). Response methods can be based on for example a coupled cluster wave function. It is, however, relatively easy to generate higher energy states by CI methods: this simply corresponds to using the $(n + 1)$ th eigenvalue from the diagonalization of the CI matrix as a description of the n th excited state (the second root is the first excited state, etc.). Such a CI procedure will normally employ a set of HF orbitals from a calculation on the lowest energy state, and the CI procedure is therefore biased against the excited states.

The simplest description of an excited state is the orbital picture where one electron has been moved from an occupied to an unoccupied orbital, i.e. an S-type determinant as illustrated in Figure 4.2. The lowest level of theory for a qualitative description of excited states is therefore a CI including only the singly excited determinants, denoted CIS. CIS gives wave functions of roughly HF quality for excited states, since no orbital optimization is involved. For valence excited states, for example those arising from excitations between π -orbitals in an unsaturated system, this may be a reasonable description. There are, however, normally also quite low-lying states that essentially

correspond to a double excitation, and those require the inclusion of at least the doubles as well, i.e. CISD.

A more balanced description requires MCSCF-based methods where the orbitals are optimized for each particular state, or optimized for a suitable average of the desired states (*state-averaged* MCSCF). It should be noted that such excited state MCSCF solutions correspond to saddle points in the parameter space for the wave function, and second-order optimization techniques are therefore almost mandatory. In order to obtain accurate excitation energies it is normally necessarily also to include dynamical correlation, for example by the CASPT2 method.

Excited states involve electrons that are more loosely bound than in the ground state, and they thus usually require basis sets with diffuse functions for a proper description. This is especially true for so-called *Rydberg* states, which may be considered as an electron orbiting a positively charged molecule. Such states resemble an atomic system, with the molecular cation playing the rule of the nucleus, and can be characterized as having s-, p-, d- etc. character. Rather than using a regular basis set with diffuse functions on each nucleus, such Rydberg states can be modelled by having a single set of diffuse functions located at the molecular centre of mass.⁵⁵

4.16 Quantum Monte Carlo Methods

Monte Carlo methods refer to techniques for obtaining the value of a multi-dimensional integral of a function by randomly probing its value within the whole variable space, and estimating the integral by statistical averaging. In the limit of an infinite number of sampling points, the result is identical to that obtained from an analytical integration, but for a finite number of points, the calculated value is given as an average with an associated standard deviation. The standard deviation, the uncertainty, depends inversely on the square root of the number of sampling points.

Since the square of the wave function represent a probability function, the associated energy can be calculated by *Quantum Monte Carlo* (QMC) methods. For a (approximate) variational wave function, the energy can be re-written as in eq. (4.86).

$$\begin{aligned}
 E &= \frac{\langle \Phi | \mathbf{H} | \Phi \rangle}{\langle \Phi | \Phi \rangle} = \frac{\int \Phi^* \mathbf{H} \Phi d\mathbf{r}}{\int \Phi^* \Phi d\mathbf{r}} \\
 E &= \frac{\int \Phi^* \Phi (\Phi^{-1} \mathbf{H} \Phi) d\mathbf{r}}{\int \Phi^* \Phi d\mathbf{r}} = \frac{\int |\Phi(\mathbf{r})|^2 (\Phi^{-1} \mathbf{H} \Phi) d\mathbf{r}}{\int |\Phi(\mathbf{r})|^2 d\mathbf{r}} \quad (4.86) \\
 E &= \int E_{\text{local}}(\mathbf{r}) P(\mathbf{r}) d\mathbf{r} \quad ; \quad P(\mathbf{r}) = \frac{|\Phi(\mathbf{r})|^2}{\int |\Phi(\mathbf{r})|^2 d\mathbf{r}} \quad ; \quad E_{\text{local}}(\mathbf{r}) = \Phi^{-1} \mathbf{H} \Phi
 \end{aligned}$$

The last equation shows that the energy can be calculated as an integral of the local energy function $\Phi^{-1} \mathbf{H} \Phi$ weighted with the probability density P . In principle, this integral could be calculated by numerical quadrature methods, such as the Simpson's trapezoidal rule, but this becomes very inefficient when the number of variables is large. For a system with N electrons, the dimensionality of the problem is $3N_{\text{elec}}$, and the integral

can be estimated much more efficiently by sampling the function point-wise within the whole function space. Estimating the functional value by a random sampling of points within the integration limits, weighted by the probability factors, is called *variational QMC*.⁵⁶ The generation of points is done using a Metropolis algorithm, as discussed in more detail in Section 14.1, and the calculated energy is simply the average of the local energies over the sampling points.

$$E = \frac{1}{M_{\text{point}}} \sum_{i=1}^{M_{\text{point}}} E_{\text{local}}(\mathbf{r}_i) \quad (4.87)$$

An improvement of the variational QMC can be obtained by the *diffusion* QMC approach. Consider the time-dependent Schrödinger equation, where the time is replaced with an imaginary time variable $\tau = it$.

$$\begin{aligned} i \frac{\partial \Phi(\mathbf{r}, t)}{\partial t} &= \mathbf{H} \Phi(\mathbf{r}, t) \\ - \frac{\partial \Phi(\mathbf{r}, \tau)}{\partial \tau} &= \mathbf{H} \Phi(\mathbf{r}, \tau) \end{aligned} \quad (4.88)$$

For a free electron, the Hamiltonian is only kinetic energy, and the resulting equation is identical to that describing a diffusion process.

$$\frac{\partial \Phi(\mathbf{r}, \tau)}{\partial \tau} = \frac{1}{2} \nabla^2 \Phi(\mathbf{r}, \tau) \quad (4.89)$$

Addition of a potential energy results in a generalized diffusion equation.

$$\frac{\partial \Phi(\mathbf{r}, \tau)}{\partial \tau} = \frac{1}{2} \nabla^2 \Phi(\mathbf{r}, \tau) - \mathbf{V}(\mathbf{r}) \Phi(\mathbf{r}, \tau) \quad (4.90)$$

The generalized diffusion equation can be solved by a random walk procedure and, in the long time limit, the resulting distribution converges to the ground state wave function. This can be seen by expanding an approximate wave function in terms of the exact wave functions (eq. (1.20)).

$$\Phi(\mathbf{r}, \tau) = \sum_k c_k \Psi_k(\mathbf{r}, \tau) = \sum_k c_k \Psi_k(\mathbf{r}) e^{-E_k \tau} \quad (4.91)$$

The exponential dependence on the energy means that the high energy states decay faster than the low energy ones and, in the long (imaginary) time limit, only the ground state wave function survives.

The main problem with QMC methods is the requirement of an antisymmetric wave function, since the electrons are fermions. The antisymmetry means that the wave function has both positive and negative regions, and consequently $3N_{\text{elec}} - 1$ dimensional surfaces where the wave function is zero. These surfaces are called *nodes*, and correspond to zero-probability regions of space. Clearly, a procedure that indiscriminately samples the nodal regions will yield inaccurate answers. QMC methods thus require a guiding function, a trial wave function, for determining how to sample the huge phase space most efficiently and in agreement with the fermion nature of the electrons.

A suitable trial wave function can be constructed from a Hartree–Fock wave function multiplied with a suitable correlation function, often taken as a *Jastrow* factor $J(\mathbf{r})$.

$$J(\mathbf{r}) = \sum_i^{N_{\text{elec}}} \chi(\mathbf{r}_i) - \sum_{i>j}^{N_{\text{elec}}} u(\mathbf{r}_i, \mathbf{r}_j) \quad (4.92)$$

The functional forms of the one- and two-electrons terms χ and u are chosen such that they model the nuclear–electron and electron–electron cusp conditions, respectively, and the parameters inherent in these functions are variationally optimized by the QMC procedure.

In order to maintain the wave function antisymmetry, the diffusion QMC is normally used within the *fixed node* approximation, i.e. the nodes are fixed by the initial trial wave function. Unfortunately, the location of nodes for the exact wave function is far from trivial to determine, although simple approximations such as HF can give quite reasonable estimates.⁵⁷ The fixed node diffusion QMC thus determines the best wave function with the nodal structure of the initial trial wave function. If the trial wave function has the correct nodal structure, the QMC will provide the exact solution to the Schrödinger equation, including the electron correlation energy. It should be noted that the region near the nuclei contributes most to the statistical error in QMC methods, and in many applications the core electrons are therefore replaced by a pseudopotential.

The scaling of QMC methods is N_{basis}^3 , but the prefactor makes these methods roughly two orders of magnitude more expensive than independent-particle models such as HF and DFT. The relatively low-order scaling, however, makes QMC competitive with for example coupled cluster methods even for relatively small systems. The main disadvantage of QMC is the statistical error in the calculated results, which only decays as the inverse square root of the number of sampling points. Generating highly accurate results is thus computationally expensive, although the calculations are well suited for running on large parallel computers. Furthermore, the statistical uncertainty makes it difficult to calculate nuclear forces and second derivatives, which are essential for optimizing structures and calculating vibrational frequencies. Finally, the accuracy of the results is tightly coupled to the form of the trial wave function, and a poor trial wave function can generate poor-quality results.

References

1. A. Szabo, N. S. Ostlund, *Modern Quantum Chemistry*, McGraw-Hill, 1982; R. McWeeny, *Methods of Molecular Quantum Mechanics*, Academic Press, 1992; W. J. Hehre, L. Radom, J. A. Pople, P. v. R. Schleyer, *Ab Initio Molecular Orbital Theory*, Wiley, 1986; J. Simons, *J. Phys. Chem.*, **95** (1991), 1017; R. J. Bartlett, J. F. Stanton, *Rev. Comp. Chem.*, **5** (1994), 65; T. Helgaker, P. Jørgensen, J. Olsen, *Molecular Electronic Structure Theory*, Wiley, 2000.
2. C. D. Sherrill, H. F. Schaefer, *Adv. Quant. Chem.*, **34** (1999), 143.
3. J. Olsen, P. Jørgensen, H. Koch, A. Balkova, R. J. Bartlett, *J. Chem. Phys.*, **104** (1996), 8007.
4. J. Olsen, P. Jørgensen, J. Simons, *Chem. Phys. Lett.*, **169** (1990), 463.
5. E. Davidson, *J. Comput. Phys.*, **17** (1975), 87.
6. B. O. Roos, in *Lecture Notes in Quantum Chemistry*, B. O. Roos, Ed., Springer-Verlag, 1992; M. W. Schmidt, M. S. Gordon, *Ann. Rev. Phys. Chem.*, **49** (1998), 233.

7. J. Olsen, B. O. Roos, P. Jørgensen, H. J. Aa. Jensen, *J. Chem. Phys.*, **89** (1988), 2185.
8. W. T. Borden, E. R. Davidson, *Acc. Chem. Res.*, **29** (1996), 67.
9. C. Møller, M. S. Plesset, *Phys. Rev.*, **46** (1934), 618.
10. J. Olsen, O. Christiansen, H. Koch, P. Jørgensen, *J. Chem. Phys.*, **105** (1996), 5082.
11. J. Olsen, P. Jørgensen, T. Helgaker, O. Christiansen, *J. Chem. Phys.*, **112** (2000), 9736.
12. D. Cremer, Z. He, *J. Phys. Chem.*, **100** (1996), 6173.
13. N. C. Handy, P. J. Knowles, K. Somasundram, *Theo. Chim. Acta*, **68** (1985), 87.
14. P. M. Kozłowski, E. R. Davidson, *J. Chem. Phys.*, **100** (1994), 3672.
15. H. B. Schlegel, *J. Phys. Chem.*, **92** (1988), 3075; P. J. Knowles, N. C. Handy, *J. Phys. Chem.*, **92** (1988), 3097.
16. B. O. Roos, K. Andersson, M. P. Fülscher, P.-Å. Malmqvist, L. Serrano-Andres, K. Pierloot, M. Merchán, *Adv. Chem. Phys.*, **93** (1996), 216.
17. R. J. Bartlett, *J. Phys. Chem.*, **93** (1989), 1697.
18. T. V. Voorhis, M. Head-Gordon, *J. Phys. Chem.*, **113** (2000), 8873.
19. J. D. Watts, R. J. Bartlett, *Int. J. Quant. Chem.*, **S27** (1993), 51.
20. J. M. L. Martin, J. P. Francois, R. Gijbels, *Chem. Phys. Lett.*, **172** (1990), 346.
21. G. E. Scuseria, T. J. Lee, *J. Chem. Phys.*, **93** (1990), 5851.
22. K. Raghavachari, J. A. Pople, E. S. Replogle, M. Head-Gordon, *J. Phys. Chem.*, **94** (1990), 5579; R. J. Bartlett, J. D. Watts, S. A. Kucharski, J. Noga, *Chem. Phys. Lett.*, **165** (1990), 513.
23. K. A. Brueckner, *Phys. Rev.*, **96** (1954), 508; J. F. Stanton, J. Gauss, R. J. Bartlett, *J. Chem. Phys.*, **97** (1992), 5554.
24. E. A. Salter, H. Sekino, R. J. Bartlett, *J. Chem. Phys.*, **87** (1987), 502.
25. C. Hampel, K. A. Peterson, H.-J. Werner, *Chem. Phys. Lett.*, **190** (1992), 1.
26. T. J. Lee, R. Kobayashi, N. C. Handy, R. D. Amos, *J. Chem. Phys.*, **96** (1992), 8931.
27. A. Balkova, R. J. Bartlett, *Chem. Phys. Lett.*, **193** (1992), 364.
28. J. A. Pople, M. Head-Gordon, K. Raghavachari, *J. Chem. Phys.*, **87** (1987), 5968.
29. G. E. Scuseria, H. F. Schaefer III, *J. Chem. Phys.*, **90** (1989), 3700.
30. T. J. Lee, A. P. Rendall, P. R. Taylor, *J. Phys. Chem.*, **94** (1990), 5463; for an exception see M. Böhme, G. Frenking, *Chem. Phys. Lett.*, **224** (1994), 195.
31. O. Christiansen, H. Koch, P. Jørgensen, *Chem. Phys. Lett.*, **243** (1995), 409; O. Christiansen, H. Koch, P. Jørgensen, *J. Chem. Phys.*, **103** (1995), 7429.
32. J. F. Stanton, *J. Chem. Phys.*, **101** (1994), 371.
33. J. P. Watts, J. Gauss, R. J. Bartlett, *J. Chem. Phys.*, **98** (1993), 8718.
34. T. J. Lee, P. R. Taylor, *Int. J. Quant. Chem.*, **S23** (1989), 199.
35. P. G. Szalay, R. J. Bartlett, *J. Chem. Phys.*, **101** (1994), 4936.
36. I. M. B. Nielsen, C. Janssen, *Chem. Phys. Lett.*, **310** (1999), 568.
37. Results courtesy of Jeppe Olsen.
38. T. Kato, *Commun. Pure Appl. Math.*, **10** (1957), 151.
39. E. A. Hylleraas, *Z. Phys.*, **65** (1930), 209; V. I. Korobov, *Phys. Rev. A*, **61** (2000), 64503.
40. W. Kolos, L. Wolniewicz, *J. Chem. Phys.*, **49** (1968), 404.
41. W. Kutzelnigg, J. D. Morgan III, *J. Chem. Phys.*, **96** (1992), 4484.
42. W. Kutzelnigg, W. Klopper, *J. Chem. Phys.*, **94** (1991), 1985.
43. W. Klopper, *J. Chem. Phys.*, **102** (1995), 6168.
44. W. Klopper, *J. Chem. Phys.*, **120** (2004), 10890.
45. D. P. Tew, W. Klopper, *J. Chem. Phys.*, **123** (2005), 074101.
46. M. Head-Gordon, J. A. Pople, M. J. Frisch, *Chem. Phys. Lett.*, **153** (1988), 503.
47. H. Koch, A. S. de Meras, T. Helgaker, O. Christiansen, *J. Chem. Phys.*, **104** (1996), 4157.
48. A. K. Wilson, J. Almlöf, *Theo. Chim. Acta*, **95** (1997), 49; M. Schütz, H. J. Werner, *J. Chem. Phys.*, **114** (2001), 661.
49. N. J. Russ, T. D. Crawford, *J. Chem. Phys.*, **121** (2004), 691.
50. F. Weigend, A. Köhn, C. Hättig, *J. Chem. Phys.*, **116** (2002), 3175

51. H. Koch, A. S. de Meras, T. B. Pedersen, *J. Chem. Phys.*, **118** (2003), 1.
52. D. S. Lambrecht, B. Doser, C. Ochsenfeld, *J. Chem. Phys.*, **123** (2005), 184102.
53. W. Li, S. Li, *J. Chem. Phys.*, **121** (2004), 6649.
54. See for example T. Helgaker, J. Gauss, P. Jørgensen, J. Olsen, *J. Chem. Phys.*, **106** (1997), 6430.
55. K. B. Wiberg, A. E. de Oliveira, G. Trucks, *J. Phys. Chem. A*, **106** (2002), 4192.
56. W. M. Foulkes, L. Mitas, R. J. Needs, G. Rajagopal, *Rev. Mod. Phys.*, **73** (2001), 33.
57. D. Bressanini, G. Morosi, S. Tarasco, *J. Chem. Phys.*, **123** (2005), 204109.

5 Basis Sets

Ab initio methods try to derive information by solving the Schrödinger equation without fitting parameters to experimental data. Actually, *ab initio* methods also make use of experimental data, but in a somewhat more subtle fashion. Many different approximate methods exist for solving the Schrödinger equation, and which one to use for a specific problem is usually chosen by comparing the performance against known experimental data. Experimental data thus guides the *selection* of the computational model, rather than directly *entering* into the computational procedure.

One of the approximations inherent in essentially all *ab initio* methods is the introduction of a basis set. Expanding an unknown function, such as a molecular orbital, in a set of known functions is not an approximation if the basis set is *complete*. However, a complete basis set means that an infinite number of functions must be used, which is impossible in actual calculations. An unknown MO can be thought of as a function in the infinite coordinate system spanned by the complete basis set. When a finite basis set is used, only the components of the MO along those coordinate axes corresponding to the selected basis functions can be represented. The smaller the basis set, the poorer the representation. The *type* of basis functions used also influence the accuracy. The better a single basis function is able to reproduce the unknown function, the fewer basis functions are necessary for achieving a given level of accuracy. Knowing that the computational effort of *ab initio* methods scales formally as at least M_{basis}^4 , it is of course of prime importance to make the basis set as small as possible, without compromising the accuracy.¹ The expansion of the molecular orbitals leads to integrals of quantum mechanical operators over basis functions, and the ease with which these integrals can be calculated also depends on the type of basis function. In some cases the *accuracy-per-function* criterion produces a different optimum function type than the *efficiency-per-function* criterion.

5.1 Slater and Gaussian Type Orbitals

There are two types of basis functions (also called *Atomic Orbitals* (AO), although they in general are not solutions to an atomic Schrödinger equation) commonly used

in electronic structure calculations: *Slater Type Orbitals* (STO) and *Gaussian Type Orbitals* (GTO). Slater type orbitals² have the functional form shown in eq. (5.1).

$$\chi_{\zeta,n,l,m}(r, \theta, \varphi) = NY_{l,m}(\theta, \varphi)r^{n-1}e^{-\zeta r} \quad (5.1)$$

Here N is a normalization constant and $Y_{l,m}$ are spherical harmonic functions. The exponential dependence on the distance between the nucleus and electron mirrors the exact orbitals for the hydrogen atom. However, the STOs do not have any radial nodes; nodes in the radial part are introduced by making linear combinations of STOs. The exponential dependence ensures a fairly rapid convergence with increasing numbers of functions, however, as noted in Section 3.5, the calculation of three- and four-centre two-electron integrals cannot be performed analytically. STOs are primarily used for atomic and diatomic systems where high accuracy is required, and in semi-empirical methods where all three- and four-centre integrals are neglected. They can also be used with density functional methods that do not include exact exchange and where the Coulomb energy is calculated by fitting the density into a set of auxiliary functions.

Gaussian type orbitals³ can be written in terms of polar or Cartesian coordinates as shown in eq. (5.2).

$$\begin{aligned} \chi_{\zeta,n,l,m}(r, \theta, \varphi) &= NY_{l,m}(\theta, \varphi)r^{2n-2-l}e^{-\zeta r^2} \\ \chi_{\zeta,l_x,l_y,l_z}(x, y, z) &= Nx^{l_x}y^{l_y}z^{l_z}e^{-\zeta r^2} \end{aligned} \quad (5.2)$$

The sum of l_x , l_y and l_z determines the type of orbital (for example $l_x + l_y + l_z = 1$ is a p-orbital). Although a GTO appears similar in the two set of coordinates, there is a subtle difference. A d -type GTO written in terms of the spherical functions has five components ($Y_{2,2}$, $Y_{2,1}$, $Y_{2,0}$, $Y_{2,-1}$, $Y_{2,-2}$), but there appear to be six components in the Cartesian coordinates (x^2 , y^2 , z^2 , xy , xz , yz). The latter six functions, however, may be transformed to the five spherical d -functions and one additional s -function ($x^2 + y^2 + z^2$). Similarly, there are ten Cartesian “ f -functions” that may be transformed into seven spherical f -functions and one set of spherical p -functions. Modern programs for evaluating two-electron integrals are geared to Cartesian coordinates and they generate pure spherical d -functions by transforming the six Cartesian components to the five spherical functions. When only one d -function is present per atom the saving by removing the extra s -function is small, but if many d -functions and/or higher angular momentum functions (f -, g -, h -, etc., functions) are present, the saving can be substantial. Furthermore, the use of only the spherical components reduces the problems of linear dependence for large basis sets, as discussed below.

The r^2 dependence in the exponential makes the GTOs inferior to the STOs in two respects. At the nucleus a GTO has a zero slope, in contrast to a STO which has a “cusp” (discontinuous derivative), and GTOs consequently have problems representing the proper behaviour near the nucleus. The other problem is that the GTO falls off too rapidly far from the nucleus compared with an STO, and the “tail” of the wave function is consequently represented poorly. Both STOs and GTOs can be chosen to form a complete basis, but the above considerations indicate that more GTOs are necessary for achieving a certain accuracy compared with STOs. A rough guideline says that three times as many GTOs as STOs are required for reaching a given level of accuracy. Figure 5.1 shows how a 1s-STO can be modelled by a linear combination of three GTOs.




Figure 5.1 A 1s-STO modelled by a linear combination of three GTOs (STO-3G)

The increase in the number of GTO basis functions, however, is more than compensated for by the ease of which the required integrals can be calculated. In terms of computational efficiency, GTOs are therefore preferred and are used almost universally as basis functions in electronic structure calculations. Furthermore, essentially all applications take the GTOs to be centred at the nuclei. For certain types of calculations the centre of a basis function may be taken not to coincide with a nucleus, for example being placed at the centre of a bond or between non-bonded atoms for improving the calculation of van der Waals interactions.

5.2 Classification of Basis Sets

Having decided on the type of function (STO/GTO) and the location (nuclei), the most important factor is the number of functions to be used. The smallest number of functions possible is a *minimum basis set*. Only enough functions are employed to contain all the electrons of the neutral atom(s). For hydrogen (and helium) this means a single *s*-function. For the first row in the periodic system it means two *s*-functions (1s and 2s) and one set of *p*-functions (2p_x, 2p_y and 2p_z). Lithium and beryllium formally only require two *s*-functions, but a set of *p*-functions is usually also added. For the second row elements, three *s*-functions (1s, 2s and 3s) and two sets of *p*-functions (2p and 3p) are used.

The next improvement of the basis sets is a doubling of all basis functions, producing a *Double Zeta* (DZ) type basis. The term zeta stems from the fact that the exponent of STO basis functions is often denoted by the Greek letter ζ . A DZ basis thus employs two *s*-functions for hydrogen (1s and 1s'), four *s*-functions (1s, 1s', 2s and 2s') and two sets of *p*-functions (2p and 2p') for first row elements, and six *s*-functions and four sets of *p*-functions for second row elements. The importance of a DZ over a minimum basis can be illustrated by considering the bonding in the HCN molecule. The H—C bond will primarily consist of the hydrogen *s*-orbital and the p_z-orbital on C.

The π -bond between C and N will consist of the p_x (and p_y) orbitals of C and N, and will have a more diffuse electron distribution than the H—C σ -bond. The optimum exponent for the carbon p-orbital will thus be smaller for the x -direction than for the z -direction. If only a single set of p-orbitals is available (minimum basis), a compromise will be necessary. A DZ basis, however, has two sets of p-orbitals with different exponents. The tighter function (larger exponent) can enter the H—C σ -bond with a large coefficient, while the more diffuse function (small exponent) can be used primarily for describing the C—N π -bond. Doubling the number of basis functions thus allows for a much better description of the fact that the electron distribution is different in different directions.




Figure 5.2 A double zeta basis allows for different bonding in different directions

The chemical bonding occurs between valence orbitals. Doubling the 1s-functions in for example carbon allows for a better description of the 1s-electrons. However, the 1s-orbital is essentially independent of the chemical environment, being very close to the atomic case. A variation of the DZ type basis only doubles the number of valence orbitals, producing a *split valence basis*. In actual calculations, a doubling of the core orbitals would rarely be considered, and the term DZ basis is used also for split valence basis sets (or sometimes VDZ, for valence double zeta).

The next step up in basis set size is a *Triple Zeta* (TZ). Such a basis contains three times as many functions as the minimum basis, i.e. six s -functions and three p -functions for the first row elements. Some of the core orbitals may again be saved by only splitting the valence, producing a *triple split valence* basis set. Again the term TZ is used to cover both cases. The names *Quadruple Zeta* (QZ) and *Quintuple* or *Pentuple Zeta* (PZ or 5Z, but not QZ) for the next levels of basis sets are also used, but large basis sets are often given explicitly in terms of the number of basis functions of each type.

So far, only the number of s - and p -functions for each atom (first or second row in the periodic table) has been discussed. In most cases, higher angular momentum functions are also important, and these are denoted *polarization functions*. Consider again the bonding in HCN in Figure 5.2. The H—C bond is primarily described by the hydrogen s -orbital(s) and the carbon s - and p_z -orbitals. It is clear that the electron distribution *along* the bond will be different than *perpendicular* to the bond. If only s -functions are present on hydrogen, this cannot be described. However, if a set of p -orbitals is added to hydrogen, the p_z component can be used for improving the description of the H—C bond. The p -orbital introduces a polarization of the s -orbital(s). Similarly, d -orbitals can be used for polarizing p -orbitals, f -orbitals for polarizing d -orbitals, etc. Once a p -orbital has been added to polarize a hydrogen s -orbital, it may be argued that the p -orbital should now be polarized by adding a d -orbital, which should be polarized by an f -orbital, etc. For independent-particle wave functions, where electron correlation is not considered, the first set of polarization functions (i.e. p -functions for

hydrogen and *d*-functions for heavy atoms) is by far the most important, and will in general describe most of the important charge polarization effects.

If methods including electron correlation are used, higher angular momentum functions are essential. Electron correlation describes the energy lowering by the electrons “avoiding” each other, beyond the average effect taken into account by Hartree–Fock methods. Two types of correlation can be identified, an “*in-out*” and an “*angular*” correlation. The *in-out* or *radial correlation* refers to the situation where one electron is close to, and the other far from, the nucleus. To describe this, the basis set needs functions of the same type, but with different exponents. The *angular correlation* refers to the situation where two electrons are on opposite sides of the nucleus. To describe this, the basis set needs functions with the same magnitude exponents, but different angular momentum. For example, to describe angular correlation of an *s*-function, *p*-functions (and *d*-, *f*-, *g*-functions, etc.) are needed. The angular correlation is of similar importance as the radial correlation, and higher angular momentum functions are consequently essential for correlated calculations. Although these should properly be labelled correlation functions, they also serve as polarization functions for HF wave functions, and it is common to denote them as polarization functions.

Normally only the correlation of the valence electrons is considered, and the exponents of the polarization functions should be of the same magnitude as the valence *s*- and *p*-functions (actually slightly larger in order to have the same maximum in the radial distribution function). In contrast to HF methods, the higher angular momentum functions (beyond the first set of polarization functions) are quite important. Or alternatively formulated, the convergence in terms of angular momentum is slower for correlated wave functions than at the HF level. For a basis set that is complete up to angular momentum *L*, numerical analysis suggests the asymptotic convergence at the HF level is exponential (i.e. $\sim \exp(-\sqrt{L})$), while it is $\sim L^{-3}$ at correlated levels.⁴

Polarization functions are added to the chosen *sp*-basis. Adding a single set of polarization functions (*p*-functions on hydrogens and *d*-functions on heavy atoms) to the DZ basis forms a *Double Zeta plus Polarization* (DZP) type basis. There is a variation where polarization functions are only added to non-hydrogen atoms. This does not mean that polarization functions are not important on hydrogen. However, hydrogen often has a “passive” role, sitting at the end of bonds that do not take active part in the property of interest. The error introduced by not including hydrogen polarization functions is often rather constant and, as the interest is usually in energy differences, tends to cancel out. As hydrogen often accounts for a large number of atoms in the system, a saving of three basis functions for each hydrogen is significant. If hydrogen plays an important role in the property of interest, it is of course not a good idea to neglect polarization functions on hydrogen.

Similarly to the *sp*-basis sets, multiple sets of polarization functions with different exponents may be added. If two sets of polarization functions are added to a TZ *sp*-basis, a *Triple Zeta plus Double Polarization* (TZ2P) type basis is obtained. For larger basis sets with many polarization functions the explicit composition in terms of number and types of functions is usually given. At the HF level there is usually little gained by expanding the basis set beyond TZ2P, and even a DZP type basis set usually gives “good” results (compared with the HF limit). Correlated methods, however, require more, and higher angular momentum, polarization functions to achieve the same level of convergence.

Before moving on we need to introduce the concept of *basis set balance*. In principle, many sets of polarization functions may be added to a small *sp*-basis, but this is a poor idea. If an insufficient number of *sp*-functions has been chosen for describing the fundamental electron distribution, the optimization procedure used in obtaining the wave function (and possibly also the geometry) may try to compensate for inadequacies in the *sp*-basis by using higher angular momentum functions, thereby producing artefacts. A rule of thumb says that the number of functions of a given type should at most be one less than the type with one lower angular momentum. A 3s2p1d basis is balanced, but a 3s2p2d2f1g is too heavily polarized. It may not be necessary to polarize the basis all the way up, thus a 5s4p3d2f1g basis is balanced, but if it is known (for example by comparison with experimental data) that *f*- and *g*-functions are unimportant, they may be left out. Furthermore, it may be that two *d*-functions are sufficient for the given purpose, although a 5s4p1d basis would be considered underpolarized.

Another aspect of basis set balance is the occasional use of *mixed* basis sets, for example a DZP quality on the atoms in the “interesting” part of the molecule and a minimum basis for the “spectator” atoms. Another example would be the addition of polarization functions for only a few hydrogens that are located “near” the reactive part of the system. For a large molecule, this may lead to a substantial saving in the number of basis functions. It should be noted that this may bias the results and can create artefacts. For example, a calculation on the H₂ molecule with a minimum basis at one end and a DZ basis at the other end will predict that H₂ has a dipole moment, since the variational principle will preferentially place the electrons near the centre with the most basis functions. The majority of calculations are therefore performed with basis sets of the same quality (minimum, DZP, TZ2P, . . .) on all atoms, possibly removing polarization and/or diffuse (small exponent) functions on hydrogen. Even so, it may be argued that small basis sets inherently tend to be unbalanced. Consider for example the LiF molecule in a minimum or DZ type basis. This will have a very ionic structure, Li⁺F⁻, with nearly all the valence electrons being located at the fluorine. In terms of number of basis functions per electron, the Li basis is thus of a much higher quality than the F basis, and thereby unbalanced. This effect of course diminishes as the size of the atomic basis set increases.

Except for very small systems, it is impractical to saturate the basis set such that the absolute error in the energy is reduced below chemical accuracy, say 4 kJ/mol. The important point in choosing a balanced basis set is to keep the error as constant as possible. The use of mixed basis sets should therefore only be done after careful consideration. Furthermore, the use of small basis sets for systems containing elements with substantially different numbers of valence electrons (such as LiF) may produce artefacts.

Having decided on the number of basis functions (from a consideration of the property of interest and the computational cost), the question becomes: how are the values for the exponents in the basis functions chosen? The values for the *s*- and *p*-functions are typically determined by performing variational HF calculations for the atoms, using the exponents as variational parameters. The exponent values that give the lowest energy are the “best”, at least for the atom. In some cases, the optimum exponents are chosen based on minimizing the energy of a wave function that includes electron correlation. The HF procedure cannot be used for determining exponents of polarization

functions for atoms. By definition these functions are unoccupied in atoms, and therefore make no contribution to the energy. Suitable polarization exponents may be chosen by performing variational calculations on molecular systems (where the HF energy does depend on polarization functions) or on atoms with correlated wave functions. Since the main function of higher angular momentum functions is to recover electron correlation, the latter approach is usually preferred. Often only the optimum exponent is determined for a single polarization function, and multiple polarization functions are generated by splitting the exponents symmetrically around the optimum value for a single function. The splitting factor is typically taken in the range 2–4. For example if a single *d*-function for carbon has an exponent value of 0.8, two polarization functions may be assigned with exponents of 0.4 and 1.6 (splitting factor of 4). The details of how the exponents are determined for various basis sets are discussed in the following sections.

5.3 Even- and Well-Tempered Basis Sets

The optimization of basis function exponents is an example of a highly non-linear optimization problem (Chapter 12). When the basis set becomes large, the optimization problem is no longer easy. The basis functions start to become linearly dependent (the basis set approaches completeness) and the energy becomes a very flat function of the exponents. Analyses of basis sets that have been optimized by variational methods reveal that the ratio between two successive exponents is approximately constant. Taking this ratio to be constant reduces the optimization problem to only two parameters for each type of basis function, *independent of the size of the basis*. Such basis sets have been labelled *even-tempered* basis sets, with the *i*th exponent given as $\zeta_i = \alpha\beta^i$, where α and β are fixed constants for a given type of function and nuclear charge. It was later discovered that the *optimum* α and β constants to a good approximation can be written as functions of the size of the basis set, M .⁵

$$\begin{aligned}\zeta_i &= \alpha\beta^i; \quad i = 1, 2, \dots, M \\ \ln(\ln \beta) &= b \ln M + b' \\ \ln \alpha &= a \ln(\beta - 1) + a'\end{aligned}\tag{5.3}$$

The constants a, a', b and b' depend only on the atom type and the type of function (*s* or *p*). Even-tempered basis sets have the advantage that it is easy to generate a sequence of basis sets that are guaranteed to converge towards a complete basis. This is useful if the attempt is to extrapolate a given property to the basis set limit. The disadvantage is that the convergence is somewhat slow, and an explicitly optimized basis set of a given size will usually give a better answer than an even-tempered basis of the same size.

Even-tempered basis sets have the same ratio between exponents over the whole range. From chemical considerations it is usually preferable to cover the valence region better than the core region. This may be achieved by *well-tempered* basis sets.⁶ The idea is similar to the even-tempered basis sets, with the exponents being generated by a suitable formula containing only a few parameters to be optimized. The exponents in a well-tempered basis of size M are generated according to eq. (5.4).

$$\zeta_i = \alpha \beta^{i-1} \left(1 + \gamma \left(\frac{i}{M} \right)^\delta \right); \quad i = 1, 2, \dots, M \quad (5.4)$$

The α , β , γ and δ parameters are optimized for each atom. The exponents are the same for all types of angular momentum functions, and s -, p - and d -functions (and higher angular momentum) consequently have the same radial part.

A well-tempered basis set has four parameters, compared with two for an even-tempered one, and is consequently capable of giving a better result for the same number of functions. Petersson *et al.*⁷ have proposed a somewhat more general parameterization based on expanding the logarithmic exponents in a polynomial of order K in the basis function number.

$$\ln \zeta_i = \sum_{k=0}^K a_k i^k; \quad i = 1, 2, \dots, M \quad (5.5)$$

Setting $K = 1$ is equivalent to generating an even-tempered basis set. The optimization of the parameters a_k becomes problematic for K larger than 2, since the polynomials are non-orthogonal, and increasing K thus significantly changes all the expansion coefficients. This problem can be alleviated by using *Legendre* polynomials instead, since these are orthogonal, and this significantly improves the optimization.

$$\begin{aligned} \ln \zeta_i &= \sum_{k=0}^K a_k P_k(i); \quad i = 1, 2, \dots, M \\ P_0(i) &= 1 \\ P_1(i) &= i \\ P_2(i) &= i^2 - 1 \\ P_3(i) &= i^3 - i \\ &\dots \end{aligned} \quad (5.6)$$

It has been found that a fourth-order polynomial ($K = 3$) expansion produces much better results than the well-tempered formula, despite having the same number of variables. Furthermore, the results from a fourth-order Legendre parameterization with M basis functions is comparable to those from a fully optimized basis set with $M - 1$ functions, i.e. the penalty in reducing the number of optimization variables from M to four is only one function. The Legendre parameterization furthermore solves the potential problem of variational collapse, i.e. two neighbouring exponents collapsing to the same value during optimization, and eq. (5.6) thus provides an efficient way of systematically approaching the basis set limit.

Optimization of basis sets is not something the common user needs to worry about. Optimized basis sets of many different sizes and qualities are available either in the forms of tables, websites⁸ or stored internally in the computer programs. The user “merely” has to select a suitable basis set. However, if the interest is in specialized properties the basis set may need to be tailored to meet the specific needs. For example if the property of interest is an accurate value for the electron density *at* the nucleus (for example for determining the Fermi contact contribution to spin–spin coupling (see Section 10.7.6)) then basis functions with very large exponents are required.

Alternatively, for calculating hyperpolarizabilities, very diffuse functions are required. In such cases, the basis function optimization is in terms of the property of interest, and not in terms of energy, i.e. basis functions are added until the change upon addition of one extra function is less than a given threshold.

5.4 Contracted Basis Sets

One disadvantage of all energy-optimized basis sets is the fact that they primarily depend on the wave function in the region of the inner-shell electrons. The 1s-electrons account for a large part of the total energy, and minimizing the energy will tend to make the basis set optimum for the core electrons, and less so for the valence electrons. However, chemistry is mainly dependent on the valence electrons. Furthermore, many properties (for example polarizability) depend mainly on the wave function “tail” (far from the nucleus), which energetically is unimportant. An energy-optimized basis set that gives a good description of the outer part of the wave function therefore needs to be very large, with the majority of the functions being used to describe the 1s-electrons with an accuracy comparable with the outer electrons in an energetic sense. This is not the most efficient way of designing basis sets for describing the outer part of the wave function. Instead energy-optimized basis sets are usually augmented explicitly with *diffuse* functions (basis functions with small exponents). Diffuse functions are needed whenever loosely bound electrons are present (for example anions or excited states) or when the property of interest is dependent on the wave function tail (for example polarizability).

The fact that many basis functions focus on describing the energetically important, but chemically unimportant, core electrons is the foundation for *contracted* basis sets. Consider for example a basis set consisting of ten *s*-functions (and some *p*-functions) for carbon. Having optimized these ten exponents by a variational calculation on a carbon atom, maybe six of the ten functions are found primarily to be used for describing the 1s-orbital, and two of the four remaining describe the “inner” part of the 2s-orbital. The important chemical region is the outer valence. Out of the ten functions, only two are actually used for describing the chemically interesting phenomena. Considering that the computational cost increases as the fourth power (or higher) of the number of basis functions, this is inefficient. As the core orbitals change very little depending on the chemical bonding situation, the MO expansion coefficients in front of these inner basis functions also change very little. The majority of the computational effort is therefore spent describing the chemically uninteresting part of the wave function, which is furthermore almost constant.

Consider now making the variational coefficients in front of the inner basis functions constant, i.e. they are no longer parameters to be determined by the variational principle. The 1s-orbital is thus described by a *fixed* linear combination of say six basis functions. Similarly, the remaining four basis functions may be contracted into only two functions, for example by fixing the coefficient in front of the inner three functions. In doing this the number of basis functions to be handled by the variational procedure has been reduced from ten to three.

Combining the full set of basis functions, known as the *primitive* GTOs (PGTOs), into a smaller set of functions by forming fixed linear combinations is known as

basis set *contraction*, and the resulting functions are called *contracted* GTOs (CGTOs).

$$\chi(\text{CGTO}) = \sum_i^k a_i \chi_i(\text{PGTO}) \quad (5.7)$$

The previously introduced acronyms DZP, TZ2P, etc., refer to the number of contracted basis functions. Contraction is especially useful for orbitals describing the inner (core) electrons, since they require a relatively large number of functions for representing the wave function cusp near the nucleus, and furthermore are largely independent of the environment. Contracting a basis set will always increase the energy, since it is a restriction of the number of variational parameters, and makes the basis set less flexible, but it will also reduce the computational cost significantly. The decision is thus how much loss in accuracy is acceptable compared with the gain in computational efficiency.

The *degree of contraction* is the number of PGTOs entering the CGTO, typically varying between one and ten. The specification of a basis set in terms of primitive and contracted functions is done by the notation (10s4p1d/4s1p) \rightarrow [3s2p1d/2s1p]. The basis in parenthesis is the number of primitives with heavy atoms (first row elements) before the slash and hydrogen after. The basis in the square brackets is the number of contracted functions. Note that this does not indicate how the contraction is done, it only indicates the size of the final basis (and thereby the size of the variational problem in HF calculations).

There are two different ways of contracting a set of primitive GTOs to a set of contracted GTOs: *segmented* and *general* contraction. Segmented contraction is the older method, and the one used in the above example. A given set of PGTOs is partitioned into smaller sets of functions that are made into CGTOs by determining suitable coefficients. A 10s basis set may be contracted to 3s by taking the inner six functions as one CGTO, the next three as the second CGTO and the one remaining PGTO as the third “contracted” GTO.

$$\begin{aligned} \chi_1(\text{CGTO}) &= \sum_{i=1}^6 a_i \chi_i(\text{PGTO}) \\ \chi_{21}(\text{CGTO}) &= \sum_{i=7}^9 a_i \chi_i(\text{PGTO}) \\ \chi_3(\text{CGTO}) &= \chi_{10}(\text{PGTO}) \end{aligned} \quad (5.8)$$

In a segmented contraction each primitive as a rule is used only in *one* contracted function, i.e. the primitive set of functions is partitioned into disjoint sets. In some cases it may be necessary to duplicate one or two PGTOs in two adjacent CGTOs. The contraction coefficients can be determined by a variational optimization of the atomic HF energy, where both the exponents and contraction coefficients are optimized simultaneously. It should be noted that this optimization often produces multiple minima, and selecting a suitable “optimum” solution may be non-trivial.⁹

In a general contraction *all* primitives (on a given atom) enter *all* the contracted functions, but with different contraction coefficients.

$$\begin{aligned}
 \chi_1(\text{CGTO}) &= \sum_{i=1}^{10} a_i \chi_i(\text{PGTO}) \\
 \chi_{21}(\text{CGTO}) &= \sum_{i=1}^{10} b_i \chi_i(\text{PGTO}) \\
 \chi_3(\text{CGTO}) &= \sum_{i=1}^{10} c_i \chi_i(\text{PGTO})
 \end{aligned} \tag{5.9}$$

One popular way of obtaining general contraction coefficients is from *Atomic Natural Orbitals* (ANOs), to be discussed in Section 5.4.5. The difference between segmented and general contraction may be illustrated as shown in Figure 5.3.

	Segmented contraction			General contraction		
	CGTO-1	CGTO-2	CGTO-3	CGTO-1	CGTO-2	CGTO-3
PGTO-1						
PGTO-2						
PGTO-3						
PGTO-4						
PGTO-5						
PGTO-6						
PGTO-7						
PGTO-8						
PGTO-9						
PGTO-10						

Figure 5.3 Illustrating segmented and general contraction

In reality, there are very few truly segmented or general contracted basis sets. General contracted basis sets normally leave the outermost function(s) uncontracted, and a Gram–Schmidt type orthogonalization can be used for partly segmenting the inner functions.¹⁰ The disjoint nature of the primitive set of functions in a segmented contraction, on the other hand, often necessitates a duplication of one or more functions, i.e. effectively a general contraction. The segmented–general classification should thus be seen as limiting cases, with actual basis sets having varying characteristics of both types.

There are many different contracted basis sets available in the literature or built into programs, and the average user usually only needs to select a suitable quality basis for the calculation. Below is a short description of some basis sets that often are used in routine calculations. The contractions are given for a first row element (such as carbon), while the corresponding ones for other elements can be found in the references.

5.4.1 Pople style basis sets

STO-*n*G basis sets These are *Slater type orbitals* consisting of *n* PGTOs.¹¹ This is a minimum type basis where the exponents of the PGTO are determined by *fitting* to the STO, rather than optimizing them by a variational procedure. Although basis sets with *n* = 2–6 have been derived, it has been found that using more than three PGTOs for representing the STO gives little improvement, and the STO-3G basis is a widely used minimum basis. This type of basis set has been determined for many elements of the periodic table. The designation of the carbon STO-3G basis is (6s3p) → [2s1p].

k-nlmG basis sets These basis sets, designed by Pople and coworkers, are of the split valence type, with the *k* in front of the dash indicating how many PGTOs are used for representing the core orbitals. The *nlm* after the dash indicate both how many functions the valence orbitals are split into, and how many PGTOs are used for their representation. Two values (*nl*) indicate a split valence, while three values (*nlm*) indicate a triple split valence. The values *before* the G (for Gaussian) indicate the *s*- and *p*-functions in the basis; the polarization functions are placed *after* the G. These types of basis sets have the further restriction that the same exponent is used for both the *s*- and *p*-functions in the valence. This increases the computational efficiency, but of course decreases the flexibility of the basis set. The exponents and contraction coefficients have been optimized by variational procedures at the HF level for atoms.

3-21G This is a split valence basis, where the core orbitals are a contraction of three PGTOs, the inner part of the valence orbitals is a contraction of two PGTOs and the outer part of the valence is represented by one PGTO.¹² The designation of the carbon 3-21G basis is $(6s3p) \rightarrow [3s2p]$. Note that the 3-21G basis contains the same number of primitive GTOs as the STO-3G, however, it is much more flexible as there are twice as many valence functions that can combine freely to make MOs.

6-31G This is also a split valence basis, where the core orbitals are a contraction of six PGTOs, the inner part of the valence orbitals is a contraction of three PGTOs and the outer part of the valence is represented by one PGTO.¹³ The designation of the carbon 6-31G basis is $(10s4p) \rightarrow [3s2p]$. In terms of contracted basis functions it contains the same number as 3-21G, but the representation of each function is better since more PGTOs are used.

6-311G This is a triple split valence basis, where the core orbitals are a contraction of six PGTOs and the valence split into three functions, represented by three, one and one PGTOs, respectively, i.e. $(11s5p) \rightarrow [4s3p]$.¹⁴

To each of these basis sets can be added diffuse¹⁵ and/or polarization functions.¹⁶ Diffuse functions are normally *s*- and *p*-functions and consequently go before the G. They are denoted by + or ++, with the first + indicating one set of diffuse *s*- and *p*-functions on heavy atoms, and the second + indicating that a diffuse *s*-function is added also to hydrogen. The argument for only adding diffuse functions on non-hydrogen atoms is the same as for only adding polarization functions on non-hydrogens (Section 5.2). Polarization functions are indicated after the G, with a separate designation for heavy atoms and hydrogen. The 6-31+G(d) is a split valence basis with one set of diffuse *sp*-functions on heavy atoms only and a single *d*-type polarization function on heavy atoms. A 6-311++G(2df,2pd) is similarly a triple split valence with additional diffuse *sp*-functions, two *d*-functions and one *f*-function on heavy atoms, and diffuse *s*- and two *p*- and one *d*-functions on hydrogen. The largest standard Pople style basis set is 6-311++G(3df,3pd). These types of basis set have been derived for hydrogen and the first row elements, and some of the basis sets have also been derived for second and higher row elements. The composition in terms of contracted and primitive functions is given in Table 5.1.

If only one set of polarization functions is used, an alternative notation in terms of * is also widely used. The 6-31G* basis is identical to 6-31G(d), and 6-31G** is identical to 6-31G(d,p). A special note should be made for the 3-21G* basis. The 3-21G

Table 5.1 Composition in terms of contracted and primitive basis functions for some Pople style basis sets

Basis	Hydrogen		First row elements		Second row elements	
	Contracted	Primitive	Contracted	Primitive	Contracted	Primitive
STO-3G	1s	3s	2s1p	6s3p	3s2p	9s6p
3-21G	2s	3s	3s2p	6s3p	4s3p	9s6p
6-31G(d,p)	2s1p	4s	3s2p1d	10s4p	4s3p1d	16s10p
6-311G(2df,2pd)	3s2p1d	5s	4s3p2d1f	11s5p	6s4p2d1f ^a	13s9p ^a

^a McLean–Chandler basis set

basis is basically too small to support polarization functions (it become unbalanced). However, the 3-21G basis by itself performs poorly for hypervalent molecules, such as sulfoxides and sulfones. This can be improved substantially by adding a set of *d*-functions. The 3-21G* basis has only *d*-functions on second row elements (it is sometimes denoted 3-21G(*) to indicate this), and should not be considered a polarized basis. Rather the addition of a set of *d*-functions is an *ad hoc* repair of a known flaw.

5.4.2 Dunning–Huzinaga basis sets

Huzinaga has determined uncontracted energy-optimized basis sets up to (10s6p) for first row elements.¹⁷ This was latter extended to (14s9p) by van Duijneveldt,¹⁸ and up to (18s13p) by Partridge.¹⁹ Dunning has used the Huzinaga primitive GTOs to derive various contraction schemes, and these are known as Dunning–Huzinaga (DH) type basis sets.²⁰ A DZ type basis can be made by a contraction of the (9s5p) PGTO to [4s2p]. The contraction scheme is 6,1,1,1 for *s*-functions and 4,1 for the *p*-functions. A widely used split valence type basis is a contraction of the same primitive set to [3s2p] where the *s*-contraction is 7,2,1 (note that one primitive enters twice). A widely used TZ type basis (actually only a triple split valence) is a contraction of the (10s6p) to [5s3p], with the contraction scheme 6,2,1,1,1 for *s*-functions and 4,1,1 for *p*-functions. Again, a duplication of one of the *s*- and *p*-primitives has been allowed.

McLean and Chandler have developed a similar set of contracted basis sets from Huzinaga primitive optimized sets for second row elements.²¹ A DZ type basis is derived by contracting (12s8p) → [5s3p], and a TZ type is derived by contracting (12s9p) → [6s5p]. The latter contraction is 6,3,1,1,1,1 for the *s*-functions (note a duplication of one function) and 4,2,1,1,1 for the *p*-functions, and is often used in connection with the Pople 6-311G when second row elements are present.

The Dunning–Huzinaga type basis sets do not have the restriction of the Pople style basis sets of equal exponents for the *s*- and *p*-functions, and they are therefore somewhat more flexible, but computationally also more expensive. The major determining factor, however, is the *number* of basis functions and less so the exact *description* of each function. Normally there is little difference in the performance between different DZ or different TZ type basis sets.

The primary reason for the popularity of the Pople and DH style basis sets is the extensive calibration available. There have been so many calculations reported with these basis sets that it is possible to get a fairly good idea of the level of accuracy that

can be attained with a given basis. This is of course a self-sustaining procedure, the more calculations that are reported with a given basis, the more popular it becomes, since the calibration set becomes larger and larger.

5.4.3 MINI, MIDI and MAXI basis sets

Tatewaki and Huzinaga have optimized minimum basis sets for a large part of the periodic table at the HF level.²² The MINI- n ($n = 1-4$) basis sets are all minimum basis sets with three PGTOS in the 2s CGTO, and a varying number of PGTOS in the 1s and 2p CGTOs. In terms of PGTOS, the MINI-1 is (3s,3s,3p), the MINI-2 is (3s,3s,4p), the MINI-3 is (4s,3s,3p) and the MINI-4 is (4s,3s,4p). These MINI basis sets in general perform better than STO-3G, but it should be kept in mind that they are still minimum basis sets. The MIDI- n basis sets are identical to MINI- n , except that the outer valence function is decontracted. The MAXI- n basis sets all employ four PGTOS for the 2s CGTO and from five to seven PGTOS for the 1s and 2p CGTOs. The valence orbitals are split into three or four functions, and MAXI-1 is (9s5p) \rightarrow [4s3p] (contraction 5,2,1,1 and 3,1,1), MAXI-3 is (10s6p) \rightarrow [5s4p] (contraction 6,2,1,1,1 and 3,1,1,1) and MAXI-5 is (11s7p) \rightarrow [5s4p] (contraction 7,2,1,1,1 and 4,1,1,1).

5.4.4 Ahlrichs type basis sets

The group centred around R. Ahlrichs has designed basis sets of DZ, TZ and QZ quality for the elements up to Kr. The *Split Valence Polarized* (SVP) basis set is a [3s2p] contraction of a (7s4p) set of primitive functions (contraction 5,1,1 and 3,1), while the *Triple Zeta Valence* (TZV) basis set is a [5s3p] contraction of an (11s6p) set of primitive functions (contraction 6,2,1,1,1 and 4,1,1).²³ More recently, the series has been extended by a *Quadruple Zeta Valence* (QZV) basis set, being a [7s4p] contraction of a (15s8p) set of primitive functions with the contraction 8,2,1,1,1,1,1 and 5,1,1,1.²⁴ Note that both the TZV and QZV basis sets employ more contracted s -functions than indicated by the TZ and QZ acronyms. The s - and p -exponents and corresponding contraction coefficients are optimized at the HF level, while the polarization functions are taken from the cc-pVxZ basis sets.

Table 5.2 Composition in terms of contracted and primitive basis functions for the Ahlrichs type basis sets

Basis	Hydrogen		First row elements		Second row elements	
	Contracted	Primitive	Contracted	Primitive	Contracted	Primitive
SVP	2s1p	4s	3s2p1d	7s4p	4s3p1d	10s7p
TZV	3s2p1d	5s	5s3p2d1f	11s6p	5s4p2d1f	14s9p
QZV	4s3p2d1f	7s	7s4p3d2f1g	15s8p	9s6p4d2f1g	20s14p

5.4.5 Atomic natural orbital basis sets

All of the above basis sets are of the segmented contraction type. Modern contracted basis sets aimed at producing very accurate wave functions often employ a general

contraction scheme. The *Atomic Natural Orbitals* (ANO) and correlation consistent basis sets below are of the general contraction type.

The idea in the ANO type basis sets is to contract a large PGTO set to a fairly small number of CGTOs by using *natural orbitals* from a correlated calculation on the free atom, typically at the CISD level.²⁵ The natural orbitals are those that diagonalize the density matrix, and the eigenvalues are called *orbital occupation numbers* (see Section 9.5). The orbital occupation number is the number of electrons in the orbital. For an RHF wave function, ANOs would be identical to the canonical orbitals with occupation numbers of exactly 0 or 2. When a correlated wave function is used, however, the occupation number may have any value between 0 and 2. The ANO contraction selects the important combinations of the PGTOs from the magnitude of the occupation numbers. A large primitive basis, typically generated as an even-tempered sequence, may generate several different contracted basis sets by gradually lowering the selection threshold for the occupation number. The nice feature of the ANO contraction is that it more or less “automatically” generates balanced basis sets, e.g. for neon the ANO procedure generates the following basis set: [2s1p], [3s2p1d], [4s3p2d1f] and [5s4p3d2f1g]. Furthermore, in such a sequence the smaller ANO basis sets are true subsets of the larger, since the same set of primitive functions is used.

5.4.6 Correlation consistent basis sets

The primary disadvantage of ANO basis sets is that a very large number of primitive GTOs are necessary for converging towards the basis set limit. Dunning and coworkers have proposed a somewhat smaller set of primitives that yields comparable results to the ANO basis sets.²⁶ The *correlation consistent* (cc; the convention is to use lower case letters as the acronym, to distinguish it from coupled cluster (CC)) basis sets are geared towards recovering the correlation energy of the valence electrons. The name correlation consistent refers to the fact that the basis sets are designed such that functions that contribute similar amounts of correlation energy are included at the same stage, independent of the function type. For example, the first *d*-function provides a large energy lowering, but the contribution from a second *d*-function is similar to that from the first *f*-function. The energy lowering from a third *d*-function is similar to that from the second *f*-function and the first *g*-function. The addition of polarization functions should therefore be done in the order: 1d, 2d1f and 3d2f1g. An additional feature of the cc basis sets is that the energy error from the *sp*-basis should be comparable with (or at least not exceed) the correlation error arising from the incomplete polarization space, and the *sp*-basis therefore also increases as the polarization space is extended. The *s*- and *p*-basis set exponents are optimized at the HF level for the atoms, while the polarization exponents are optimized at the CISD level, and the primitive functions are contracted by a general contraction scheme using natural orbital coefficients.

Several different sizes of cc basis sets are available in terms of final number of contracted functions. These are known by their acronyms: cc-pVDZ, cc-pVTZ, cc-pVQZ, cc-pV5Z and cc-pV6Z (*correlation consistent polarized Valence Double/Triple/Quadruple/Quintuple/Sextuple Zeta*). The composition in terms of contracted and primitive (for the *s*- and *p*-part) functions is shown in Table 5.3. Note that each step up in terms of quality increases each type of basis function by one, and adds a new type of higher

Table 5.3 Composition in terms of contracted and primitive basis functions for the correlation consistent basis sets

Basis	Hydrogen		First row elements		Second row elements	
	Contracted	Primitive	Contracted	Primitive	Contracted	Primitive
cc-pVDZ	2s1p	4s	3s2p1d	9s4p	4s3p2d	12s8p
cc-pVTZ	3s2p1d	5s	4s3p2d1f	10s5p	5s4p3d1f	15s9p
cc-pVQZ	4s3p2d1f	6s	5s4p3d2f1g	12s6p	6s5p4d2f1g	16s11p
cc-pV5Z	5s4p3d2f1g	8s	6s5p4d3f2g1h	14s8p	7s6p5d3f2g1h	20s12p
cc-pV6Z	6s5p4d3f2g1h	10s	7s6p5d4f3g2h1i	16s10p	8s7p6d4f3g2h1i	21s14p

order polarization function. For second row systems it has been found that the performance is significantly improved by adding an extra tight *d*-function.²⁷

The energy-optimized cc-basis sets can be augmented with diffuse functions, indicated by adding the prefix *aug-* to the acronym.²⁸ The augmentation consists of adding one extra function with a smaller exponent for each angular momentum, i.e. the *aug-cc-pVDZ* has additionally one *s*-, one *p*- and one *d*-function, the *cc-pVTZ* has 1s1p1d1f extra for non-hydrogens and so on. The cc-basis sets may also be augmented with additional *tight* functions (large exponents) if the interest is in recovering core–core and core–valence electron correlation, producing the acronyms *cc-pCVXZ* (*X* = D, T, Q, 5). The *cc-pCVDZ* has additionally one tight *s*- and one *p*-function, the *cc-pCVTZ* has 2s2p1d tight functions, the *cc-pCVQZ* has 3s3p2d1f and the *cc-pCV5Z* has 4s4p3d2f1g for non-hydrogens.²⁹

5.4.7 Polarization consistent basis sets

The basis set convergence of electron correlation methods is inverse polynomial in the highest angular momentum functions included in the basis set, while the convergence of the independent-particle HF and DFT methods is exponential.³⁰ This difference in convergence properties suggests that the optimum basis sets for the two cases will also be different, especially should low angular momentum functions be more important for HF/DFT methods than for electron correlation methods as the basis set becomes large. Since DFT methods (Chapter 6) are rapidly becoming the preferred method for routine calculations, it is of interest to have basis sets that are optimized for DFT type calculations, and that are capable of systematically approaching the basis set limit. The *polarization consistent* (pc) basis sets are developed analogously to the correlation consistent basis sets except that they are optimized for DFT methods.³¹ The name indicates that they are geared towards describing the polarization of the (atomic) electron density upon formation of a molecule, rather than describing the correlation energy. Since there is little difference between HF and DFT, and even less difference between different DFT functionals, these basis sets are suitable for independent-particle methods in general.

The polarization consistent basis sets again employ an energetic criterion for determining the importance of each type of basis function. The level of polarization beyond the isolated atom is indicated by a value after the acronym, i.e. a pc-0 basis set is

Table 5.4 Composition in terms of contracted and primitive basis functions for the polarization consistent basis sets

Basis	Hydrogen		First row elements		Second row elements	
	Contracted	Primitive	Contracted	Primitive	Contracted	Primitive
pc-0	2s	3s	3s2p	5s3p	4s3p	8s6p
pc-1	2s1p	4s	3s2p1d	7s4p	4s3p1d	11s8p
pc-2	3s2p1d	6s	4s3p2d1f	10s6p	5s4p2d1f	13s10p
pc-3	5s4p2d1f	9s	6s5p4d2f1g	14s9p	6s5p4d2f1g	17s13p
pc-4	7s6p3d2f1g	11s	8s7p6d3f2g1h	18s11p	7s6p6d3f2g1h	20s16p

unpolarized, pc-1 contains a single polarization function with one higher angular momentum, pc-2 contains polarization functions up to two beyond that required for the atom, etc. In contrast to the cc-pVxZ basis sets, the importance of the polarization functions must be determined at the molecular level, since the atomic energies only depend on *s*- and *p*-functions (at least for elements in the first two rows in the periodic table). For the DZ and TZ type basis sets (pc-1 and pc-2), the consistent polarization is the same as for the cc-pVxZ basis sets (1d and 2d1f), but at the QZ and 5Z levels (pc-3 and pc-4) there are one and two additional *d*-functions (4d2f1g and 6d3f2g1h), respectively. The *s*- and *p*-basis set exponents are optimized at the DFT level for the atoms, while the polarization exponents are selected as suitable average values from optimizations for a selection of molecules. The primitive functions are subsequently contracted by a general contraction scheme by using the atomic orbital coefficients.

For properties dependent on the wave function tail, such as electric moments and polarizabilities, the convergence towards the basis set limit can be improved by explicitly adding a set of diffuse functions, producing the acronym aug-pc-*n*.

5.4.8 Basis set extrapolation

The main advantage of the ANO, correlation consistent and polarization consistent basis sets is the ability to generate a sequence of basis sets that converges toward the basis set limit in a systematic fashion. For example, from a series of calculations with the 3-21G, 6-31G(d,p), 6-311G(2d,2p) and 6-311++G(3df,3pd) basis sets it may not be obvious whether the property of interest is “converged” with respect to further increases in the basis, and it is difficult to estimate what the basis set limit would be. This is partly due to the fact that different primitive GTOs are used in each of these segmented basis sets, and partly due to the lack of higher angular momentum functions. From the same (large) set of primitive GTOs, however, increasingly large ANO basis sets may be generated by a general contraction scheme that allows an estimate of the basis set limiting value. Similarly, the cc-pVxZ basis sets consistently reduce errors (both HF and correlation) for each step up in quality. In test cases it has been found that the cc-pVDZ basis can provide ~65% of the total (valence) correlation energy, the cc-pVTZ ~85%, cc-pVQZ ~93%, cc-pV5Z ~96% and cc-pV6Z ~98%, with similar reductions of the HF error.

Given the systematic nature of the cc basis sets, several different schemes have been proposed for extrapolation to the infinite basis set limit, using the highest angular momentum L_{\max} included in the basis set as the extrapolating parameter.³² At the HF and DFT levels the convergence is expected to be exponential, and indeed functions of the form shown in eq. (5.10) in connection with the cc-pVxZ basis sets usually provide a good fit.³³

$$\begin{aligned} E(L_{\max}) &= E(\infty) + A e^{-B L_{\max}} \\ E(L_{\max}) &= E(\infty) + A(L_{\max} + 1) e^{-B \sqrt{L_{\max}}} \end{aligned} \quad (5.10)$$

An alternative fitting function (eq. (5.11)) for use with the pc- n basis sets has been shown to improve the accuracy of absolute energies by almost an order of magnitude, although relative energies are only marginally improved.³⁴ The number of s -functions (N_s) in the basis set is here used as the main extrapolating parameter.

$$E(L_{\max}) = E(\infty) + A(L_{\max} + 1) e^{-B \sqrt{N_s}} \quad (5.11)$$

Exponential forms like eq. (5.10) have also been used for extrapolating the total energy at correlated levels of theory with the cc-pVxZ basis sets. Theoretical analysis, however, suggest that the *correlation* energy itself (i.e. not the *total* energy, which includes the HF contribution) should converge with an inverse power dependence, with the leading term for singlet electron pairs being $(L + 1)^{-3}$ while the leading term for triplet pairs is $(L + 1)^{-5}$.³⁵ The theoretical assumption underlying these results is that the basis set is saturated in the radial part (e.g. a TZ type basis set should be *complete* in the s -, p -, d - and f -function space). This is not the case for the correlation consistent basis sets: even for the cc-pV6Z basis set, the errors due to insufficient numbers of s - to i -functions are comparable with that from neglect of functions with angular momentum higher than i -functions. Nevertheless, it has been found that extrapolations based on only the leading L^{-3} term give good results when compared with accurate results generated by for example R12 methods.³⁶ This has the advantage that the infinite basis set result can be estimated from only two calculations with basis sets having maximum angular momentum N and M according to eq. (5.12).

$$\Delta E_{\text{corr},\infty} = \frac{N^3 \Delta E_{\text{corr},N} - M^3 \Delta E_{\text{corr},M}}{N^3 - M^3} \quad (5.12)$$

It has been suggested that a separate extrapolation of the singlet (opposite spin) and triplet (same spin) correlation energies with $A + B(L + 1/2)^{-3}$ and $A + B(L + 1/2)^{-5}$ function forms, respectively, may provide better results.³⁷

The main difficulty in using the cc-pVxZ or pc- n basis sets is that each step up in quality roughly doubles the number of basis functions. The fitting functions in eqs (5.10) and (5.11) contain three parameters, and therefore require at least three calculations with increasingly larger basis sets. The simplest sequence is cc-pVDZ, cc-pVTZ and cc-pVQZ, but the cc-pVDZ basis is too small to give good extrapolated values for the correlation energy, and a better sequence is cc-pVTZ, cc-pVQZ and cc-pV5Z. The requirement of performing calculations with at least the cc-pVQZ basis places severe constraints on the size of the systems that can be treated. The extrapolation based on eq. (5.12) has the advantage of requiring only two reference calculations. It should be noted

that the B parameter in eq. (5.11) varies little from system to system, and taking this to be a universal constant also reduces eq. (5.11) to a two-parameter fitting function.

Perhaps the most interesting aspect of the analyses that led to the development of the correlation consistent basis sets is the fact that high angular momentum functions are necessary for achieving high accuracy. While d -polarization functions are sufficient for a DZ type basis, a TZ type should also include f -functions. Similarly, it is questionable to use a QZ type basis for the sp -functions without also including three d -, two f - and one g -function in order to systematically reduce the errors. It can therefore be argued that an extension of for example the 6-31G(d,p) to 6-311G(d,p) is inconsistent as the second set of d -orbitals (and second set of p -orbitals for hydrogen) and a set of f -functions (d -functions for hydrogen) will give similar contributions as the extra set of sp -functions. Similarly, the extension of the 6-311G(2df,2pd) basis to 6-311G(3df,3pd) may be considered inconsistent, as the third d -function is expected to be as important as the fourth valence set of sp -functions, the second set of f -functions and the first set of g -functions, all of which are neglected.

In the search for a basis set converged value, other approximations should be kept in mind. Basis sets with many high angular momentum functions are normally designed for recovering a large fraction of the correlation energy. In the majority of cases, only the electron correlation of the valence electrons is considered (frozen-core approximation), since the core orbitals usually are insensitive to the molecular environment. As the valence space approaches completeness in terms of basis functions, the error from the frozen-core approximation will at some point become comparable to the remaining valence error. From studies of small molecules, where good experimental data are available, it is suggested that the effect of core electron correlation for unproblematic systems is comparable with the change observed upon enlarging the cc-pV5Z basis, i.e. of a similar magnitude as the introduction of h -functions.³⁸ Improvements beyond the cc-pV6Z basis set have been argued to produce changes of similar magnitude to those expected from relativistic corrections for first row elements, and further increases to cc-pV7Z and cc-pV8Z type basis sets would be comparable with corrections due to breakdown of the Born–Oppenheimer approximation for systems with hydrogen. Within the non-relativistic realm, it would therefore appear that basis sets larger than cc-pV6Z would be of little use, except for extrapolating to the non-relativistic, clamped nuclei limit for testing purposes. In attempts at obtaining results of “spectroscopic accuracy” (~ 0.01 kJ/mol), a brute force calculation with for example the cc-pV7Z quality basis set combined with explicit extrapolation has been shown to become problematic,³⁷ and such high-quality results must probably be obtained by explicit correlated techniques, such as the R12 method discussed in Section 4.11.

There is a practical aspect of using large basis sets, especially those including diffuse functions, that requires special attention, namely the problem of *linear dependence*. Linear dependence means that one (or more) of the basis functions can be written as a linear combination of the other, i.e. the basis set is overcomplete. A diffuse function has a small exponent and consequently extends far away from the nucleus on which it is located. An equally diffuse function located on a nearby atom will therefore span almost the same space. A measure of the degree of linear dependence in a basis set can be obtained from the eigenvalues of the overlap matrix \mathbf{S} (eq. (3.51)). A truly linearly dependent basis will have at least one eigenvalue of exactly zero, and the smallest eigenvalue of the \mathbf{S} matrix is therefore an indication of how close the actual basis

set is to linear dependence. As described in Section 16.2.3, solution of the SCF equations requires orthogonalization of the basis by means of the $\mathbf{S}^{-1/2}$ matrix (or a related matrix that makes the basis orthogonal). If one of the \mathbf{S} matrix eigenvalues is close to zero, this means that the $\mathbf{S}^{-1/2}$ matrix is essentially singular, which in turn will cause numerical problems if trying to carry out an actual calculation. In practice, there is therefore an upper limit on how close to completeness a basis set can be chosen to be, and this limit is determined by the finite precision with which the calculations are carried out. If the selected basis set turns out to be too close to linear dependence to be handled, the linear combinations of basis functions with low eigenvalues in the \mathbf{S} matrix may be discarded.

5.5 Plane Wave Basis Functions

Rather than starting with basis functions aimed at modelling the atomic orbitals (STOs or GTOs), and forming linear combination of these to describe orbitals for the whole system, one may use functions aimed directly at the full system. For modelling extended (infinite) systems, for example a unit cell with periodic boundary conditions, this suggests the use of functions with an “infinite” range. The outer valence electrons in metals behave almost like free electrons, which leads to the idea of using solutions for the free electron as basis functions. The solutions to the Schrödinger equation for a free electron in one dimension can be written either in terms of complex exponentials or cosine and sine functions.

$$\begin{aligned}\phi(x) &= Ae^{ikx} + Be^{-ikx} \\ \phi(x) &= A \cos(kx) + B \sin(kx) \\ E &= \frac{1}{2}k^2\end{aligned}\tag{5.13}$$

Note that the energy depends quadratically on the k factor. For infinite systems, the molecular orbitals coalesce into *bands*, since the energy spacing between distinct levels vanishes. The electrons in a band can be described by orbitals expanded in a basis set of *plane waves*, which in three dimensions can be written as a complex function.

$$\chi_k(\mathbf{r}) = e^{i\mathbf{k} \cdot \mathbf{r}}\tag{5.14}$$

The wave vector \mathbf{k} plays the same role as the exponent ζ in a GTO (eq. (5.2)), and is related to the energy by means of eq. (5.13) (conventionally given in units of eV). As seen in eq. (5.14), \mathbf{k} can also be thought of as a frequency factor, with high \mathbf{k} values indicating a rapid oscillation. The permissible \mathbf{k} values are given by the unit cell translational vector \mathbf{t} , i.e. $\mathbf{k} \cdot \mathbf{t} = 2\pi m$, with m being a positive integer. This leads to a typical spacing between \mathbf{k} vectors of ~ 0.01 eV, and the size of the basis set is thus uniquely characterized by the highest energy \mathbf{k} vector included. A typical energy cutoff of 200 eV thus corresponds to a basis set with ~ 20000 functions, i.e. plane wave basis sets tend to be significantly larger than typical Gaussian basis sets. Note, however, that the size of a plane wave basis set depends only on the size of the periodic cell, not on the actual system described within the cell. This is in contrast to the linear increase with

system size for nuclear-centred Gaussian functions, i.e. plane wave basis sets become more favourable for large systems.

While plane wave basis sets have primarily been used for periodic systems, they can also be used for molecular species by using a *supercell* approach, where the molecule is placed in a sufficiently large unit cell such that it does not interact with its own image in the neighbouring cells.³⁹ Placing a relatively small molecule in a large supercell to avoid self-interaction consequently requires many plane wave functions, and such cases are handled more efficiently by localized Gaussian functions. A three-dimensional periodic system, on the other hand, may be better described by a plane wave basis than with nuclear-centred basis functions.

Plane wave basis functions are ideal for describing delocalized slowly varying electron densities, such as the valence bands in a metal. The core electrons, however, are strongly localized around the nuclei, and the valence orbitals have a number of rapid oscillations in the core region to maintain orthogonality. Describing the core region adequately thus requires a large number of rapidly oscillating functions, i.e. a plane wave basis with very large k_{max} . The singularity of the nucleus–electron potential is furthermore almost impossible to describe in a plane wave basis, and this type of basis set is therefore used in connection with pseudopotentials (Section 5.9) for smearing the nuclear charge and modelling the effect of the core electrons. Note that a pseudopotential is also required for smearing the potential near the nucleus in hydrogen, even though hydrogen does not have core electrons.

5.6 Recent Developments and Computational Issues

Recent developments have attempted to combine localized and plane wave basis functions, i.e. describing the core region by radial polynomials or Gaussian functions, and the valence region by plane waves.⁴⁰ This price of this approach is increased computational complexity, since new integrals involving different types of basis function are required.

Harrison and coworkers have recently proposed a multi-resolution procedure where the molecular orbitals are expanded into a set of *wavelets*.⁴¹ The essence of the method is to place the molecule in a suitably sized box, and to repeatedly subdivide the space by a factor of 2, leading to 2^n boxes at level n . The part of the orbitals within each box is expanded into a set of Legendre polynomials of order k , the first few of which are given in eq. (5.6). By increasing the number of boxes and the number of Legendre polynomials, the accuracy can be tuned to any desired degree. For small systems, this approach has been shown to be able to provide HF energies accurate to within 10^{-6} au.⁴² There are a few technical issues regarding how to treat the singularity of the nuclear-electron potential and whether the nuclei are located at box boundaries or not, but these do not appear to be especially problematic. One significant advantage of the multi-resolution method is that it is by construction a linear scaling method.

The computational problem is formally the same whether a Gaussian, plane wave or polynomial basis is used – calculate matrix elements of quantum mechanical operators over basis functions and solve the variational problem by an iterative procedure – but the nature of the functions results in some differences. With a GTO basis the matrix elements are calculated directly, while with a plane wave basis the matrix elements involving the potential energy can be generated by simple multiplication, as long

as the operator is local. The Coulomb operator is local, as is the exchange operator in density functional theory, but the exchange energy at the Hartree–Fock level involves a non-local operator. Incorporating HF exchange with plane wave basis sets is therefore somewhat more difficult than with GTO basis sets,⁴³ and this difference at least partly explains why density functional methods have dominated in solid-state physics, while HF traditionally has been preferred for molecular systems. Another reason is of course that HF theory cannot describe metallic systems – the large band gap predicted by HF makes all periodic systems insulators or semiconductors.

A GTO basis will typically have 10–20 functions per atom, with perhaps a few hundred functions for the whole system. A plane wave basis, on the other hand, will often have tens of thousands of functions. In a traditional implementation, the variational problem is solved by repeated diagonalization of a Fock-type matrix but this becomes problematic when the number of basis functions exceeds a few thousand owing to the cubic scaling of matrix diagonalization. For large plane wave basis sets, the variational problem is therefore often solved by other methods, such as conjugate gradient optimization, quenched dynamics methods or DIIS type extrapolations.⁴⁴ In Car–Parrinello type dynamics (Section 14.2.5), the variational problem is solved by propagating the orbital parameters with fictive masses along with the nuclear degrees of freedom. The computational scaling of solving the variations problem with these methods is $N_{\text{electron}}^2 M_{\text{basis}}$, i.e. the computational time increases linearly with the number of plane wave functions.

A significant advantage of plane wave basis sets is that they are independent of the nuclear positions. This means that the problem of basis set superposition error (Section 5.10) does not occur, and the calculation of the gradient of the energy is easy, as it is given directly by the Hellman–Feynman force, i.e. there are no components associated with the change of basis function position (“Pulay forces”).

5.7 Composite Extrapolation Procedures

In principle, the large majority of systems can be calculated with a high accuracy by using a highly correlated method such as CCSD(T) and performing a series of calculations with systematically larger basis sets in order to extrapolate to the basis set limit. In practice, even a single water molecule is demanding to treat in this fashion (Chapter 11). Various approximate procedures have therefore been developed for estimating the “infinite correlation, infinite basis” limit (Figure 4.3) as efficiently as possible. These models rely on the fact that different properties converge with different rates as the level of sophistication increases, and that effects from extending the basis set to a certain degree are additive. There are four main steps in these procedures:

- (1) Selecting the geometry.
- (2) Selecting a basis set for calculating the Hartree–Fock energy.
- (3) Estimating the electron correlation energy.
- (4) Estimating the energy from translation, rotation and vibrations.

Given a predefined target accuracy, the error from each of these four steps should be reduced below the desired tolerance. The error at a given level may be defined as the change that would occur if the calculation were taken to the “infinite correlation, infinite basis” limit. A typical target accuracy is ~4 kJ/mol (~1 kcal/mol), the so-called

“chemical accuracy”, although some of the more recent methods aim for an accuracy of ~ 1 kJ/mol.

Geometries converge relatively fast: at the HF level with a DZP type basis the “geometry error” is often already ~ 4 kJ/mol or less, and an MP2 or DFT geometry optimized with a DZP basis set is normally sufficient for most applications. The translational and rotational contributions are trivial to calculate, as they depend only on the molecular mass and the geometry (Sections 13.5.1 and 13.5.2), and are very small in absolute values. The error from these can be neglected. The vibrational effect is mainly the zero point energy, and it requires calculation of the frequencies. An accurate prediction of frequencies is fairly difficult. However, since the absolute value of the zero point energy is small, a large relative error is tolerable. Furthermore, the errors in calculated frequencies are to a certain extent systematic and can therefore be improved by a uniform scaling.⁴⁵

The HF error depends only on the size of the basis set. The energy, however, behaves asymptotically as $\sim \exp(-\sqrt{L})$, where L is the highest angular momentum in the basis set. For example, with a basis set of TZP quality (4s3p2d1f for first row elements) the results are already quite stable. Combined with the fact that an HF calculation is the least expensive *ab initio* method, this means that the HF error is rarely the limiting factor.

The main problem is estimating the correlation effect. All electron correlation methods have a rather steep increase in computational cost as the size of the basis is enlarged, and the convergence in terms of the highest angular momentum in the basis is quite slow ($\sim L^{-3}$). The main contribution to the correlation energy is from pairs of electrons in the same spatial MO. This effect is reasonably well described at the MP2 level, but requires a large basis set in order to recover a large fraction of the absolute value. The remaining correlation energy is much harder to calculate: coupled cluster is the preferred method here but, since the absolute value is substantially smaller than the MP2 correlation energy, a smaller basis can be employed. This means that the relative error is quite large but the absolute error is of the same magnitude as the correlation error from the MP2 calculation with the large basis.

In the *Gaussian-1* (G1), *Gaussian-2* (G2),⁴⁶ *Gaussian-3* (G3)⁴⁷ and *Complete Basis Set* (CBS) models, calculations from different levels of theory are combined with the goal of producing energy differences accurate to about 4 kJ/mol, as compared with experimental results. They have been calibrated on a reference set of 125 atomic and molecular properties (atomization energies, ionization potentials, electron and proton affinities) that is often referred to as the G2 or G2-1 data set.⁴⁸ A somewhat larger set of data, called G2-2, has been used more recently.⁴⁹ The ability to accurately calculate atomization energies (corresponding to dissociating molecules into isolated atoms) enables the prediction of absolute values of heat of formation, since the atomic values are known experimentally. The main difference between the G_n and CBS methods is the way they try to extrapolate the correlation energy, as described below. Both the G_n and CBS methods come in different flavours, depending on the exact combinations of methods used for obtaining the above four contributions.

As an example, the G2(MP2) method involves the following steps:⁵⁰

- (1) The geometry is optimized at the HF/6-31G(d) level and the vibrational frequencies are calculated. To correct for the known deficiencies at the HF level, these are scaled by 0.893 to produce zero point energies.

- (2) The geometry is re-optimized at the MP2/6-31G(d) level, which is used as the reference geometry.
- (3) An MP2/6-311+G(3df,2p) calculation is carried out, which automatically yields the corresponding HF energy.
- (4) The energy is calculated at the QCISD(T)/6-311G(d,p) level. This automatically generates the MP2 value as an intermediate result, and the difference between the QCISD(T) and MP2 energies is taken as an estimate of the higher order correlation energy. The G2 method (not G2(MP2)) performs additional MP4 calculations with larger basis sets to get a better estimate of the higher order correlation energy.
- (5) To correct for electron correlation beyond QCISD(T) and basis set limitations, an empirical correction is added to the total energy, $\Delta E_{\text{emp}} = -0.00481N_{\alpha} - 0.00019N_{\beta}$, where it is assumed that the number of β electrons is larger than or equal to the number of α electrons. The numerical constants are determined by fitting to the reference data. It should be noted that this correction makes the G2 methods non-size extensive.

The net effect of steps (3)–(5) is that a single calculation at the QCISD(T)/6-311+G(3df,2p) is replaced by a series of calculations at lower levels, which in combination yields a comparable accuracy in significantly less computer time.⁵¹

The main difference among the G2 models is the way the electron correlation beyond MP2 is estimated. The G2 method itself performs a series of MP4 and QCISD(T) calculations, G2(MP2) only does a single QCISD(T) calculation with the 6-311G(d,p) basis, while G2(MP2,SVP) (SVP stands for split valence polarization) reduces the basis set to only 6-31G(d).⁵² An even more pruned version, G2(MP2,SV), uses the unpolarized 6-31G basis for the QCISD(T) part, which increases the mean absolute deviation (MAD) to 9 kJ/mol. That it is possible to achieve such good performance with this small a basis set for QCISD(T) partly reflects the importance of the large basis MP2 calculation and partly the absorption of errors in the empirical correction.

A comparison between G1, G2, G2(MP2), G2(MP2,SVP) and G3⁵³ is shown in Table 5.5, and for the reference G2 data set the MADs vary from 4.3 to 6.3 kJ/mol. There are other variations of the G2 methods in use, for example involving DFT methods for geometry optimization and frequency calculation⁵⁴ or CCSD(T) instead of QCISD(T),⁵⁵ with slightly varying performance and computational cost. The errors with the G2 method are comparable with those obtained directly from calculations at the CCSD(T)/cc-pVTZ level, at a significantly lower computational cost.⁵⁶

The main difference between the *Gn* and CBS models is the extrapolation of the correlation energy. The *Gn* methods assume basis set additivity and add an empirical correction to recover some of the remaining correlation energy. The CBS procedures, on the other hand, attempt to perform an explicit extrapolation of the calculated values. The main part of the correlation energy is due to electron pairs, i.e. described by doubly excited configurations. In terms of perturbation theory, this may again be divided into contributions from different orders, the most important being from second order (MP2). By using pair natural orbitals (being eigenvectors of the density matrix, see Section 9.5) as the expansion parameter, and assuming that enough pairs have been included to reach the asymptotic limit, it may be shown that the MP2 energy

Table 5.5 Computational levels in the G1/G2/G3 models

Method	Geometry	HF and MP2	Higher order correlation	Thermo [scale factor]	MAD (kJ/mol)
G1	MP2/6-31G(d)	6-311G(2df,p)	MP4/6-311G(d,p) MP4/6-311+G(d,p) MP4/6-311G(2df,p) QCISD(T)/6-311G(d,p)	HF/6-31G(d) [0.893]	6.3
G2	MP2/6-31G(d)	6-311+G(3df,2p)	MP4/6-311G(d,p) MP4/6-311+G(d,p) MP4/6-311G(2df,p) QCISD(T)/6-311G(d,p)	HF/6-31G(d) [0.893]	6.2
G2(MP2)	MP2/6-31G(d)	6-311+G(3df,2p)	QCISD(T)/6-311G(d,p)	HF/6-31G(d) [0.893]	6.3
G3	MP2/6-31G(d)	6-311++G(2df,2p)	QCISD(T)/6-31G(d) MP4/6-31G(d) MP4/6-31+G(d) MP4/6-31G(2df,p)	HF/6-31G(d) [0.893]	4.3

Geometry = level at which the structure is optimized

Higher order correlation = method(s) for estimating higher order correlation effects

Thermo = level at which the thermodynamic corrections are calculated [vibrational scale factor]

MAD = mean absolute deviation relative to the reference data set

calculated by a limited natural orbitals expansion (of the size N_{ij}) behaves as $1/N_{ij}$, and can therefore be extrapolated to the complete basis set limit.

There are several different CBS methods, each having their own set of prescriptions and resulting computational cost and accuracy, and they are known by the acronyms CBS-4, CBS-q, CBS-Q and CBS-APNO.⁵⁷

As an explicit example, we will take the CBS-Q model,⁵⁸ which computationally is similar to the G2(MP2) method:

- (1) The geometry is optimized at the HF/6-31G(d⁺) level (d⁺ denotes that the exponents for the *d*-functions are taken from the 6-311G(d) basis), and the vibrational frequencies are calculated. To correct for the known deficiencies at the HF level, these are scaled by 0.918 to produce zero point energies.
- (2) The geometry is re-optimized at the MP2/6-31G(d⁺) level, which is used as the reference geometry.
- (3) An MP2/6-311+G(2df,2p) calculation is carried out, which automatically yields the corresponding HF energy. The MP2 result is extrapolated to the basis set limit by the pair natural orbital method.
- (4) The energy is calculated at the MP4(SDQ)/6-31G(d,p) and QCISD(T)/6-31+G(d⁺) level to estimate the effect from higher order electron correlation.
- (5) Corrections due to remaining correlation effects are estimated by an empirical expression,

$$\Delta E_{\text{emp}} = -0.00533 \sum_i \left(\sum_{\mu} C_{\mu i} \right)^2 |S|_{ii}^2$$

where the sum over $C_{\mu ii}$ is the trace of the first-order wave function coefficients for the natural orbital pair ii , $|S_{ii}|$ is the absolute value of the spatial overlap between the α and β spin components of the i th MO, and the factor 0.00533 is determined by fitting to the reference data. This empirical correction is size extensive.

- (6) For open-shell species the UHF method is used, which in some cases suffers from spin contamination. To correct for this an empirical correction based on the deviation of $\langle S^2 \rangle$ from the theoretical value is added for the CBS-4 and CBS-Q methods, $\Delta E_{\text{emp}} = -0.0092[\langle S^2 \rangle - S_z(S_z - 1)]$, where the factor of -0.0092 is derived by fitting.

The use of the smaller basis for the QCISD(T) calculation means that the CBS-Q model is computationally faster than G2(MP2), but nevertheless gives slightly lower errors. A comparison among the four CBS models is shown in Table 5.6 (p. 218).

It should be noted that the G2-1 data set, with two exceptions (SO_2 and CO_2), only includes data for molecules containing one or two heavy (non-hydrogen) atoms. It is likely that the typical error for a given model to a certain extent depends on the size of the system, i.e. the G2 method is presumably not able to predict the heat of formation of say C_{60} (if it were computationally feasible) with an accuracy of ~ 6 kJ/mol. Furthermore, the properties included (atomization energies, ionization potentials, electron and proton affinities) all correspond to energy differences between well-separated systems: atomization energies are energy differences between a molecule and isolated atoms, and the other three properties correspond to removal or addition of a single electron or proton. As illustrated in Chapter 11, such energy differences are easier to calculate than between systems containing half broken/formed bond. As with any scheme that has been parameterized on experimental data, it is questionable to assume that the typical accuracy for a selected set of properties will be true in general. A good performance for the G2 data set does not necessarily indicate that the same level of accuracy can be obtained over a wide variety of geometries, for example including transition structures. A modified version of the G2 method, denoted G2Q, involving geometry optimization and frequency calculation at the QCISD/6-311G(d,p) level, has been advocated by Durant and Rohlfing for use with transition structures.⁵⁹

The G3 and CBS-APNO methods are capable of calculating average atomization energies to within 2–4 kJ/mol, but the maximum error for the reference data set is often significantly larger. Since it is difficult to know in advance whether the particular system of interest behaves as the average or the exceptional case, the predicted value must realistically be assumed to have an uncertainty of perhaps 10–20 kJ/mol. Part of the reason for the relatively large spread in the errors is the assumption of additivity in basis sets effect, which has little theoretical foundation, although the empirical corrections at least partly absorb some of these errors.

If higher accuracy is desired, for example “sub-chemical” (~ 0.5 kJ/mol) or “spectroscopic accuracy” (~ 1 cm⁻¹, ~ 0.01 kJ/mol), a number of other factors must also be considered:

- (1) Correlation of the core and core-valence electrons. This becomes progressively more important as heavier elements are considered.

Table 5.6 Computational levels in the CBS models

Method	Geometry	HF	MP2	Higher order correlation	Thermo [scale factor]	MAD (kJ/mol)
CBS-4	HF/3-21G(*)	6-311++G(2df,p)	6-31+G(d')	MP4(SDQ)/6-31G	HF/3-21G [0.917]	8.8
CBS-q	HF/3-21G(*)	6-311++G(2df,p)	6-31+G(d')	MP4(SDQ)/6-31G(d')	HF/3-21G [0.917]	6.7
CBS-Q	MP2/6-31G(d')	6-311++G(2df,2p)	6-311++G(2df,2p)	QCISD(T)/6-31G MP4(SDQ)/6-31+G(d,p) QCISD(T)/6-31+G(d')	HF/6-31G(d')	4.2
CBS-APNO	QCISD/6-311G(d,p)	6s6p3d2f/4s2p1d	6s6p3d2f/4s2p1d	QCISD(T)/6-311+G(2df,p)	[0.918] HF/6-311G(d,p) [0.925]	2.1

Geometry = level at which the structure is optimized

Higher order correlation = method(s) for estimating higher order correlation effects

Thermo = level at which the thermodynamic corrections are calculated [vibrational scale factor]

MAD = mean absolute deviation relative to the reference data set

- (2) Inclusion of high-order correlation effects, such as connected triple, quadruple and quintuple excitations.
- (3) Relativistic effects, such as mass–velocity, Darwin and spin–orbit coupling perturbative corrections, or more sophisticated relativistic treatments. Obviously, these corrections become important even at the chemical accuracy level if elements from the lower part of the periodic table are considered.
- (4) Non-Born–Oppenheimer corrections. These will be most important for systems containing hydrogen.
- (5) Basis set superposition corrections.
- (6) Anharmonic vibrations.
- (7) Vibrational–rotational coupling.

At present, there is no standard procedure for achieving “spectroscopic accuracy”, but Martin and coworkers have developed W1, W2⁶⁰ and W3 methods⁶¹ aimed at a target accuracy of ~1 kJ/mol on the average for atomization energies, with worst-case systems having errors below ~5 kJ/mol. The *Wn* methods all rely on an explicit extrapolation to the infinite basis set limit for the HF and correlation energies, and addition of relativistic effects. The W3 method is the most sophisticated of these, and consists of the following steps:

- (1) The geometry is optimized at the CCSD(T)/cc-pVQZ level.
- (2) Anharmonic frequencies are calculated at the CCSD(T)/cc-pVQZ level.
- (3) The HF limit is estimated by extrapolating the results from the aug-cc-pVQZ and aug-cc-pV5Z basis sets.
- (4) The valence correlation energy is estimated from two-point extrapolation for CCSD, CCSD(T) and CCSDT calculations. The effects of higher order correlation effects are estimated from CCSDTQ/cc-pVDZ results scaled by an empirical factor of 1.25.
- (5) Relativistic corrections are estimated by comparing the results from a Douglas–Kroll CCSD(T)/aug-pRVQZ (relativistic version of the aug-cc-pVQZ basis) calculation with non-relativistic CCSD(T)/aug-pVQZ results.

The W3 method provides a mean absolute error of 0.8 kJ/mol for 30 molecules, with the worst case having an error of ~2 kJ/mol, and these values can be compared with the average experimental error of 0.6 kJ/mol for the same set of data.⁶¹ It can be noted that the experimental data were carefully selected to have small experimental uncertainties, a more typical experimental error is 5–10 kJ/mol. Such explicit extrapolation procedures are thus capable of yielding results with accuracies comparable to experimental methods, and may soon surpass experiments as the preferred method for obtaining geometry and stability data for small- and medium-sized systems.

There are a few other correction procedures that may be considered as extrapolation schemes. The *Scaled External Correlation* (SEC) and *Scaled All Correlation* (SAC) methods scale the correlation energy by a factor such that calculated dissociation energy agrees with the experimental value.⁶²

Table 5.7 Computational levels in the Wn models; only the largest basis set calculations are indicated, results from smaller basis sets are used for extrapolating to the basis set limit

Method	Geometry	HF	Valence correlation	Core correlation	Thermo [scale factor]	Relativistic corrections	MAD (kJ/mol) (data 1) (data 2)
W1	B3LYP/ cc-pVTZ	aug-cc- pVQZ	CCSD/aug-cc-pVQZ CCSD(T)/aug-cc-pVTZ	CCSD(T)/ cc-pVTZ ^a	B3LYP/cc- pVTZ [0.985]	Mass-velocity, Darwin, spin-orbit	1.4 –
W2	CCSD(T)/ cc-pVQZ	aug-cc- pV5Z	CCSD/aug-cc-pV5Z CCSD(T)/aug-cc-pVQZ	CCSD(T)/ cc-pVTZ ^a	CCSD(T)/cc- pVQZ anharmonic	Mass-velocity, Darwin, spin-orbit	1.0 1.7
W3	CCSD(T)/ cc-pVQZ	aug-cc- pV5Z	CCSD/aug-cc-pV5Z CCSD(T)/aug-cc-pVQZ CCSDT/cc-pVTZ CCSDTQ/cc-pVDZ	CCSD(T)/ cc-pVTZ ^a	CCSD(T)/cc- pVQZ anharmonic	Douglas-Kroll CCSD(T)/aug- pRVQZ ^b	– 0.8

^a Augmented with additional tight polarization functions

^b Relativistic version of the cc-pVQZ basis

Geometry = level at which the structure is optimized

Thermo = level at which the thermodynamic corrections are calculated [vibrational scale factor]

MAD = mean absolute deviation relative to the reference data set

$$E_{\text{SEC/SAC}} = E_{\text{ref}} + \frac{E_{\text{corr}} - E_{\text{ref}}}{F} \quad (5.15)$$

$$F = \frac{D_e(\text{corr}) - D_e(\text{ref})}{D_e(\text{exp}) - D_e(\text{ref})}$$

The SEC acronym refers to the case where the reference wave function is of the MCSCF type and the correlation energy is calculated by a MR-CISD procedure. When the reference is a single determinant (HF), the SAC nomenclature is used. In the latter case the correlation energy may be calculated for example by MP2, MP4 or CCSD, producing acronyms such as MP2-SAC, MP4-SAC and CCSD-SAC. In the SEC/SAC procedure, the scale factor F is assumed to be constant over the whole surface. If more than one dissociation channel is important, a suitable average F may be used.

The *Parameterized Configuration Interaction* (PCI-X) method⁶³ simply takes the correlation energy and scales it by a constant factor X (typical value ~ 1.2), i.e. it is assumed that the given combination of method and basis set recovers a constant fraction of the correlation energy.

The introduction of various empirical corrections, such as scale factors for frequencies and energy corrections based on the number of electrons and degree of spin contamination, blurs the distinction between whether they should be considered *ab initio*, or as belonging to the semi-empirical class of methods such as AM1 and PM3. Nevertheless, the accuracy that these methods are capable of delivering makes it possible to calculate absolute stabilities (heat of formation) for small- and medium-sized systems that rival (or surpass) experimental data, often at a substantially lower cost than that for actually performing the experiments.

5.8 Isogyric and Isodesmic Reactions

The most difficult part in calculating absolute stabilities (heat of formation) is the correlation energy. For calculating energies relative to isolated atoms, the goal of the G_n /CBS models, essentially all the correlation energy of the bond being broken must be recovered. This in turn necessitates large basis sets and sophisticated correlation methods. This is also the reason why *ab initio* energies are not converted into heat of formation, as is normally done for semi-empirical methods (eq. (3.97)), since the resulting values are poor unless a very high level of theory is employed.

In many cases, however, it is possible to choose less demanding reference systems than the isolated atoms. Consider for example calculating the C—H dissociation energy of CH_4 . In a direct calculation this is given as the difference in total energy of CH_4 and $\text{CH}_3 + \text{H}$.




Figure 5.4 Dissociation of CH_4

In order to calculate an accurate value for this energy difference, essentially all the electron correlation (and HF) energy for the C—H bond must be recovered.

Consider now the reaction in Figure 5.5.



Figure 5.5 An example of an isogyric reaction

The difference between the two reactions in Figures 5.4 and 5.5 is that the latter has the same number of electron pairs on both sides, and such reactions are called *isogyric*. The task of calculating *all* the correlation energy of a C—H bond is replaced by calculating the *difference* in correlation energy between a C—H and an H—H bond. The latter will benefit from cancellation of errors, and therefore stabilize much earlier in terms of theoretical level. Isogyric reactions can thus be used for obtaining *relative* values. In the above example the CH₄ dissociation energy is given relative to that of H₂. By using the experimental value for H₂, the CH₄ dissociation energy may be calculated quite accurately even at relatively low levels of theory.

The concept may be taken one step further. It is often possible to set up reactions where not only the number of electron pairs is constant, but also the formal type of bonds is the same on both sides. Consider for example calculating the stability of propene by the reaction in Figure 5.6.




Figure 5.6 An example of an isodesmic reaction

In this case the number of C=C, C—C and C—H bonds is the same on both sides, and the “reaction” energy is therefore relatively easy to calculate since the electron correlation is to a large extent the same on both sides. Such reactions that conserve both the number and types of bonds are called *isodesmic* reactions. Combining the calculated energy difference for the left- and right-hand sides with experimental values for H₂C=CH₂, H₃C—CH₃ and CH₄, the (absolute) stability of propene can be obtained reasonably accurate at a quite low level of theory. It does, however, require that the experimental values for the chosen reference compounds be available. Furthermore, there are several possible ways of constructing isodesmic or isogyric reactions (e.g. replacing H with Cl in Figure 5.5), i.e. such methods are not unique.

5.9 Effective Core Potentials

Systems involving elements from the lower part of the periodic table have a large number of core electrons. These are, as already mentioned, unimportant in a chemical sense, but it is necessary to use a large number of basis functions to expand the corresponding orbitals, otherwise the valence orbitals will not be properly described (due to a poor description of the electron–electron repulsion). In the lower half of the periodic table relativistic effects further complicate matters (see Chapter 8). These two problems may be “solved” simultaneously by modelling the core electrons by a suitable function, and treating only the valence electrons explicitly.

The function modelling the core electrons is usually called an *Effective Core Potential* (ECP) in the chemical community,⁶⁴ while the physics community uses the term *Pseudopotential* (PP).⁴⁴ The neglect of an explicit treatment of the core electrons, analogous to the semi-empirical methods in Section 3.10, often gives quite good results at a fraction of the cost of a calculation involving all electrons, and part of the relativis-

tic effects (especially the scalar effects) may also be taken care of, without having to perform the full relativistic calculation.

There are four major steps in designing a pseudopotential:

- (1) Generate a good-quality all-electron wave function for the atom. This will typically be from a numerical Hartree–Fock, a relativistic Dirac–Hartree–Fock or a density functional calculation.
- (2) Replace the valence orbitals by a set of nodeless pseudo-orbitals. The regular valence orbitals will have radial nodes in order to make them orthogonal to the core orbitals, and the pseudo-orbitals are designed such that they behave correctly in the outer part, but without the nodal structure in the core region.
- (3) Replace the core electrons by a potential parameterized by expansion into a suitable set of analytical functions of the nuclear–electron distance, for example a polynomial or a set of spherical Bessel or Gaussian functions. Since relativistic effects are mainly important for the core electrons, this potential can effectively include relativity. The potential may be different for each angular momentum.
- (4) Fit the parameters of the potential such that the solutions of the Schrödinger (or Dirac) equation produce pseudo-orbitals matching the all-electron valence orbitals.

Molecular systems have traditionally been described by Gaussian type basis sets, while plane waves have been favoured for extended (periodic) systems, and this difference has resulted in some differences for the corresponding pseudopotentials. When using Gaussian functions for describing the valence orbitals, it is natural to also use Gaussian functions to describe the ECP. Since Gaussian functions are continuous, there is no fixed distance to characterize the extent of the core potential and the quality of the ECP is determined by the number of electrons chosen to be represented by the ECP. For transition metals, it is clear that the outer $(n+1)s$ -, $(n+1)p$ - and $(n)d$ -orbitals constitute the valence space. While such “full-core” potentials give reasonable geometries, it has been found that the energetics are not always satisfactory. Better results can be obtained by also including the orbitals in the next lower shell in the valence space, albeit at an increase in the computational cost. For silver with an atomic number of 47, for example, one may consider two different choices of core size, where the electrons replaced by an ECP are indicated in *italic* and the remaining electrons in bold:

- “Large-core” ECP: 11 electrons considered explicitly:
 $(1s)^2 (2s)^2 (2p)^6 (3s)^2 (3p)^6 (4s)^2 (3d)^{10} (4p)^6$ **(4d)¹⁰ (5s)¹**
- “Small-core” ECP: 19 electrons considered explicitly:
 $(1s)^2 (2s)^2 (2p)^6 (3s)^2 (3p)^6$ **(4s)² (3d)¹⁰ (4p)⁶ (4d)¹⁰ (5s)¹**
- All-electron ECP: 47 electrons considered explicitly:
(1s)² (2s)² (2p)⁶ (3s)² (3p)⁶ (4s)² (3d)¹⁰ (4p)⁶ (4d)¹⁰ (5s)¹

The shape of the resulting 5s-(pseudo)-orbital for these choices is shown in Figure 5.7.

The gain by using ECPs is largest for atoms in the lower part of the periodic table, especially those where relativistic effects are important. Since fully relativistic results are scarce, the performance of ECPs is somewhat difficult to evaluate by comparing with other calculations,⁶⁵ but they often reproduce known experimental results, thereby justifying the approach. ECPs have also been designed for first row elements




Figure 5.7 The 5s-orbital for Ag with either an all-electron, large- or small-core effective core potential

(Li—Ne),⁶⁶ although the savings in these cases are marginal relative to all electron calculations.

The size of a plane wave basis set is given by the highest energy wave, which is inversely related to the smallest variation of the wave function that can be described. The singularity of the nuclear potential (V_{ne}) and the resulting strongly localized core electrons are essentially impossible to describe by any reasonable-sized plane wave basis set. Pseudopotentials are therefore used for smearing the nuclear charge and modelling the core electrons. These potentials are typically characterized by a “core radius” r_c (which may depend on the angular momentum of the valence orbitals), i.e. the pseudopotentials used in connection with plane waves have a finite physical extent. The potential for distances smaller than r_c is described by a suitable analytical function, typically a polynomial or spherical Bessel function, and the pseudo-wave function and its first and second derivatives are required to match those of the reference wave function at r_c . It is clear that a “hard” (small r_c) pseudopotential will require more plane wave basis functions for describing the region beyond r_c than a “soft” (large r_c) pseudopotential, but a too large r_c will deteriorate the quality of the calculated results and also make the pseudopotential less transferable.

The *norm-conserving* pseudopotentials proposed by Hamann, Schlüter and Chiang require in addition to the above matching conditions at r_c that the integral of the square of the reference and pseudo-wave from 0 to r_c agrees, i.e. conservation of the wave function norm.⁶⁷ For the late first row elements (C—F) and the 3d transition-metals (Sc—Zn), these pseudopotentials are rather “hard” and therefore require a relatively large energy cutoff for the plane waves. Vanderbilt proposed to relax the norm-conserving requirement to give the so-called *ultrasoft* pseudopotentials,⁶⁸ thereby reducing the necessary number of plane wave for expanding the valence orbitals by roughly a factor of 2.

While there is essentially no basis set error when using plane waves for expanding the orbitals, the requirement of using a pseudopotential to describe the core region means that there is a fundamental limitation in how accurate the results can be. For systems composed of elements from the first two rows in the periodic table, the error in DFT calculations is roughly equivalent to that imposed by using a Gaussian basis set of TZP quality in an all-electron calculation.⁶⁹ An implicit limitation of pseudopotential methods is of course the inability to describe molecular properties that depend directly on the core electrons (as in X-ray photoelectron spectroscopy) or the electron density near the nucleus (as in NMR shielding and coupling constants).

The *Projector Augmented Wave* (PAW) method is usually also considered a pseudopotential method, although it formally retains all the core electrons.⁷⁰ Indeed, the Vanderbilt ultrasoft pseudopotential can be derived by linearization of two terms in the PAW expression. The PAW wave function is written as a valence term expanded in a plane wave basis plus a contribution from the region within the core radius of each nucleus, evaluated on a grid. The contribution from a core region is expanded as a difference between two sets of densities, one arising from the (all-electron) atomic orbitals, the other from a set of nodeless pseudo-atomic orbitals, i.e. this term allows the wave function within the core region to adjust for different environments. In all the applications so far, however, the (all-electron) atomic orbitals have been kept fixed at their form for the isolated atoms, i.e. a “frozen-core” approach. If the atomic orbitals were fully optimized, the PAW would be essentially equivalent to the all-electron mixed basis sets methods discussed in Section 5.6.⁴⁰ While the PAW method in principle should be better than using pseudopotential method, very few explicit comparisons have so far been performed.

A closely related idea arising from the chemical community is the use of the *frozen-core* approximation.⁷¹ The core electrons are here included in the treatment, but the corresponding orbitals are fixed at their atomic values and represented by a fixed expansion in a suitable basis set. This preserves the full electron–electron interaction but ignores the change in the core orbitals due to the molecular environment. For first row systems the savings are marginal but for heavier elements the computational cost may be significantly reduced. The frozen-core approximation may furthermore be useful for calculations using relativistic wave functions, as it effectively prevents a variational collapse.

A common feature of all pseudopotential methods is that the parameters depend on the employed method, i.e. the potential derived for e.g. the *Local Spin Density Approximation* (LSDA) functional (Section 6.5.1) is different from that derived from a generalized gradient functional such as *Perdew–Burke–Ernzerhof* (PBE) (Section 6.5.2). In practice, the difference is relatively small and pseudopotentials optimized for one functional are often used for other functionals without re-optimization.

5.10 Basis Set Superposition Errors

By far the most common type of basis set for molecular applications is centred on the nuclei. As a complete basis set cannot be used in practice, the M_{basis}^4 (or worse) increase in computational effort limits practical calculations to hundreds or a few thousand basis functions at best. For most systems this means that absolute errors in the energy from basis set incompleteness are quite large, maybe several au (thousands of kJ/mol).

The interest is usually in relative energies, however, and the primary goal is therefore to make the error as constant as possible. This is one of the reasons why it is important to choose a “balanced” basis set. The first, perhaps obvious, step is that the same basis set must be used when comparing energies: comparing energies of two isomers where the 6-31G basis set has been used for one of them, and the DH basis set for the other, is meaningless, although both basis sets are of double zeta quality.

Fixing the position of the basis functions to the nuclei allows for a compact basis set, otherwise sets of basis functions positioned at many points in the geometrical space would be needed. When comparing energies at different geometries, however, the nuclear fixed basis set introduces an error. The quality of the basis set is not the same at all geometries, owing to the fact that the electron density around one nucleus may be described by functions centred at another nucleus. This is especially troublesome when calculating small effects, such as energies of van der Waals complexes and hydrogen bonds. Consider for example the hydrogen bond between two water molecules. The simplest approach consists of calculating the energy of the dimer and subtracting two times the energy of an isolated molecule (assuming a size extensive method). The electron distribution within each water molecule in the dimer is very close to that of the monomer. In the dimer, however, basis functions from one molecule can help compensate for the basis set incompleteness on the other molecule, and vice versa. The dimer will therefore be artificially lowered in energy, and the strength of the hydrogen bond overestimated. This effect is known as the *Basis Set Superposition Error* (BSSE). In the limit of a complete basis set, the BSSE will be zero, and adding more basis functions will not give any improvement. The conceptually simplest approach for eliminating BSSE is therefore to add more and more basis functions, until the interaction energy no longer changes. Unfortunately, this requires very large basis sets. Since non-bonded interactions are weak, the desired accuracy is often ~ 0.5 kJ/mol. Using the correlation consistent basis sets, the water dimer interaction energy stabilizes at this level with the aug-cc-pVTZ basis (184 basis functions for H₂O) at the HF level, but requires (at least) the aug-cc-pV5Z basis (574 basis functions) at the MP2 level.⁷² As inclusion of electron correlation is mandatory for calculating the dispersion interaction between molecules, even the water dimer potential is computationally challenging.

An *approximate* way of assessing BSSE is the *Counterpoise* (CP) correction.⁷³ In this method the BSSE is estimated as the difference between monomer energies with the regular basis and the energies calculated with the full set of basis functions for the whole complex. Consider two molecules A and B, each having regular nuclear-centred basis sets denoted with subscripts a and b, and the complex AB having the combined basis set ab. The geometries of the two isolated molecules and of the complex are first optimized or otherwise assigned. The geometries of the A and B molecules in the complex will usually be slightly different than for the isolated species, and the complex geometry will be denoted with a*. The dimer energy minus the monomer energies is the directly calculated complexation energy.

$$\Delta E_{\text{complexation}} = E(\text{AB})_{\text{ab}}^* - E(\text{A})_{\text{a}} - E(\text{B})_{\text{b}} \quad (5.16)$$

To estimate how much of this complexation energy is due to BSSE, four additional energy calculations are needed. Using the a basis set for A, and the b basis set for B, the energies of each of the two fragments are calculated *with the geometry they have*

in the complex. Two additional energy calculations of the fragments at the complex geometry are then carried out with the full ab basis set. This means that the energy of A is calculated in the presence of both the normal a basis functions *and* with the b basis functions of fragment B located at the corresponding nuclear positions, but *without* the B nuclei present, and vice versa. Such basis functions located at fixed points in space are often referred to as *ghost orbitals*. The fragment energy for A will be lowered due to these ghost functions, since the a basis becomes more complete. The CP correction is defined in eq. (5.17).

$$\Delta E_{\text{CP}} = E(\text{A})_{\text{ab}}^* + E(\text{B})_{\text{ab}}^* - E(\text{A})_{\text{a}}^* - E(\text{B})_{\text{b}}^* \quad (5.17)$$

The counterpoise-corrected complexation energy is then given as $\Delta E_{\text{complexation}} - \Delta E_{\text{CP}}$. For regular basis sets, this typically stabilizes at the basis set limiting value much earlier than uncorrected values, but this is not necessarily the case if diffuse functions are included in the basis set. Note that ΔE_{CP} is an approximate correction: it gives an *estimate* of the BSSE effect but does not provide either an upper or lower limit.

There are variations of this method. For example may it be argued that the full set of ghost orbitals should not be used, since some of the functions in the complex are used for describing the electrons of the other component, and only the virtual orbitals are available for “artificial” stabilization. However, it appears that the method of full counterpoise correction (using all basis functions as ghost orbitals) gives the best results.

It is usually observed that the CP correction for methods including electron correlation is larger and more sensitive to the size of the basis set than at the HF (or DFT) level. This is in line with the fact that the HF wave function converges much faster with respect to the size of the basis set than correlated wave functions.

There have also been attempts at developing methods where the BSSE is excluded explicitly in the computational expressions. An example of this is the *Chemical Hamiltonian Approach* (CHA),⁷⁴ but such methods are not yet commonly used.

The BSSE is always present, also in calculating energies of “normal” species, for example the differential stability of ethanol and dimethyl ether, or the conformational difference between staggered and eclipsed ethane.⁷⁵ Indeed, part of what is often referred to as the “basis set effect” (the change in relative energies when the basis is enlarged) should more correctly be considered as intramolecular BSSE. For intermolecular (non-bonded) interactions the CP correction is well defined, although it may not be as accurate as desired. For intramolecular cases, however, it is difficult to define a unique procedure for estimating the BSSE, and it is almost always ignored.

5.11 Pseudospectral Methods

The goal of *pseudospectral* methods⁷⁶ is to reduce the formal M^4 dependence of the Coulomb and exchange operators in the basis set representation (two-electron integrals, eq. (3.52)) to M^3 . This can be accomplished by switching between a *grid* representation in the *physical* space (the three-dimensional Cartesian space) and the *spectral* representation in the *function* space (the basis set).

Consider the following Coulomb contribution to the $F_{\alpha\beta}$ element of the Fock matrix (eq. (3.52)); similar considerations hold for the exchange contribution.

$$\langle \chi_\alpha | \mathbf{F} | \chi_\beta \rangle \leftarrow \sum_j^{N_{\text{elec}}} \langle \chi_\alpha | \mathbf{J}_j | \chi_\beta \rangle = \sum_j^{N_{\text{elec}}} \langle \chi_\alpha \phi_j | \mathbf{g} | \chi_\beta \phi_j \rangle \quad (5.18)$$

Written out in terms of the actual integral, this is given by eq. (5.19).

$$\langle \chi_\alpha \phi_j | \mathbf{g} | \chi_\beta \phi_j \rangle = \iint \chi_\alpha(\mathbf{r}_1) \phi_j(\mathbf{r}_2) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \chi_\beta(\mathbf{r}_1) \phi_j(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2 \quad (5.19)$$

For a specific point in space for coordinate 1, \mathbf{r}_g , the integration over coordinate 2 may be carried out as shown in eq. (5.20).

$$\langle \chi_\alpha \phi_j | \mathbf{g} | \chi_\beta \phi_j \rangle = \int \chi_\alpha(\mathbf{r}_g) \chi_\beta(\mathbf{r}_g) d\mathbf{r}_1 \left(\int \frac{\phi_j^2(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_g|} d\mathbf{r}_2 \right) \quad (5.20)$$

The last integral may be written in terms of atomic quantities, as in eq. (5.21).

$$\int \frac{\phi_j^2(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_g|} d\mathbf{r}_2 = \sum_{\gamma\delta}^{M_d} c_{\gamma j} c_{\delta j} \int \chi_\gamma(\mathbf{r}_2) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_g|} \chi_\delta(\mathbf{r}_2) d\mathbf{r}_2 = \sum_{\gamma\delta}^{M_d} c_{\gamma j} c_{\delta j} A_{\gamma\delta}(\mathbf{r}_g) \quad (5.21)$$

The $A_{\gamma\delta}$ integral is just a three-centre one-electron integral, which can be evaluated analytically. The integration over coordinate 1 may then be approximated as a sum over a finite set of grid points in the physical space.

$$\langle \chi_\alpha \phi_j | \mathbf{g} | \chi_\beta \phi_j \rangle \cong \sum_{g=1}^G \chi_\alpha(\mathbf{r}_g) \chi_\beta(\mathbf{r}_g) \sum_{\gamma\delta=1}^{M_{\text{basis}}} c_{\gamma j} c_{\delta j} A_{\gamma\delta}(\mathbf{r}_g) \quad (5.22)$$

As the number of grid points increases, this approximation becomes better. The reduction in the formal scaling from M^4 to $\sim M^3$ comes from the fact that the summations involve GM^2 operations, G being the number of grid points, which will typically be linearly dependent on the number of basis functions M , i.e. $GM^2 \sim M^3$.

Unfortunately, the above formula does not work well unless a very large number of grid points are used. This is due to an effect known as *aliasing*, where the physical space Coulomb operator $\mathbf{J}(\mathbf{r}_g)$ acting on the basis function χ_β produces a result that has components outside the basis set. In practice the $\mathbf{J}(\mathbf{r}_g)\chi_\beta$ product is therefore *fitted* to a larger *dealiasing* basis set, typically constructed from the original basis set by adding functions with exponents intermediate to those already present, and polarization functions with one higher angular momentum than already present.

$$\mathbf{J}_j(\mathbf{r}_g) \chi_\beta(\mathbf{r}_g) \cong \sum_{\sigma}^{M^*} W_{\alpha\beta}^* \chi_{\sigma}^*(\mathbf{r}_g) \quad (5.23)$$

The full set of dealiasing basis is denoted χ_{σ}^* , and contains M^* functions. The weights $W_{\alpha\beta}^*$ are assigned based on a least squares fitting procedure. A similar scheme may be constructed for the exchange operator \mathbf{K} . By careful control of the grid size and the dealiasing basis, and by analytical evaluation of the one-centre (and sometimes also the two-centre) Coulomb and exchange contributions, which are computationally insignificant compared with the three- and four-centre integrals, pseudospectral methods can provide energies at the same accuracy as fully analytical methods.

While the formal scaling of pseudospectral methods is M^3 , compared with M^4 for all-integral methods, the effective scaling has been found to be very similar for the two approaches in actual calculations where integral screening is employed to eliminate small contributions to the Fock matrix.⁷⁷ The prefactor for pseudospectral methods, however, is somewhat smaller and leads to an order of magnitude faster computational time for medium-sized systems. It is unclear how the timing of pseudospectral methods compares with linear scaling methods discussed in Section 3.8.6 for large systems.

References

1. D. Feller, E. R. Davidson, *Rev. Comp. Chem.*, **1** (1990), 1; T. Helgaker, P. R. Taylor, *Modern Electronic Structure Theory*, Part II, D. Yarkony, Ed., World Scientific, 1995, pp. 727–856; T. Helgaker, P. Jørgensen, J. Olsen, *Molecular Electronic Structure Theory*, Wiley, 2000.
2. J. C. Slater, *Phys. Rev.*, **36** (1930), 57.
3. S. F. Boys, *Proc. R. Soc. (London) A*, **200** (1950), 542.
4. W. Klopper, W. Kutzelnigg, *J. Mol. Struct.*, **135** (1986), 339; W. Kutzelnigg, J. D. Morgan III, *J. Chem. Phys.*, **96** (1992), 4484.
5. M. W. Schmidt, K. Ruedenberg, *J. Chem. Phys.*, **71** (1979), 3951.
6. S. Huzinaga, M. Klobukowski, H. Tatewski, *Can. J. Chem.*, **63** (1985), 1812.
7. G. A. Petersson, S. Zhong, J. A. Montgomery Jr, M. J. Frisch, *J. Chem. Phys.*, **118** (2003), 1101.
8. <http://www.emsl.pnl.gov/forms/basisform.html>.
9. F. Jensen, *J. Chem. Phys.*, **122** (2005), 074111.
10. E. R. Davidson, *Chem. Phys. Lett.*, **260** (1996), 514.
11. W. J. Hehre, R. F. Stewart, J. A. Pople, *J. Chem. Phys.*, **51** (1969), 2657.
12. J. S. Binkley, J. A. Pople, *J. Am. Chem. Soc.*, **102** (1980), 939.
13. W. J. Hehre, R. Ditchfield, J. A. Pople, *J. Chem. Phys.*, **56** (1972), 2257.
14. R. Krishnan, J. S. Binkley, R. Seeger, J. A. Pople, *J. Chem. Phys.*, **72** (1980), 650.
15. M. J. Frisch, J. A. Pople, J. S. Binkley, *J. Chem. Phys.*, **809** (1984), 3265.
16. M. M. Francl, W. J. Pietro, W. J. Hehre, J. S. Binkley, M. S. Gordon, D. J. DeFrees, J. A. Pople, *J. Chem. Phys.*, **77** (1982), 3654.
17. S. Huzinaga, *J. Chem. Phys.*, **42** (1965), 1293.
18. F. B. van Duijneveldt, *IBM Tech. Res. Rep RJ945* (1971).
19. H. Partridge, *J. Chem. Phys.*, **90** (1989), 1043.
20. T. H. Dunning, *J. Chem. Phys.*, **55** (1971), 716.
21. A. D. McLean, G. S. Chandler, *J. Chem. Phys.*, **72** (1980), 5639.
22. H. Tatewaki, S. Huzinaga, *J. Comp. Chem.*, **1** (1980), 205.
23. A. Schafer, H. Horn, R. Ahlrichs, *J. Chem. Phys.*, **97** (1992), 2571; A. Schafer, C. Huber, R. Ahlrichs, *J. Chem. Phys.*, **100** (1994), 5829.
24. F. Weigend, F. Furche, R. Ahlrichs, *J. Chem. Phys.*, **119** (2001), 12753.
25. J. Almlöf, P. R. Taylor, *J. Chem. Phys.*, **92** (1990), 551; J. Almlöf, P. R. Taylor, *Adv. Quant. Chem.*, **22** (1991), 301.
26. T. H. Dunning Jr, *J. Chem. Phys.*, **90** (1989), 1007; A. K. Wilson, T. van Mourik, T. H. Dunning Jr, *J. Mol. Struct.*, **388** (1996), 339.
27. T. H. Dunning Jr, K. A. Peterson, A. K. Wilson, *J. Chem. Phys.*, **114** (2001), 9244.
28. R. A. Kendall, T. H. Dunning Jr, R. J. Harrison, *J. Chem. Phys.*, **96** (1992), 6796.
29. D. E. Woon, T. H. Dunning Jr, *J. Chem. Phys.*, **103** (1995), 4572.
30. F. Jensen, *Theor. Chem. Acc.*, **104** (2000), 484; K. Aa. Christensen, F. Jensen, *Chem. Phys. Lett.*, **317** (2000), 400.
31. F. Jensen, *J. Chem. Phys.*, **115** (2001), 9113; F. Jensen, *J. Chem. Phys.*, **116** (2002), 3502.
32. A. K. Wilson, T. H. Dunning Jr, *J. Chem. Phys.*, **106** (1997), 8718.

33. A. Karton, J. M. L. Martin, *Theor. Chem. Acc.*, **115** (2006), 330.
34. F. Jensen, *Theor. Chem. Acc.*, **113** (2005), 267.
35. W. Kutzelnigg, J. D. Morgan III, *J. Chem. Phys.*, **96** (1992), 4484; W. Kutzelnigg, J. D. Morgan III, *J. Chem. Phys.*, **97** (1992), 8821.
36. T. Helgaker, W. Klopper, H. Koch, J. Noga, *J. Chem. Phys.*, **106** (1997), 9639.
37. E. F. Valeev, W. D. Allen, R. Hernandez, C. D. Sherrill, H. F. Schaefer III, *J. Chem. Phys.*, **118** (2003), 8594.
38. K. A. Peterson, A. K. Wilson, D. E. Woon, T. H. Dunning Jr, *Theor. Chem. Acc.*, **97** (1997), 251.
39. M. Preuss, W. G. Schmidt, K. Seino, J. Furthmüller, F. Bechstedt, *J. Comp. Chem.*, **25** (2004), 112.
40. F. A. Pahl, N. C. Handy, *Mol. Phys.*, **100** (2002), 3199; L. Fusti-Molnar, P. Pulay, *J. Chem. Phys.*, **116** (2002), 7795.
41. R. J. Harrison, G. I. Fann, T. Yanai, Z. Gan, G. Beylkin, *J. Chem. Phys.*, **121** (2004), 11587.
42. T. Yanai, G. I. Fann, Z. Gan, R. J. Harrison, G. Beylkin, *J. Chem. Phys.*, **121** (2004), 6680.
43. S. Chawla, G. A. Voth, *J. Phys. Chem.*, **108** (1998), 4697.
44. M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, J. D. Joannopoulos, *Rev. Mod. Phys.*, **64** (1992), 1045.
45. A. P. Scott, L. Radom, *J. Phys. Chem.*, **100** (1996), 16502.
46. L. A. Curtiss, K. Raghavachari, G. W. Trucks, J. A. Pople, *J. Chem. Phys.*, **94** (1991), 7221.
47. L. A. Curtiss, K. Raghavachari, P. C. Redfern, V. Rassolov, J. A. Pople, *J. Chem. Phys.*, **109** (1998), 7764.
48. L. A. Curtiss, K. Raghavachari, G. W. Trucks, J. A. Pople, *J. Chem. Phys.*, **94** (1991), 7221.
49. L. A. Curtiss, K. Raghavachari, P. C. Redfern, J. A. Pople, *J. Chem. Phys.*, **106** (1997), 1063.
50. L. A. Curtiss, K. Raghavachari, J. A. Pople, *J. Chem. Phys.*, **98** (1993), 1293.
51. L. A. Curtiss, J. E. Carpenter, K. Raghavachari, J. A. Pople, *J. Chem. Phys.*, **96** (1992), 9030.
52. L. A. Curtiss, P. C. Redfern, B. J. Smith, L. Radom, *J. Chem. Phys.*, **104** (1996), 5148.
53. L. A. Curtis, K. Raghavachari, P. C. Redfern, V. Rassolov, J. A. Pople, *J. Chem. Phys.*, **109** (1998), 7764.
54. C. W. Bauschlicher, H. Partridge, *J. Chem. Phys.*, **103** (1995), 1788; A. M. Mebel, K. Morokuma, M. C. Lin, *J. Chem. Phys.*, **103** (1995), 7414.
55. L. A. Curtiss, K. Raghavachari, J. A. Pople, *J. Chem. Phys.*, **103** (1995), 4192.
56. J. M. L. Martin, *J. Chem. Phys.*, **100** (1994), 8186.
57. J. A. Montgomery Jr, J. W. Ochterski, G. A. Petersson, *J. Chem. Phys.*, **101** (1994), 5900.
58. J. W. Ochterski, G. A. Petersson, J. A. Montgomery Jr, *J. Chem. Phys.*, **104** (1996), 2598.
59. J. L. Durant Jr, G. M. Rohlfing, *J. Chem. Phys.*, **98** (1993), 8031.
60. J. M. L. Martin, G. Oliveira, *J. Chem. Phys.*, **111** (1999), 1843.
61. A. D. Boese, M. Oren, O. Atasoylu, J. M. L. Martin, M. Kallay, J. Gauss, *J. Chem. Phys.*, **120** (2004), 4129.
62. J. Rossi, D. G. Truhlar, *Chem. Phys. Lett.*, **234** (1995), 64.
63. P. E. M. Siegbahn, M. Svensson, P. J. E. Boussard, *J. Chem. Phys.*, **102** (1995), 5377.
64. G. Frenking, I. Antes, M. Böhme, S. Dapprich, A. W. Ehlers, V. Jonas, A. Nauhaus, M. Otto, R. Stegmann, A. Veldkamp, S. F. Vyboishchikov, *Rev. Comp. Chem.*, **8** (1996), 63; T. R. Cundari, M. T. Benson, M. L. Lutz, S. O. Sommerer, *Rev. Comp. Chem.*, **8** (1996), 145.
65. See however K. Dyall, *J. Chem. Phys.*, **96** (1991), 1210.
66. W. J. Stevens, H. Basch, M. Krauss, *J. Chem. Phys.*, **81** (1984), 6026.
67. D. R. Hamann, M. Schlüter, C. Chiang, *Phys. Rev. Lett.*, **43** (1979), 1494.
68. D. Vanderbilt, *Phys. Rev. B*, **41** (1990), 7892; G. Kresse, J. Hafner, *J. Phys. Condens. Matter*, **6** (1994), 8245.
69. C. Janfelt, F. Jensen, *Chem. Phys. Lett.*, **406** (2005), 501; F. Jensen, C. Janfelt, *Chem. Phys. Lett.*, **412** (2005), 12.

70. P. E. Blöchl, *Phys. Rev. B*, **50** (1994), 17953; G. Kresse, D. Joubert, *Phys. Rev. B*, **59** (1999), 1758.
71. E. J. Baerends, D. E. Ellis, P. Ros, *Chem. Phys.*, **2** (1973), 41; G. te Velde, F. M. Bickelhaupt, E. J. Baerends, C. Fonseca Guerra, S. J. A. van Gisbergen, J. G. Snijders, T. Ziegler, *J. Comp. Chem.*, **22** (2001), 931.
72. A. Halkier, H. Koch, P. Jørgensen, O. Christiansen, I. M. B. Nielsen, T. Helgaker, *Theor. Chem. Acc.*, **97** (1997), 150.
73. F. B. van Duijneveldt, J. G. C. M. van Duijneveldt-van de Rijdt, J. H. van Lenthe, *Chem. Rev.*, **94** (1994), 1873.
74. I. Mayer, A. Vibok, *Mol. Phys.*, **92** (1997), 503.
75. F. Jensen, *Chem. Phys. Lett.*, **261** (1996), 633.
76. B. H. Greedy, T. V. Russo, D. T. Mainz, R. A. Friesner, J.-M. Langlois, W. A. Goddard III, R. E. Donnally, M. N. Ringalda, *J. Chem. Phys.*, **101** (1994), 4028; R. A. Friesner, R. B. Murphy, M. N. Ringalda, *Encyclopedia Comp. Chem.*, **3** (1998), 2290.
77. Y. Cao, R. A. Friesner, *J. Chem. Phys.*, **122** (2005), 104102.

6 Density Functional Methods

The basis for *Density Functional Theory* (DFT) is the proof by Hohenberg and Kohn¹ that the ground state electronic energy is determined completely by the electron density ρ (see Appendix B for details). In other words, there exists a one-to-one correspondence between the electron density of a system and the energy. The “intuitive” proof of why the density completely defines the system is due to E. B. Wilson,² who argued that:

- The integral of the density defines the number of electrons.
- The cusps in the density define the position of the nuclei.
- The heights of the cusps define the corresponding nuclear charges.

The significance of the Hohenberg–Kohn theorem is perhaps best illustrated by comparing it with the wave function approach. A wave function for an N electron system contains $4N$ variables, three spatial and one spin coordinate for each electron. The electron density is the square of the wave function, integrated over $N - 1$ electron coordinates, and each spin density only depends on three spatial coordinates, *independent* of the number of electrons. While the complexity of a wave function increases exponentially with the number of electrons, the electron density has the same number of variables, independent of the system size. The “only” problem is that although it has been proven that each different density yields a different ground state energy, the functional connecting these two quantities is not known. The goal of DFT methods is to design functionals connecting the electron density with the energy.^{3,4}

A note on semantics: a *function* is a prescription for producing a number from a set of *variables* (coordinates). A *functional* is a prescription for producing a number from a *function*, which in turn depends on variables. A wave function and the electron density are thus functions, while the energy depending on a wave function or an electron density is a functional. We will denote a function depending on a set of variables with parenthesis, $f(\mathbf{x})$, while a functional depending on a function is denoted with brackets, $F[f]$.

Early attempts at designing DFT models (actually predating wave mechanics) tried to express all the energy components as a functional of the electron density but these methods had poor performance, and wave function-based methods were consequently preferred. The success of modern DFT methods is based on the suggestion by Kohn and Sham in 1965 that the electron kinetic energy should be calculated from an auxiliary set of orbitals used for representing the electron density.⁵ The exchange–correlation energy, which is a rather small fraction of the total energy, is then the only unknown functional, and even relatively crude approximations for this term provide quite accurate computational models. The simplest model is the local density approximation, where the electron density is assumed to be slowly varying, such that the exchange–correlation energy can be calculated using formulas derived for a uniform electron density. A significant improvement in the accuracy can be obtained by making the exchange–correlation functional dependent also on the first derivative of the density, and further refinements also add the second derivative and mix Hartree–Fock exchange into the functional. Density functional theory is conceptually and computationally very similar to Hartree–Fock theory, but provides much better results and has consequently become a very popular method. The main problem in DFT is the inability to systematically improve the results, and the known failure to describe certain important features, such as van der Waals interactions.

6.1 Orbital-Free Density Functional Theory

Compared with the wave mechanics approach, it seems clear that the energy functional may be divided into three parts, kinetic energy, $T[\rho]$, attraction between the nuclei and electrons, $E_{\text{ne}}[\rho]$, and electron–electron repulsion, $E_{\text{ee}}[\rho]$ (the nuclear–nuclear repulsion is a constant within the Born–Oppenheimer approximation). Furthermore, with reference to Hartree–Fock theory (eq. (3.32)), the $E_{\text{ee}}[\rho]$ term may be divided into Coulomb and exchange parts, $J[\rho]$ and $K[\rho]$, implicitly including correlation energy in all the terms. The $E_{\text{ne}}[\rho]$ and $J[\rho]$ functionals are given by their classical expressions, where the factor of $1/2$ in $J[\rho]$ allows the integration to be over all space for both variables.

$$\begin{aligned} E_{\text{ne}}[\rho] &= - \sum_a^{N_{\text{nuclei}}} \int \frac{Z_a(\mathbf{R}_a)\rho(\mathbf{r})}{|\mathbf{R}_a - \mathbf{r}|} d\mathbf{r} \\ J[\rho] &= \frac{1}{2} \iint \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \end{aligned} \quad (6.1)$$

Early attempts of deducing functionals for the kinetic and exchange energies considered a uniform electron gas, where it may be shown that $T[\rho]$ and $K[\rho]$ are given by eq. (6.2).

$$\begin{aligned} T_{\text{TF}}[\rho] &= C_{\text{F}} \int \rho^{5/3}(\mathbf{r}) d\mathbf{r} \\ K_{\text{D}}[\rho] &= -C_{\text{x}} \int \rho^{4/3}(\mathbf{r}) d\mathbf{r} \\ C_{\text{F}} &= \frac{3}{10} (3\pi^2)^{2/3} \\ C_{\text{x}} &= \frac{3}{4} \left(\frac{3}{\pi} \right)^{1/3} \end{aligned} \quad (6.2)$$

The energy functional $E_{\text{TF}}[\rho] = T_{\text{TF}}[\rho] + E_{\text{ne}}[\rho] + J[\rho]$ is known as *Thomas–Fermi* (TF) theory, while inclusion of the $K_{\text{D}}[\rho]$ exchange part (first derived by Bloch,⁶ but commonly associated with the name of Dirac⁷) constitutes the *Thomas–Fermi–Dirac* (TFD) model.

The assumption of a uniform electron gas is fair for the valence electrons in certain metallic (periodic) systems, but is poor for atoms and molecules. A serious flaw from a chemical point of view is it that neither TF nor TFD theories predict bonding: molecules simply do not exist.

The kinetic and exchange functionals can be improved by the addition of terms depending on the derivative(s) of the electron density. This is equivalent to considering a non-uniform electron gas and performing a Taylor-like expansion with the density as a variable.^{8,9} The expansion for the kinetic energy is given in eq. (6.3), where odd terms vanish owing to the rotational invariance of the energy with respect to \mathbf{r} .

$$\begin{aligned} T[\rho] &= T_{\text{TF}}[\rho] + T_2[\rho] + T_4[\rho] + T_6[\rho] + \dots \\ T_2[\rho] &= \lambda \tau_{\text{w}}[\rho]; \quad \tau_{\text{w}}[\rho] = \int \frac{|\nabla \rho(\mathbf{r})|^2}{8\rho(\mathbf{r})} d\mathbf{r} \\ T_4[\rho] &= \left(540(3\pi)^{2/3}\right)^{-1} \int \rho^{1/3}(\mathbf{r}) \left\{ \left(\frac{\nabla^2 \rho(\mathbf{r})}{\rho(\mathbf{r})} \right)^2 - \frac{9}{8} \left(\frac{\nabla^2 \rho(\mathbf{r})}{\rho(\mathbf{r})} \right) \left(\frac{\nabla \rho(\mathbf{r})}{\rho(\mathbf{r})} \right)^2 + \frac{1}{3} \left(\frac{\nabla \rho(\mathbf{r})}{\rho(\mathbf{r})} \right)^4 \right\} d\mathbf{r} \end{aligned} \quad (6.3)$$

The T_2 correction contains the *von Weizsacker* kinetic energy, τ_{w} , where λ has a value of 1/9. Various empirical values for λ have been used in cases where the expansion is terminated after the T_2 term, motivated by the fact that the von Weizsacker expression is equivalent to the Hartree–Fock kinetic energy for one- and two-electron systems. The kinetic energy at the Thomas–Fermi level is typically underestimated by ~10%, which is reduced to ~1% by addition of the T_2 term, while inclusion of the T_4 correction leads to an overestimation of slightly larger magnitude. In terms of absolute energies, this is comparable to the HF method, but energy differences (e.g. atomization energies) are calculated with much lower accuracy than with the HF model. Unfortunately, the sixth- and higher order T terms diverge in regions far from the nuclei, preventing further improvements.

The second-order exchange term K_2 is given in eq. (6.4) and the K_4 term has an expression similar to T_4 , except that not all the expansion coefficients are known.¹⁰

$$\begin{aligned} K[\rho] &= K_{\text{D}}[\rho] + K_2[\rho] + K_4[\rho] + \dots \\ K_2[\rho] &= -\frac{5}{216} (3\pi^5)^{-1/3} \int \frac{|\nabla \rho(\mathbf{r})|^2}{\rho^{4/3}(\mathbf{r})} d\mathbf{r} \end{aligned} \quad (6.4)$$

Addition of gradient correction terms improves the Thomas–Fermi results, for example bonding is then allowed, but the lack of sufficient error cancellation and the divergence of higher order corrections means that this is not a viable approach for constructing DFT models capable of yielding results comparable with those obtained by wave mechanics methods.

Although there have been some recent attempts at constructing such *orbital-free* (as opposed to the Kohn–Sham version discussed in the next section) T and K function-

als depending directly on the electron density, the accuracy is still too low to be of general use.¹¹ If such functionals could be derived, however, the full potential of DFT in having only three variables independent of system size could be fully realized.

6.2 Kohn-Sham Theory

The foundation for the use of DFT methods in computational chemistry is the introduction of orbitals, as suggested by Kohn and Sham (KS).⁵ The main flaw in orbital-free models is the poor representation of the kinetic energy, and the idea in the KS formalism is to split the kinetic energy functional into two parts, one which can be calculated exactly, and a small correction term. The price to be paid is that orbitals are re-introduced, thereby increasing the complexity from 3 to $3N$ variables, and that electron correlation re-emerges as a separate term. The KS model is closely related to the HF method, sharing identical formulas for the kinetic, electron-nuclear and Coulomb electron-electron energies.

The division of the electron kinetic energy into two parts, with the major contribution being equivalent to the HF kinetic energy, can be justified as follows. Assume for the moment a Hamiltonian operator of the form in eq. (6.5) with $0 \leq \lambda \leq 1$.

$$\mathbf{H}_\lambda = \mathbf{T} + \mathbf{V}_{\text{ext}}(\lambda) + \lambda \mathbf{V}_{\text{ee}} \quad (6.5)$$

The external potential operator \mathbf{V}_{ext} is equal to \mathbf{V}_{ne} for $\lambda = 1$, but for intermediate λ values it is assumed that $\mathbf{V}_{\text{ext}}(\lambda)$ is adjusted such that the *same* density is obtained for $\lambda = 1$ (the real system), for $\lambda = 0$ (a hypothetical system with non-interacting electrons) and for all intermediates λ values. For the $\lambda = 0$ case, the electrons are non-interacting, and the *exact* solution to the Schrödinger equation is given as a Slater determinant composed of (molecular) orbitals, ϕ_i , and the *exact* kinetic energy functional is given in eq. (6.6).

$$T_S = \sum_{i=1}^{N_{\text{elec}}} \langle \phi_i | -\frac{1}{2} \nabla^2 | \phi_i \rangle \quad (6.6)$$

The subscript S denotes that it is the kinetic energy calculated from a Slater determinant. The $\lambda = 1$ case corresponds to *interacting* electrons, and eq. (6.6) is therefore only an approximation to the real kinetic energy, but a substantial improvement over the TF formula (eq. (6.2)).

Another way of justifying the use of eq. (6.6) for calculating the kinetic energy is by reference to natural orbitals (eigenvectors of the density matrix, Section 9.5). The *exact* kinetic energy can be calculated from the natural orbitals (NO) arising from the *exact* density matrix.

$$\begin{aligned} T[\rho_{\text{exact}}] &= \sum_{i=1}^{\infty} n_i \langle \phi_i^{\text{NO}} | -\frac{1}{2} \nabla^2 | \phi_i^{\text{NO}} \rangle \\ \rho_{\text{exact}} &= \sum_{i=1}^{\infty} n_i |\phi_i^{\text{NO}}|^2 \\ N_{\text{elec}} &= \sum_{i=1}^{\infty} n_i \end{aligned} \quad (6.7)$$

The orbital occupation numbers n_i (eigenvalues of the density matrix) will be between 0 and 1, corresponding to the number of electrons in the (spin) orbital. Representing the exact density will require an infinite number of natural orbitals, with the N_{elec} first having occupation numbers close to 1, and the remaining close to 0. Since the exact density matrix is not known, an (approximate) density can be written in terms of a set of auxiliary one-electron functions, i.e. orbitals.

$$\rho_{\text{approx}} = \sum_{i=1}^{N_{\text{elec}}} |\phi_i|^2 \quad (6.8)$$

This corresponds to eq. (6.7) with occupation numbers of exactly 1 or 0. The “missing” kinetic energy from eq. (6.6) is thus due to the occupation numbers deviating from being exactly 1 or 0. Since the occupation numbers of an HF (single-determinant) wave function are also exactly 1 or 0, the missing kinetic energy can also be considered as the (kinetic) correlation energy.

The key to Kohn–Sham theory is to calculate the kinetic energy under the assumption of non-interacting electrons (in the same sense that HF orbitals in wave mechanics describe non-interacting electrons) from eq. (6.6). In reality, the electrons *are* interacting, and eq. (6.6) does not provide the total kinetic energy. However, just as HF theory provides ~99% of the correct answer, the difference between the exact kinetic energy and that calculated by assuming non-interacting orbitals is small. The remaining kinetic energy is absorbed into an exchange–correlation term, and a general DFT energy expression can be written as in eq. (6.9)

$$E_{\text{DFT}}[\rho] = T_{\text{s}}[\rho] + E_{\text{ne}}[\rho] + J[\rho] + E_{\text{xc}}[\rho] \quad (6.9)$$

By equating E_{DFT} to the exact energy, this expression *defines* E_{xc} , i.e. it is the part that remains after subtraction of the non-interacting kinetic energy, and the E_{ne} and J potential energy terms.

$$E_{\text{xc}}[\rho] = (T[\rho] - T_{\text{s}}[\rho]) + (E_{\text{ee}}[\rho] - J[\rho]) \quad (6.10)$$

The first parenthesis in eq. (6.10) may be considered as the kinetic correlation energy, while the last contains both potential correlation and exchange energy.

The task in developing orbital-free models is to derive approximations to the kinetic, exchange and correlation energy functionals, while the corresponding task in Kohn–Sham theory is to derive approximations to the exchange–correlation energy functional only. For the neon atom, for example, the kinetic energy is 128.9 au, the exchange energy is –12.1 au, and the correlation energy is –0.4 au (as calculated by wave mechanics methods). Since the exchange–correlation energy is roughly a factor of 10 smaller than the kinetic energy, Kohn–Sham theory is much less sensitive to inaccuracies in the functional(s) than orbital-free theory. While orbital-free theory is a true density functional theory (three variables), Kohn–Sham methods are independent-particle models ($3N$ variables), analogous to Hartree–Fock theory, but are still much less complicated than many-particle (correlation) wave function models.

6.3 Reduced Density Matrix Methods

Before embarking on a more detailed analysis of how to design exchange–correlation energy functionals in Kohn–Sham theory, it may be instructive to take a slight detour

and consider methods using *reduced density matrices*, rather than the electron density itself. We will start by defining the first- and second-order reduced density matrices γ_1 and γ_2 .

$$\begin{aligned}\gamma_1(\mathbf{r}_1, \mathbf{r}_1') &= N_{\text{elec}} \int \Psi^*(\mathbf{r}_1', \mathbf{r}_2 \dots \mathbf{r}_{N_{\text{elec}}}) \Psi(\mathbf{r}_1, \mathbf{r}_2 \dots \mathbf{r}_{N_{\text{elec}}}) d\mathbf{r}_2 \dots d\mathbf{r}_{N_{\text{elec}}} \\ \gamma_2(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_1', \mathbf{r}_2') &= N_{\text{elec}}(N_{\text{elec}} - 1) \int \Psi^*(\mathbf{r}_1', \mathbf{r}_2', \mathbf{r}_3 \dots \mathbf{r}_{N_{\text{elec}}}) \Psi(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3 \dots \mathbf{r}_{N_{\text{elec}}}) d\mathbf{r}_3 \dots d\mathbf{r}_{N_{\text{elec}}}\end{aligned}\quad (6.11)$$

We will ignore electron spin, except when required. The corresponding reduced spin density matrices are defined completely analogously to eq. (6.11), if the \mathbf{r} variables are taken to represent both spatial and spin coordinates.

The diagonal components of the first-order density matrix (setting $\mathbf{r}_1' = \mathbf{r}_1$) gives the *electron density function* ρ_1 , often written without the subscript 1 when higher order densities are not involved.

$$\rho_1(\mathbf{r}_1) = \gamma_1(\mathbf{r}_1, \mathbf{r}_1) = N_{\text{elec}} \int \Psi^*(\mathbf{r}_1 \dots \mathbf{r}_{N_{\text{elec}}}) \Psi(\mathbf{r}_1 \dots \mathbf{r}_{N_{\text{elec}}}) d\mathbf{r}_2 \dots d\mathbf{r}_{N_{\text{elec}}} \quad (6.12)$$

The integral is the probability of finding an electron (it does not matter which, since they are indistinguishable) at position \mathbf{r}_1 , and the N_{elec} prefactor ensures that the density integrates to the number of electrons.

The corresponding second-order density matrix yields the *electron pair-density* upon setting $\mathbf{r}_1' = \mathbf{r}_1$ and $\mathbf{r}_2' = \mathbf{r}_2$.

$$\begin{aligned}\rho_2(\mathbf{r}_1, \mathbf{r}_2) &= \gamma_2(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_1, \mathbf{r}_2) = \\ &N_{\text{elec}}(N_{\text{elec}} - 1) \int \Psi^*(\mathbf{r}_1, \mathbf{r}_2 \dots \mathbf{r}_{N_{\text{elec}}}) \Psi(\mathbf{r}_1, \mathbf{r}_2 \dots \mathbf{r}_{N_{\text{elec}}}) d\mathbf{r}_3 \dots d\mathbf{r}_{N_{\text{elec}}}\end{aligned}\quad (6.13)$$

The integral is the probability of finding an electron at position \mathbf{r}_1 and another electron at position \mathbf{r}_2 , and the $N_{\text{elec}}(N_{\text{elec}} - 1)$ prefactor ensures that ρ_2 integrates to the number of electron pairs (note that the number of *unique* electron pairs is only $1/2 N_{\text{elec}}(N_{\text{elec}} - 1)$).

The *exact* kinetic and potential energies are given by the (exact) first- and second-order density matrices, with an implicit summation over electron spin.

$$\begin{aligned}T &= -\frac{1}{2} \int \nabla^2 \gamma_1(\mathbf{r}_1, \mathbf{r}_1') \Big|_{\mathbf{r}=\mathbf{r}'} d\mathbf{r}_1 \\ V_{\text{ne}} &= - \sum_a^{N_{\text{nuclei}}} \int \frac{Z_a \gamma_1(\mathbf{r}_1, \mathbf{r}_1)}{|\mathbf{R}_a - \mathbf{r}_1|} d\mathbf{r}_1 = - \sum_a^{N_{\text{nuclei}}} \int \frac{Z_a \rho_1(\mathbf{r}_1)}{|\mathbf{R}_a - \mathbf{r}_1|} d\mathbf{r}_1 \\ V_{\text{ee}} &= \frac{1}{2} \int \frac{\gamma_2(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2 = \frac{1}{2} \int \frac{\rho_2(\mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2\end{aligned}\quad (6.14)$$

Note that the potential energy terms in eq. (6.14) can be written in terms of ρ_1 and ρ_2 ; only the kinetic energy requires γ_1 .

For the special case of a single-determinant wave function, the first- and second-order reduced density matrices are given by eq. (6.15).

$$\Phi_{\text{SD}} = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(1) & \phi_2(1) & \cdots & \phi_N(1) \\ \phi_1(2) & \phi_2(2) & \cdots & \phi_N(2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(N) & \phi_2(N) & \cdots & \phi_N(N) \end{vmatrix} \quad (6.15)$$

$$\gamma_1(\mathbf{r}_1, \mathbf{r}'_1) = \sum_{i=1}^{N_{\text{occ}}} \phi_i(\mathbf{r}'_1) \phi_i(\mathbf{r}_1)$$

$$\gamma_1(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}'_1, \mathbf{r}'_2) = \sum_{i,j=1}^{N_{\text{occ}}} \{ \phi_i(\mathbf{r}'_1) \phi_j(\mathbf{r}'_2) \phi_i(\mathbf{r}_1) \phi_j(\mathbf{r}_2) - \phi_i(\mathbf{r}'_1) \phi_j(\mathbf{r}'_2) \phi_j(\mathbf{r}_1) \phi_i(\mathbf{r}_2) \}$$

Using eq. (6.15) in eq. (6.14) is readily seen to provide the Hartree–Fock energy expression, eq. (3.31).

Since γ_1 can be obtained by straightforward integration of γ_2 , an appealing idea is to use the elements of γ_2 as variables for solving the Schrödinger equation or its equivalent density formulation by a variational procedure. Unfortunately, the γ_2 elements cannot be varied freely, since they must correspond to an antisymmetric wave function.¹² Although a formal solution of this *N-representability* problem exists,¹³ it does not lend itself to an efficient computational implementation. Recent work by D. Mazziotti, however, has shown that good approximations to the *N-representability* can be obtained by enforcing positive semi-definiteness (non-negative eigenvalues) of three matrices during the optimization of γ_2 .¹⁴ These three matrices are the two-particle density matrix γ_2 itself and the corresponding representations in terms of hole–hole and particle–hole creations, commonly denoted D, Q and G. The semi-definite condition arises since these matrices describe probabilities. Performing such constrained optimizations is a non-trivial computational task, but recently a method has been proposed that (only) scales as M_{basis}^6 , making such calculation tractable for general systems.¹⁵ The accuracy of the results can be improved by imposing semipositivity of higher order density matrices, albeit at a significantly higher computational cost. It is at present unclear exactly what the limitations of these methods are in terms of accuracy and computational costs.

The electron density matrix ρ_1 can be diagonalized to produce eigenvalues and eigenvectors, called occupation numbers n_i and natural orbitals ϕ_i^{NO} . The γ_1 can be written in term of these quantities (compare with eq. (6.15), but note that the summation now includes all orbitals, since $n_i \neq 0$ in general).

$$\gamma_1(\mathbf{r}_1, \mathbf{r}'_1) = \sum_i^{N_{\text{orb}}} n_i \phi_i^{\text{NO}}(\mathbf{r}'_1) \phi_i^{\text{NO}}(\mathbf{r}_1) \quad (6.16)$$

As an alternative to using γ_2 as the fundamental variable, the first-order reduced density matrix, or its parameterization in terms of natural orbitals and occupation numbers, can be used. The *N-representability* problem for γ_1 is easy to fulfil, since the only requirements are that the eigenvalues are between 0 and 1, and that they sum to N_{elec} . The task in this case is to derive a suitable approximation for the V_{ee} energy term in terms of γ_1 , rather than ρ_2 .¹⁶ The Coulomb and exchange parts of V_{ee} are given by the analogous formula (eqs (6.14) and (6.15)) using natural orbitals and occupation numbers, leaving the correlation energy as the only unknown function of γ_1 . The cor-

relation part can be incorporated by multiplying the orbital products with functions of the occupation numbers as shown in eq. (6.17).

$$V_{ee} = \frac{1}{2} \int \frac{\rho_2(\mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2$$

$$\rho_2(\mathbf{r}_1, \mathbf{r}_2) = \sum_{ij}^{N_{orb}} \left\{ \begin{aligned} & f(n_i, n_j) \phi_i^{NO}(\mathbf{r}_1) \phi_i^{NO}(\mathbf{r}_1) \phi_j^{NO}(\mathbf{r}_2) \phi_j^{NO}(\mathbf{r}_2) - \\ & g(n_i, n_j) \phi_i^{NO}(\mathbf{r}_1) \phi_j^{NO}(\mathbf{r}_1) \phi_i^{NO}(\mathbf{r}_2) \phi_j^{NO}(\mathbf{r}_2) \end{aligned} \right\} \quad (6.17)$$

The choice of $f(n_i, n_j) = g(n_i, n_j) = n_i n_j$ implies the Hartree–Fock model, and optimization of the occupation numbers and natural orbitals with this choice indeed returns the HF wave function, i.e. the HF wave function cannot be improved by allowing fractional occupation numbers. The $f(n_i, n_j)$ function is usually set equal to $n_i n_j$, since this is just the Coulomb interaction, and the exchange–correlation part is modelled by $g(n_i, n_j)$.¹⁷ Modelling the exchange–correlation energy by a $g(n_i, n_j)$ function is still in its infancy.

The Hohenberg–Kohn theorem, which states that the energy is uniquely determined by the one-electron density ρ_1 , forms the basis for what is commonly called density functional theory. As discussed in Section 6.1, it is difficult to construct a sufficiently accurate total energy functional depending only on ρ_1 . The Kohn–Sham version, where only the exchange–correlation part of the energy must be estimated as a functional of ρ_1 , provides viable models, and these will be discussed in more detail in Section 6.5.

Perhaps the most surprising result of the Hohenberg–Kohn theorem is that the correlation energy is completely determined by the one-electron density function ρ_1 . Electron correlation is inherently a two-electron phenomenon, and it is difficult to envision how an accurate correlation functional depending on only the one-electron density should be constructed from theoretical arguments, although one certainly can understand that the correlation will affect the electron density. Indeed, the interpretation of electron correlation in terms of correlation holes, as discussed in the next section, requires the two-electron density ρ_2 . P. M. W. Gill has suggested that one could consider a quantity similar to the second-order reduced density matrix, except that the arguments are the position and momentum of the two electrons.¹⁸

$$W(\mathbf{r}_1, \dots, \mathbf{r}_{N_{elec}}, \mathbf{p}_1, \dots, \mathbf{p}_{N_{elec}}) = \pi^{-3N_{elec}} \times$$

$$\int \Psi^*(\mathbf{r}_1 + \mathbf{q}_1, \dots, \mathbf{r}_{N_{elec}} + \mathbf{q}_{N_{elec}}) \Psi(\mathbf{r}_1 - \mathbf{q}_1, \dots, \mathbf{r}_{N_{elec}} - \mathbf{q}_{N_{elec}}) e^{2i(\mathbf{p}_1 \cdot \mathbf{q}_1 + \dots + \mathbf{p}_{N_{elec}} \cdot \mathbf{q}_{N_{elec}})} d\mathbf{q}_1 \dots d\mathbf{q}_{N_{elec}}$$

$$W_2(\mathbf{r}_1, \mathbf{r}_2, \mathbf{p}_1, \mathbf{p}_2) = \int W(\mathbf{r}_1, \dots, \mathbf{r}_{N_{elec}}, \mathbf{p}_1, \dots, \mathbf{p}_{N_{elec}}) d\mathbf{r}_3 \dots d\mathbf{r}_{N_{elec}} d\mathbf{p}_3 \dots d\mathbf{p}_{N_{elec}} \quad (6.18)$$

The W_2 is called the second-order *Wigner intracule* and represents a quasi-probability function for finding two electrons at positions \mathbf{r}_1 and \mathbf{r}_2 with momentum \mathbf{p}_1 and \mathbf{p}_2 . It cannot be interpreted as a genuine probability function, as it can achieve negative values. Since one would expect the correlation energy to depend on the distance between the two electrons, the relative momentum between them and the orientation of the distance and momentum difference vectors, Gill has suggested that the corresponding Ω intracule depending on the “internal” coordinates may contain information sufficient for determining the correlation energy.

$$\begin{aligned}
W_2(\mathbf{r}_1, \mathbf{r}_2, \mathbf{p}_1, \mathbf{p}_2) &\rightarrow \Omega(u, v, \omega) \\
u = |\mathbf{r}_1 - \mathbf{r}_2|; \quad v = |\mathbf{p}_1 - \mathbf{p}_2|; \quad \cos \omega &= \frac{(\mathbf{r}_1 - \mathbf{r}_2) \cdot (\mathbf{p}_1 - \mathbf{p}_2)}{|\mathbf{r}_1 - \mathbf{r}_2| |\mathbf{p}_1 - \mathbf{p}_2|} \\
\Delta E_{\text{corr}} &= \int \Omega(u, v, \omega) G(u, v, \omega) du dv d\omega
\end{aligned} \tag{6.19}$$

Analogously to density functional methods, the task is to construct a correlation kernel G that yields the correlation energy upon integration with the Ω intracule, and use this in connection with the exchange energy calculated at the Hartree–Fock level. Given that the intracule is based on physical principles underlying the correlation phenomenon, it may be possible to use theoretical arguments for deriving useful approximations to the correlation kernel. Preliminary results for atoms using simple correlation kernels are sufficiently accurate that this may represent a viable approach.

The differences between wave mechanics and density-based methods can be summarized as follows:

- Wave mechanics employs the exact Hamiltonian operator, but makes approximations in the form of the wave function.
- Density functional methods make approximations in the energy functional (Hamiltonian), but allow a free variation of the electron density ρ_1 . The functional must therefore *implicitly* enforce the N -representability. This is difficult to achieve in orbital-free methods, but the Kohn–Sham approach with a determinantal orbital product takes care of the majority of this problem. Furthermore, in orbital-free methods the kinetic energy functional is unknown, and since this is equal in magnitude to the total energy, even minor inaccuracies cause large errors. In the Kohn–Sham version only the exchange–correlation functional is unknown, and since this is a relatively minor component of the total energy, the results are less sensitive to inaccuracies in the functional.
- Methods using the first-order reduced density matrix as variable can be chosen to strictly enforce the N -representability of γ_1 , and employ the exact energy functional for all the terms except the correlation energy. The latter, however, inherently depends on γ_2 , which in this approach must be approximated as a function of γ_1 . One can argue that Hartree–Fock belongs to this class of methods, with implicit neglect of the electron correlation.
- Methods using the second-order reduced density matrix as variable employ the exact energy functional in terms of γ_2 , but must make approximations for enforcing the N -representability of γ_2 .

Solving the Schrödinger equation by means of reduced density matrices has many appealing features, such as being able to describe the whole potential energy curve with equal accuracy and to account for a very large fraction of the correlation energy. The methods using the second-order reduced density matrix look especially interesting, although so far results have only been reported for small systems.

6.4 Exchange and Correlation Holes

We now return to the problem of expressing the exchange–correlation energy as a functional of ρ ($=\rho_1$). Since the exchange energy is by far the largest contributor to E_{xc}

(cf. the values for the neon atom in Section 6.2), one may reasonably ask why not calculate this term “exactly” from orbitals (analogous to the kinetic energy), by the formula known from wave mechanics (eq. (3.31)), and only calculate the computationally difficult part, the correlation energy, by DFT. Although this has been tried, it gives poor results. The basic problem is that the DFT definitions of exchange and correlation energies are not completely equivalent to their wave mechanics counterparts.¹⁹ The DFT exchange energy may be defined by the same formula as in HF theory (eq. (3.30)), except that Kohn–Sham orbitals are used. This leads to a non-local potential, i.e. the exchange potential at a given point is strongly dependent on the density at distant points. The correlation energy in wave mechanics is defined as the difference between the exact energy and the corresponding Hartree–Fock value. Both the exchange and correlation energies have a short- and long-range part (in terms of the distance between two electrons). The long-range correlation is essentially the “static” correlation energy (i.e. the “multi-reference” part, see Section 4.6) while the short-range part is the “dynamical” correlation. The long-range part of the correlation energy in wave mechanics effectively cancels the delocalized part of the exchange energy. The definitions of exchange and correlation in DFT (at least in current implementations) are local (short range), since they only depend on the density at a given point and the immediate vicinity (via derivatives of the density). The cancellation at long range is (or should be) implicitly built into the exchange–correlation functional. Calculating the exchange energy by wave mechanics and the correlation by DFT thus destroys the cancellation, although recent work has attempted to address this problem by separating the correlation functional into a long- and short-range part, and use this in connection with HF exchange.²⁰

A more detailed discussion of these features is most easily given in terms of *exchange and correlation holes*. Electrons avoid each other owing to their electric charges, and the energy associated with this repulsion is given classically by the Coulomb equation (eq. 6.1). Quantum mechanically, however, this repulsion must be modified to take into account that electrons have spins of $1/2$. The Pauli principle states that two fermions (particles with half-integer spin) cannot occupy the same spatial position, or equivalently, that the total wave function must be antisymmetric upon interchange of any two particles. This leads to the exchange energy (see Section 3.3), which can be considered as a quantum correction to the classical Coulomb repulsion. The exchange term is already present in Hartree–Fock theory, and must also be incorporated into DFT. In addition, there is a dynamical effect where electrons tend to avoid each other more than given by an HF wave function, and this is the correlation energy calculated by wave mechanics methods.

These qualitative considerations can be put into quantitative terms by probability *holes*. If electrons did not have charge or spin, the probability of finding an electron at a given position would be independent of the position of a second electron, and the electron pair-density ρ_2 would be given as a simple product of two one-electron densities ρ_1 , with a proper normalization factor.

$$\rho_2^{\text{indep}}(\mathbf{r}_1, \mathbf{r}_2) = \frac{N_{\text{elec}} - 1}{N_{\text{elec}}} \rho_1(\mathbf{r}_1) \rho_1(\mathbf{r}_2) = \left(1 - \frac{1}{N_{\text{elec}}}\right) \rho_1(\mathbf{r}_1) \rho_1(\mathbf{r}_2) \quad (6.20)$$

Since electrons have both charge and spin, however, there is a reduced probability of finding an electron near another electron. We can write this formally in terms of a con-

ditional probability factor $h_{xc}(\mathbf{r}_1, \mathbf{r}_2)$ that includes the $1/N_{\text{elec}}$ self-interaction factor in eq. (6.20).

$$\rho_2(\mathbf{r}_1, \mathbf{r}_2) = \rho_1(\mathbf{r}_1)\rho_1(\mathbf{r}_2) + \rho_1(\mathbf{r}_1)h_{xc}(\mathbf{r}_1, \mathbf{r}_2) \quad (6.21)$$

The reduced probability is called the *exchange–correlation hole*, and can be written in terms of ρ_2 and ρ_1 by solving eq. (6.21).

$$h_{xc}(\mathbf{r}_1, \mathbf{r}_2) = \frac{\rho_2(\mathbf{r}_1, \mathbf{r}_2)}{\rho_1(\mathbf{r}_1)} - \rho_1(\mathbf{r}_2) \quad (6.22)$$

The exchange–correlation hole represents the reduced probability of finding electron 2 at a position \mathbf{r}_2 given that electron 1 is located at \mathbf{r}_1 . The exchange part of h_{xc} is called the *Fermi hole*, while the dynamical correlation gives rise to the *Coulomb hole*. Since exchange only occurs between electrons of the same spin, the total hole can also be written in terms of individual spin contributions.

$$\begin{aligned} h_{xc} &= h_x + h_c \\ h_x &= h_x^{\alpha\alpha} + h_x^{\beta\beta} \\ h_c &= h_c^{\alpha\alpha} + h_c^{\beta\beta} + h_c^{\alpha\beta} \end{aligned} \quad (6.23)$$

From the definitions of ρ_2 and ρ_1 , it follows that the integral of h_{xc} over \mathbf{r}_2 equals -1 .

$$\int h_{xc}(\mathbf{r}_1, \mathbf{r}_2) d\mathbf{r}_2 = \int \frac{\rho_2(\mathbf{r}_1, \mathbf{r}_2)}{\rho_1(\mathbf{r}_1)} d\mathbf{r}_2 - \int \rho_1(\mathbf{r}_2) d\mathbf{r}_2 = \frac{N_{\text{elec}}(N_{\text{elec}} - 1)}{N_{\text{elec}}} - N_{\text{elec}} = -1 \quad (6.24)$$

A similar argument for the separate spin densities shows that the Fermi hole itself is negative everywhere and integrates to -1 , which means that the integral of the Coulomb hole is 0. The Fermi (exchange) hole describes a static reduction in the probability function corresponding to one electron. The Coulomb (correlation) function, on the other hand, reduces the probability of finding an electron near the reference electron, but increases the probability of finding it far from the reference electron.

The exchange energy in Hartree–Fock theory is a non-local function, i.e. the HF exchange hole is delocalized over the whole system (or at least a large part of it). For a diatomic system, for example, the exchange hole is delocalized over both nuclei. When electron correlation is added explicitly, the left–right correlation to a large extent serves to cancel the delocalized nature of the HF exchange hole. The exchange functional in DFT, on the other hand, is local, i.e. the cancellation of the delocalized HF exchange hole by the left–right correlation in wave function approaches should be inherent in the functional, and this is the main difference between the definitions of exchange and correlation in wave function and current density functional descriptions.

A closely related phenomenon is the electron self-interaction energy. The Coulomb energy functional given in eq. (6.1) only depends on ρ_1 . This means that the density arising from a single electron will interact with itself (e.g. there will be a non-zero electron–electron Coulomb repulsion even for a one-electron system), and this self-repulsion is clearly non-physical. For a multi-electron system, there will be such a self-interaction term for the density associated with each electron, although it is difficult to define this rigorously. The HF model takes care of this problem elegantly, since

the expression for the exchange energy exactly cancels the Coulomb self-interaction (eq. (3.32)).

Within the DFT model, a one-electron system should have an exchange energy exactly opposing the Coulomb energy, $E_x = -J$, and the correlation energy should be zero. In a multi-electron system, one may thus consider part of the exchange energy as a correction for the self-interaction energy, with the remaining being the “true” Fermi hole. In this view, the self-interaction correction is the major part of the exchange energy, with the “true” Fermi hole being comparable to the Coulomb (correlation) hole. It should be noted, however, that such a partitioning is not invariant to a unitary transformation of the occupied orbitals.²¹ The self-interaction cancellation by the exchange energy is not guaranteed in DFT, and very few of the current exchange–correlation functionals are completely self-interaction-free. It has been proven that a completely self-interaction-free *local* potential does not exist.²² Perdew and Zunger have suggested an approximate correction scheme, where each orbital becomes self-interaction-free.²³ The procedure formally changes the underlying functional and destroys the invariance of the energy with respect to mixing of the occupied orbitals. Since it is furthermore computationally quite expensive and often degrades the quality of the results by overcorrecting the error, it has seen little use.²⁴ Other versions have been proposed, each having different computational and theoretical disadvantages.²⁵

These concepts can be illustrated for the H_2 molecule for increasing internuclear distances.²⁶ In wave mechanics, the ground state for H_2 has two electrons of opposite spin in the same spatial orbital, and the exchange hole is thus entirely the self-interaction correction (no same-spin exchange).

$$\begin{aligned} h_x &= -\frac{1}{2} \rho_1 \\ \rho_1 &= 2\phi^2 \\ \phi &= N(\chi_A + \chi_B) \end{aligned} \tag{6.25}$$

For the case of H_2 , this is just the negative of the occupied molecular orbital and, by symmetry arguments, the HF exchange hole at each nucleus thus integrates to $-1/2$, independent of the internuclear separation. This delocalization is clearly non-physical in the dissociation limit, since the correct wave function must have one electron localized at each nucleus. This means that for a reference electron near nucleus A, the total probability hole is localized at nucleus A. The correlation hole must therefore exactly cancel the exchange hole at nucleus B, while increasing the hole at nucleus A in order to integrate to -1 (Figure 6.1), and this is the reason why the wave function correlation energy increases as a function of internuclear distance.

Within wave mechanics, the exchange hole is static and delocalized over the whole molecule, while the long-range part of the electron correlation is dynamical and serves to cancel the exchange hole away from the reference electron.

6.5 Exchange–Correlation Functionals

The difference between various DFT methods is the choice of functional form for the exchange–correlation energy. It can be proven that the exchange–correlation potential is a unique functional, valid for all systems, but an explicit functional form of this potential has been elusive, except for special cases such as a uniform electron gas. It is




Figure 6.1 Illustrating the exchange and correlation holes for the H_2 molecule at the dissociation limit, with the reference electron located near nucleus A and the vertical axis representing probability

possible, however, to derive a number of properties that the exact functional should have, of which some of the more important ones are:²⁷

- (1) The energy functional should be self-interaction-free, i.e. the exchange energy for a one-electron system, such as the hydrogen atom, should exactly cancel the Coulomb energy, and the correlation energy should be zero. Although these seem like obvious requirements, none of the common functionals have this property.
- (2) When the density becomes constant, the uniform electron gas result should be recovered. While this surely is a valid mathematical requirement, and important for applications in solid-state physics, it may not be as important for chemical applications, as molecular densities are relatively poorly described by uniform electron gas methods.
- (3) The coordinate scaling of the exchange energy should be linear, i.e. multiplying the electron coordinates with a constant factor should result in a similar linear scaling of the exchange energy.²⁸

$$\begin{aligned}\rho_\lambda(x, y, z) &= \lambda^3 \rho(\lambda x, \lambda y, \lambda z) \\ E_x[\rho_\lambda] &= \lambda E_x[\rho]\end{aligned}\tag{6.26}$$

- (4) No direct scaling law applies for the correlation energy, but scaling the electron coordinates by a factor larger than 1 should increase the magnitude of the correlation (and vice versa).²⁸ In the low density limit, the scaling becomes linear, as for the exchange energy.

$$-E_c[\rho_\lambda] > -\lambda E_c[\rho]; \quad \lambda > 1\tag{6.27}$$

- (5) As the scaling parameter goes to infinity, the correlation energy for a finite system approaches a negative constant.

- (6) The *Lieb–Oxford* condition places an upper bound for the exchange–correlation energy relative to the *Local Density Approximation* (LDA) (see Section 6.5.1) exchange energy.²⁹

$$E_x[\rho] \geq E_{xc}[\rho] \geq 2.273 E_x^{\text{LDA}}[\rho] \quad (6.28)$$

- (7) The exchange potential should show an asymptotic $-r^{-1}$ behaviour as $r \rightarrow \infty$.³⁰ Furthermore, the exchange–correlation potential is discontinuous as a function of the number of electrons, by an amount corresponding to the difference between the ionization potential and electron affinity.³¹
- (8) The correlation potential should show an asymptotic $-1/2\alpha r^{-4}$ behaviour, with α being the polarizability of the $N_{\text{elec}} - 1$ system.

The difference in scaling behaviour (points 3 and 4) is a strong argument for separating the corresponding exchange and correlation functionals but, on the other hand, it implies a difficult task for getting the correlation component to exactly cancel the long-range exchange component.

Exchange–correlation functionals have, in analogy with other (partly) empirical methods, a mathematical form containing parameters. There are two main philosophies for assigning values to these parameters, either by requiring the functional to fulfil the above criteria (or a suitable selection thereof), or by fitting the parameters to experimental data, although in practice a combination of these approaches is often used. The quality of exchange–correlation functionals will ultimately have to be settled by comparing the performance with experiments or high-level wave mechanics calculations. Such calibration studies, however, only evaluate the quality for the chosen selection of systems and properties. It has indeed been found that the “best” functionals depend on the system and properties, some being good for molecular systems, others for delocalized (periodic) systems, and others again for properties such as excitation energies or NMR chemical shifts. At present, there are no clear “standard” methods, like MP2 and CCSD in traditional *ab initio* theory, although the hybrid methods discussed below usually give good performance. Since DFT is an active area of research, new and improved functionals are likely to emerge. It should be noted that many of the proposed functionals have never made it past the research stage and are not available in commonly available programs. Below we will give a short summary of functionals that have been proposed by different research groups.³² We will give the explicit forms for some of the more commonly used functionals for illustration, although they do not contain much physical insight by themselves.

It is customary to separate E_{xc} into two parts, a pure exchange E_x and a correlation part E_c , which seems reasonable based on the discussion of the exchange and correlation holes above, and their different scaling properties. It should be noted, however, that only the combined exchange–correlation hole has a physical meaning, and it could be argued that this calls for a combined E_{xc} . Early work tended to focus on only one of the components, and subsequently combined these, while the current trend is to construct the two parts in a combined fashion.

Each of the exchange and correlation energies is often written in term of the energy per particle (energy density), ϵ_x and ϵ_c .

$$E_{xc}[\rho] = E_x[\rho] + E_c[\rho] = \int \rho(r) \epsilon_x[\rho(r)] dr + \int \rho(r) \epsilon_c[\rho(r)] dr \quad (6.29)$$

As mentioned at the start of Chapter 4, the correlation between electrons of parallel spin is different from that between electrons of opposite spin. The exchange energy is “by definition” given as a sum of contributions from the α and β spin densities, as exchange energy only involves electrons of same spin. The kinetic energy, the nuclear–electron attraction and Coulomb terms are trivially separable in terms of electron spin.

$$\begin{aligned} E_x[\rho] &= E_x^\alpha[\rho_\alpha] + E_x^\beta[\rho_\beta] \\ E_c[\rho] &= E_c^{\alpha\alpha}[\rho_\alpha] + E_c^{\beta\beta}[\rho_\beta] + E_c^{\alpha\beta}[\rho_\alpha, \rho_\beta] \end{aligned} \quad (6.30)$$

The total density is the sum of the α and β contributions, $\rho = \rho_\alpha + \rho_\beta$, and these are identical ($\rho_\alpha = \rho_\beta$) for a closed shell singlet. Functionals for the exchange and correlation energies may be formulated in terms of separate spin densities, however, they are often given instead as functions of the spin polarization ζ (normalized difference between ρ_α and ρ_β), and the radius of the effective volume containing one electron, r_s ,

$$\begin{aligned} \zeta &= \frac{\rho_\alpha - \rho_\beta}{\rho_\alpha + \rho_\beta} \\ \frac{4}{3}\pi r_s^3 &= \rho^{-1} \end{aligned} \quad (6.31)$$

In the formulas below it is implicitly assumed that the exchange and correlation energies are summed over both α and β densities.

The difference between various wave function-based methods is how the electron correlation is included, and the quality of these methods can be characterized by an ordering parameter, such as the perturbation order or the level of excitations included. There are no similar theoretically founded ordering parameters for DFT methods, as the exchange–correlation functionals to a large extent are empirical. A heuristic characterization can be done by considering the fundamental variables used for defining the exchange–correlation functional. J. P. Perdew has suggested such a “Jacob’s ladder” approach, where one can expect or at least hope for an improvement in the accuracy for each step up the ladder,³³ and this is the approach taken here to systematize the plethora of functionals that has been proposed.

6.5.1 Local Density Approximation

In the *Local Density Approximation* (LDA) it is assumed that the density locally can be treated as a uniform electron gas, or equivalently that the density is a slowly varying function. The exchange energy for a uniform electron gas is given by the Dirac formula (eq. (6.2)).

$$\begin{aligned} E_x^{\text{LDA}}[\rho] &= -C_x \int \rho^{4/3}(\mathbf{r}) d\mathbf{r} \\ \epsilon_x^{\text{LDA}} &= -C_x \rho^{1/3} \end{aligned} \quad (6.32)$$

In the more general case, where the α and β densities are not equal, LDA (where the *sum* of the α and β densities is raised to the $\frac{4}{3}$ power) has been virtually abandoned and replaced by the *Local Spin Density Approximation* (LSDA) (which is given as the sum of the *individual* densities raised to the $\frac{4}{3}$ power, eq. (6.33)).

$$E_x^{\text{LSDA}}[\rho] = -2^{1/3} C_x \int (\rho_\alpha^{4/3} + \rho_\beta^{4/3}) d\mathbf{r} \quad (6.33)$$

LSDA may also be written in terms of the total density and a spin-polarization function.

$$\begin{aligned} \epsilon_x^{\text{LSDA}} &= -C_x f_1(\zeta) \rho^{1/3} \\ f_1(\zeta) &= \frac{1}{2} \left[(1+\zeta)^{4/3} + (1-\zeta)^{4/3} \right] \end{aligned} \quad (6.34)$$

For closed shell systems, LSDA is equal to LDA and, since this is the most common case, LDA is often used interchangeably with LSDA, although this is not true in the general case. The X_α method proposed by Slater in 1951³⁴ can be considered as an LDA method where the correlation energy is neglected and the exchange term is as given in eq. (6.35).

$$\epsilon_{X_\alpha} = -\frac{3}{2} \alpha C_x \rho^{1/3} \quad (6.35)$$

With $\alpha = 2/3$ this is identical to the Dirac expression. The original X_α method used $\alpha = 1$, but a value of $3/4$ has been shown to give better agreement for atomic and molecular systems. The name Slater is often used as a synonym for the L(S)DA exchange energy involving the electron density raised to the $4/3$ power.

The analytical form for the correlation energy of a uniform electron gas, which is purely dynamical correlation, has been derived in the high and low density limits.³⁵ For intermediate densities, the correlation energy has been determined to a high precision by quantum Monte Carlo methods (Section 4.16). In order to use these results in DFT calculations, it is desirable to have a suitable analytic interpolation formula, and such formulas have been constructed by *Vosko, Wilk and Nusair* (VWN) and by *Perdew and Wang* (PW), and are considered to be accurate fits.³⁶ The VWN parameterization is given in eq. (6.36), where a slightly different spin-polarization function has been used.

$$\begin{aligned} \epsilon_c^{\text{VWN}}(r_s, \zeta) &= \epsilon_c(r_s, 0) + \epsilon_a(r_s) \left[\frac{f_2(\zeta)}{f_2''(0)} \right] \left(1 - \zeta^4 \right) + [\epsilon_c(r_s, 1) - \epsilon_c(r_s, 0)] f_2(\zeta) \zeta^4 \\ f_2(\zeta) &= \frac{(f_1(\zeta) - 2)}{(2^{1/3} - 1)} \end{aligned} \quad (6.36)$$

The $\epsilon_c(r_s, \zeta)$ and $\epsilon_a(r_s)$ functions are parameterized as in eq. (6.37), with A , x_0 , b and c being suitable fitting constants. Several slightly different parameterizations were proposed in the original paper, which has caused some confusion, since different implementations have used different parameterizations, and therefore produce slightly different numerical results.

$$\begin{aligned} \epsilon_{c/a}(x) &= A \left\{ \ln \frac{x^2}{X(x)} + \frac{2b}{Q} \tan^{-1} \left(\frac{Q}{2x+b} \right) - \right. \\ &\quad \left. \left[\frac{bx_0}{X(x_0)} \left[\ln \frac{(x-x_0)^2}{X(x)} + \frac{2(b+2x_0)}{Q} \tan^{-1} \left(\frac{Q}{2x+b} \right) \right] \right] \right\} \\ x &= \sqrt{r_s} \\ X(x) &= x^2 + bx + c \\ Q &= \sqrt{4c - b^2} \end{aligned} \quad (6.37)$$

The PW parameterization for $\varepsilon_{c/a}$ is given in eq. (6.38), with a , α , β_1 , β_2 , β_3 and β_4 again being fitting parameters.

$$\varepsilon_{c/a}^{\text{PW}}(x) = -2ap(1 + \alpha x^2) \ln \left(1 + \frac{1}{2a(\beta_1 x + \beta_2 x^2 + \beta_3 x^3 + \beta_4 x^4)} \right) \quad (6.38)$$

The LSDA method is an *exact* DFT method for the special case of a uniform electron gas, except for small differences depending on the interpolation formula chosen for the correlation energy. For molecular systems, the LSDA approximation underestimates the exchange energy by $\sim 10\%$, thereby creating errors that are larger than the whole correlation energy. Electron correlation is overestimated, often by a factor close to 2, and bond strengths are as a consequence overestimated, often by ~ 100 kJ/mol. Despite the simplicity in the fundamental assumptions, LSDA methods are often found to provide results with an accuracy similar to that obtained by wave mechanics Hartree–Fock methods. It has furthermore been used extensively in the physics community for describing extended systems, such as metals, where the approximation of a slowly varying electron density is quite valid.

6.5.2 Gradient-corrected methods

Improvements over the LSDA approach must consider a non-uniform electron gas. A step in this direction is to make the exchange and correlation energies dependent not only on the electron density but also on derivatives of the density. The first-order correction for the exchange energy is given in eq. (6.4), and the corresponding quantity for the correlation energy is also known.³⁷ While inclusion of the first-order exchange term improves the exchange energy, inclusion of the first-order correlation correction often makes the correlation energy *positive*. A straightforward inclusion of these first-order terms leads to a model that performs worse than the simple LSDA model. The main reason for the success of the LSDA approach is that it fulfils the requirements of the Fermi hole integrating to -1 , and the Coulomb hole to 0, while the addition of gradient terms destroys these important properties.

In *Generalized Gradient Approximation* (GGA) methods, the first derivative of the density is included as a variable, and in addition it is required that the Fermi and Coulomb holes integrate to the required values of -1 and 0. GGA methods are also sometimes referred to as *non-local* methods, although this is somewhat misleading since the functionals only depend on the density (and derivative) at a given point, not on a space volume as the Hartree–Fock exchange energy.

One of the earliest and most popular GGA exchange functionals was proposed by A. D. Becke (B or B88) as a correction to the LSDA exchange energy.³⁸

$$\begin{aligned} \varepsilon_x^{\text{B88}} &= \varepsilon_x^{\text{LDA}} + \Delta \varepsilon_x^{\text{B88}} \\ \Delta \varepsilon_x^{\text{B88}} &= -\beta p^{1/3} \frac{x^2}{1 + 6\beta x \sinh^{-1} x} \\ x &= \frac{|\nabla \rho|}{\rho^{4/3}} \end{aligned} \quad (6.39)$$

The β parameter is determined by fitting to known data for the rare gas atoms using the dimensionless gradient variable x . The B88 exchange functional has the correct

asymptotic behaviour for the energy density (but not for the exchange potential).³⁹ It reduces the error in the exchange energy by almost two orders of magnitude relative to the LSDA result, and thus represents a substantial improvement for a simple functional form containing only one adjustable parameter.

Handy and Cohen have investigated several forms related to eq. (6.39) where the parameters were optimized with respect to exchange energies calculated at the HF level. The best resulting model had two parameters and was labelled *OPTX* (*OPTimized eXchange*).⁴⁰ It was also found that no significant improvement could be made by including higher order derivatives (discussed in the next section). *Hamprecht, Cohen, Tozer and Handy* have further extended the B97 model discussed in the next section, but using only the pure density components (i.e. no exact exchange) to a functional containing 15 parameters which were fitted to experimental and *ab initio* data, giving the acronyms *HCTH93*, *HCTH147* and *HCTH407*, where the number referring to the number of molecules in the fitting data set.⁴¹

There have similarly been various GGA functionals proposed for the correlation energy. One popular functional is due to *Lee, Yang and Parr* (LYP),⁴² which has the rather intimidating form shown in eq. (6.40).

$$\begin{aligned} \epsilon_c^{\text{LYP}} = & -4a \frac{\rho_\alpha \rho_\beta}{\rho^2(1+d\rho^{-1/3})} - \\ & ab\omega \left\{ \frac{\rho_\alpha \rho_\beta}{18} \left[144(2^{2/3})C_F(\rho_\alpha^{8/3} + \rho_\beta^{8/3}) + (47-7\delta)|\nabla\rho|^2 - \right. \right. \\ & \left. \left[(45-\delta)(|\nabla\rho_\sigma|^2 + |\nabla\rho_\beta|^2) + 2\rho^{-1}(11-\delta)(\rho_\sigma|\nabla\rho_\sigma|^2 + \rho_\beta|\nabla\rho_\beta|^2) \right] \right\} \\ & \left. + \frac{2}{3}\rho^2(|\nabla\rho_\sigma|^2 + |\nabla\rho_\beta|^2 - |\nabla\rho|^2) - (\rho_\alpha^2|\nabla\rho_\beta|^2 + \rho_\beta^2|\nabla\rho_\alpha|^2) \right\} \\ & e^{-c\rho^{-1/3}} \\ \omega = & \frac{\rho^{14/3}}{\rho^{14/3}(1+d\rho^{-1/3})} \\ \delta = & c\rho^{-1/3} + \frac{d\rho^{-1/3}}{(1+d\rho^{-1/3})} \end{aligned} \quad (6.40)$$

The a , b , c and d parameters are determined by fitting to data for the helium atom. Although not obvious from the form shown in eq. (6.40), the LYP functional does not include parallel spin correlation when all the spins are aligned (e.g. the LYP correlation energy for ^3He is zero). The LYP correlation functional is often combined with the B88 or OPTX exchange functional to produce the BLYP and OLYP acronyms.

J. P. Perdew and coworkers have proposed several related exchange-correlation functionals based on removing spurious oscillations in the Taylor-like expansion to first order and ensuring that the exchange and correlation holes integrate to the required values of -1 and 0 . The associated acronyms are *PW86* (*Perdew-Wang 1986*),⁴³ *PW91* (*Perdew-Wang 1991*)⁴⁴ and *PBE* (*Perdew-Burke-Ernzerhof*).⁴⁵ These three functionals should be considered as refinements of the same underlying model, i.e. the PBE version should be used in favour of the PW86 and PW91 versions. The exchange part is written as an enhancement factor multiplied onto the LSDA functional, where the dimensionless gradient variable x is defined in eq. (6.39).

$$\begin{aligned}\epsilon_x^{\text{PBE}} &= \epsilon_x^{\text{LDA}} F(x) \\ F(x) &= 1 + a - \frac{a}{1 + bx^2}\end{aligned}\tag{6.41}$$

The correlation part is similarly written as an enhancement factor added to the LSDA functional, where the t variable is related to the x variable by means of yet another spin-polarization function.

$$\begin{aligned}\epsilon_c^{\text{PBE}} &= \epsilon_c^{\text{LDA}} + H(t) \\ H(t) &= cf_3^3 \ln \left[1 + dt^2 \left(\frac{1 + At^2}{1 + At^2 + A^2 t^4} \right) \right] \\ A &= d \left[\exp \left(-\frac{\epsilon_c^{\text{LDA}}}{cf_3^3} \right) - 1 \right]^{-1} \\ f_3(\zeta) &= \frac{1}{2} \left[(1 + \zeta)^{2/3} + (1 - \zeta)^{2/3} \right] \\ t &= \left[2(3\pi^3)^{1/3} f_3 \right]^{-1} x\end{aligned}\tag{6.42}$$

The a , b , c and d parameters in these functionals are non-empirical, i.e. they are not obtained by fitting to experimental data, but derived from some of the conditions in Section 6.5. The PW91 functional has been tuned to improve the performance for weak interactions, producing the *mPW91* acronym.⁴⁶ The PBE functional has similarly been slightly modified (*RPBE*) to improve the performance for periodic systems,⁴⁷ but this modification actually destroys the hole condition (eq. (6.24)) for the exchange energy. An alternative modification using one additional parameter to give the acronym *mPBE* has also been proposed.⁴⁸

The *KT3* (*Keal–Tozer*) functional has been constructed as a combination of LDA and OPTX exchange combined with the LYP correlation functional, and modified with an additional gradient term, all multiplied with fitting coefficients, as shown in eq. (6.43).⁴⁹

$$E_{\text{xc}}^{\text{KT3}} = aE_x^{\text{LDA}} + bE_x^{\text{OPTX}} + cE_c^{\text{LYP}} + d \int \frac{|\nabla \rho|^2}{\rho^{4/3} + e}\tag{6.43}$$

The a , b and c coefficients have been optimized with respect to experimental quantities such as atomization energies and geometries, while the d and e parameters are fitted to NMR nuclear shielding constants. The primary focus of KT3 and earlier versions (KT1 and KT2) is to provide a functional suitable for calculating shielding constants, which other standard functionals have difficulties with.

6.5.3 Higher order gradient or meta-GGA methods

The logical extension of GGA methods is to allow the exchange and correlation functionals to depend on higher order derivatives of the electron density, with the Laplacian ($\nabla^2 \rho$) being the second-order term. Alternatively, the functional can be taken to depend on the orbital kinetic energy density τ , which for a single orbital is identical to the von Weizsäcker kinetic energy τ_w (eq. (6.3)).

$$\begin{aligned}\tau(\mathbf{r}) &= \frac{1}{2} \sum_i^{\text{occ}} |\nabla \phi_i(\mathbf{r})|^2 \\ \tau_w(\mathbf{r}) &= \frac{|\nabla \rho(\mathbf{r})|^2}{8\rho(\mathbf{r})}\end{aligned}\quad (6.44)$$

The orbital kinetic energy density and the Laplacian of the density essentially carry the same information, since they are related via the orbitals and the effective potential (all potential terms in the KS equation).

$$\tau(\mathbf{r}) = \frac{1}{2} \sum_i^{\text{occ}} \epsilon_i |\phi_i(\mathbf{r})|^2 - v_{\text{eff}}(\mathbf{r})\rho(\mathbf{r}) + \frac{1}{2} \nabla^2 \rho(\mathbf{r}) \quad (6.45)$$

This may also be seen from the gradient expansion of τ for slowly varying densities.⁵⁰

$$\tau(\mathbf{r}) = \frac{3}{10} (6\pi^2)^{2/3} \rho(\mathbf{r})^{5/3} + \frac{1}{72} \frac{|\nabla \rho(\mathbf{r})|^2}{\rho(\mathbf{r})} + \frac{1}{6} \nabla^2 \rho(\mathbf{r}) + O(\nabla^4 \rho(\mathbf{r})) \quad (6.46)$$

Inclusion of either the Laplacian or orbital kinetic energy density as a variable leads to the so-called *meta-GGA* functionals, and functionals which in general use orbital information may also be placed in this category. Calculation of the orbital kinetic energy density is numerically more stable than calculation of the Laplacian of the density, and the two τ functions in eq. (6.44) are common components of meta-GGA functionals.

One of the earliest attempts to include kinetic energy functionals was by *Becke* and *Roussel* (BR), who proposed the exchange functional shown in eq. (6.47).⁵¹

$$\begin{aligned}\epsilon_x^{\text{BR}} &= -\frac{2 - (2 + ab)e^{-ab}}{4b} \\ a^3 e^{-ab} &= 8\pi\rho \\ a(ab - 2) &= b \frac{\nabla^2 \rho - 4(\tau - \tau_w)}{\rho}\end{aligned}\quad (6.47)$$

A similar correlation functional shown in eq. (6.48) was proposed somewhat later by A. D. Becke (B95) and is one of the few functionals that does not have the self-interaction problem.⁵²

$$\begin{aligned}\epsilon_c^{\text{B95}} &= \epsilon_c^{\alpha\beta} + \epsilon_c^{\alpha\alpha} + \epsilon_c^{\beta\beta} \\ \epsilon_c^{\alpha\beta} &= [1 + a(x_\alpha^2 + x_\beta^2)]^{-1} \epsilon_c^{\text{PW},\alpha\beta} \\ \epsilon_c^{\sigma\sigma} &= [1 + bx_\sigma^2]^{-2} \frac{(\tau - \tau_w)_\sigma}{2^{5/3} C_F \rho_\sigma^{5/3}} \epsilon_c^{\text{PW},\sigma\sigma}\end{aligned}\quad (6.48)$$

Here σ runs over α and β spins, x_σ is defined in eq. (6.39) with the implicit spin dependence denoted by the subscript σ , a and b are fitting parameters, and ϵ_c^{PW} is the Perdew–Wang parameterization of the LSDA correlation functional (eq. (6.38)).

The HCTH functional has been extended to also include the kinetic energy density as a variable, producing the acronym τ -HCTH.⁵³ The *VSXC* (*Voorhis–Scuseria eXchange–Correlation*) functional similarly includes the kinetic energy density and contains 21 parameters that are fitted to experimental data.⁵⁴ The *TPSS*

(*Tao–Perdew–Staroverov–Scuseria*) exchange–correlation functional, on the other hand, is a non-empirical version that represents a further development of the *PKZB* (*Perdew–Kurth–Zupan–Blaha*) functional,⁵⁵ and can be considered as the next improvement over the PBE functional.⁵⁶

6.5.4 Hybrid or hyper-GGA methods

From the Hamiltonian in eq. (6.5) and the definition of the exchange–correlation energy in eq. (6.10), an exact connection can be made between the exchange–correlation energy and the corresponding hole potential connecting the non-interacting reference and the actual system (see appendix B for details). The resulting equation is called the *Adiabatic Connection Formula* (ACF)⁵⁷ and involves integration over the parameter λ , which “turns on” the electron–electron interaction.

$$E_{xc} = \int_0^1 \langle \Psi_\lambda | \mathbf{V}_{xc}^{\text{hole}}(\lambda) | \Psi_\lambda \rangle d\lambda \quad (6.49)$$

In the crudest approximation (taking $\mathbf{V}_{xc}^{\text{hole}}$ to be linear in λ), the integral is given as the average of the values at the two end-points.

$$E_{xc} \approx \frac{1}{2} (\langle \Psi_0 | \mathbf{V}_{xc}^{\text{hole}}(0) | \Psi_0 \rangle + \langle \Psi_1 | \mathbf{V}_{xc}^{\text{hole}}(1) | \Psi_1 \rangle) \quad (6.50)$$

In the $\lambda = 0$ limit, the electrons are non-interacting and there is consequently no correlation energy, only exchange energy. Furthermore, since the *exact* wave function in this case is a single Slater determinant composed of KS orbitals, the exchange energy is *exactly* that given by Hartree–Fock theory (eq. (3.33)). If the KS orbitals were identical to the HF orbitals, the exchange energy would be precisely the energy calculated by HF wave mechanics methods. The last term in eq. (6.50) is still unknown. Approximating it by the LSDA result defines the *Half-and-Half* (H + H) method.⁵⁸

$$E_{xc}^{\text{H+H}} = \frac{1}{2} E_x^{\text{exact}} + \frac{1}{2} (E_x^{\text{LSDA}} + E_c^{\text{LSDA}}) \quad (6.51)$$

Since the GGA methods give a substantial improvement over LDA, a generalized version of the H+H method may be defined by writing the exchange energy as a combination of LSDA, exact exchange and a gradient correction term. The correlation energy may similarly be taken as the LSDA formula plus a gradient correction term. Models that include exact exchange are often denoted *hybrid* methods, the *Adiabatic Connection Model* (ACM) and *Becke 3 parameter functional* (B3) methods are examples of such hybrid models, with the popular B3LYP method defined by eq. (6.52).⁵⁹ An alternative version uses the PW91 correlation functional and has the acronym B3PW91, and an O3LYP combination has also been used.

$$E_{xc}^{\text{B3LYP}} = (1-a)E_x^{\text{LSDA}} + aE_x^{\text{exact}} + b\Delta E_x^{\text{B88}} + (1-c)E_c^{\text{LSDA}} + cE_c^{\text{LYP}} \quad (6.52)$$

The a , b and c parameters are determined by fitting to experimental data and depend on the chosen forms for E_x^{GGA} and E_c^{GGA} , with typical values being $a \sim 0.2$, $b \sim 0.7$ and $c \sim 0.8$. Subsequent versions denoted B97 and B98 employed ten fitting parameters,⁶⁰ but the improvements were rather marginal relative to the three parameters version.

The τ -HCTH functional has been augmented with exact exchange to produce the acronym τ -HCTH-hybrid.⁶¹ The PBE functional has also been improved by addition

of exact exchange to give the PBE0 functional (also denoted PBE1PBE in the literature),⁶² where the mixing coefficient for the exact exchange is argued to have a value of 0.25 from perturbation arguments.⁶³ Similarly, the third-rung TPSS functional has been augmented with ~10% exact exchange to give the TPSSH method.⁶⁴

Inclusion of exact HF exchange is often found to improve the calculated results, although the optimum fraction to include depends on the specific property of interest. The improvement of new functionals by inclusion of a suitable fraction of exact exchange is now a standard feature. At least part of the improvement may arise from reducing the self-interaction error, since HF theory is completely self-interaction-free.

6.5.5 Generalized random phase methods

At the fifth level of the Jacob's ladder classification, the full information of the KS orbitals is employed, i.e. not only the occupied but also the virtual orbitals are included. The formalism here becomes similar to those used in the random phase approximation (Section 10.9), but very little work has appeared on such methods. Inclusion of the virtual orbitals is expected to significantly improve on, for example, dispersion (such as van der Waals) interactions, which is a significant problem for almost all current functionals.

One approach that can be considered as falling into this category is the class of *Optimized Effective Potential* (OEP) methods.⁶⁵ The central idea is that the energy as a functional of the density is unknown (or at least the exchange–correlation part is), but the energy as a function of the orbitals is well known from wave function theory to a given order in the correlation, as defined for example by a perturbation expansion. Since the density is given by the sum of the square of the orbitals, this *implicitly* defines the energy as a function of the density. By requiring that the density derived from a Kohn–Sham calculation using a single-determinant wave function exactly matches the density derived from a (correlated) wave function, this implicitly defines the exchange–correlation potential.

The reference wave functions have so far been based on an MBPT type expansion (Section 4.8). The OEP1 method is defined by terminating the reference density at first order in the perturbation series. Since correlation only enters the perturbation expansion at order two, this yields the exchange-only potential. Terminating the expansion at second order defines the OEP2 method and corresponds to constructing a KS determinant that yields the (generalized) MP2 density. From the condition that the MP2-like density matrix matches that from the KS determinant, one may derive a set of coupled equations at the orbital level that provides the exchange–correlation potential correct to second order in the correlation. The OEP2 method is computationally equivalent to an iterative MP2 calculation, i.e. such calculations are computationally more expensive than standard DFT methods. Furthermore the OEP2 method has basis set requirements similar to other correlated wave function methods and thus cannot benefit from the faster basis set convergence of other DFT methods. Not surprisingly, OEP2 provides results of roughly MP2 quality although, in favourable cases, the performance may approach those from coupled cluster calculations. It does have the desirable feature that it can describe for example dispersion interactions, which are problematic with almost all traditional functionals.

Whether one should consider the OEP method as a density or wave functional theory is an open question, as it clearly tries to combine the best of both worlds. It has the advantage of being able to systematically improve the results towards the exact limit, but inherits also the wave function disadvantages of a slow convergence with respect to basis set size.

6.5.6 Functionals overview

The introduction of GGA and hybrid functionals during the early 1990s yielded a major improvement in terms of accuracy for chemical applications, and resulted in the Nobel prize being awarded to W. Kohn and J. A. Pople in 1998. Progress since this initial exciting developments has been slower, and the (in)famous B3LYP functional⁵⁹ proposed in 1993 still represents one of the most successful in terms of overall performance. Unfortunately, neither the addition of more fitting parameters, the addition of more variables in the functionals, nor imposing more fundamental restrictions for the functional form have (yet) provided models with a significantly better overall performance.⁶⁶ Although the performance for a given property can be improved by tailoring the functional form or parameters, such measures often result in the deterioration of the results for other properties.

It should be noted that the implicit cancellation of the long-range part of the exchange and correlation energies implies that the two functional parts should be at the same level of the ladder, and preferably developed in an integrated fashion. A popular topic in the literature is to search for a magic combination of exchange and correlation functionals, perhaps with a few adjustable scaling parameters and a choice of basis set, in order to reproduce a selected set of experimental data. This is not a theoretically justified procedure and should be considered merely as data fitting without much physical relevance. Nevertheless, such a procedure can of course be taken as an “experimental” fitting function that can be useful for predicting specific properties for a series of compounds.

Table 6.1 shows an overview of commonly used functionals given by their acronym, and placed in the Jacob’s ladder classification. One may furthermore differentiate the functionals based on their use (or lack) of experimental data for assigning values to the parameters in the functional forms. The non-empirical ones such as the PW86,

Table 6.1 Perdew classification of exchange–correlation functionals

Level	Name	Variables	Examples
1	Local density	ρ	LDA, LSDA, X_α
2	GGA	$\rho, \nabla\rho$	BLYP, OPTX, OLYP, PW86, PW91, PBE, HCTH
3	Meta-GGA	$\rho, \nabla\rho, \nabla^2\rho$ or τ	BR, B95, VSXC, PKZB, TPSS, τ -HCTH
4	Hyper-GGA	$\rho, \nabla\rho, \nabla^2\rho$ or τ <i>HF exchange</i>	H+H, ACM, B3LYP, B3PW91, O3LYP, PBE0, TPSSh, τ -HCTH-hybrid
5	Generalized RPA	$\rho, \nabla\rho, \nabla^2\rho$ or τ <i>HF exchange</i> <i>Virtual orbitals</i>	OEP2

PW91, PBE and TPSS functionals use the free parameters to fulfil as many of the requirements in Section 6.5 as possible at each level. Empirical ones such as the BLYP, B3LYP, HTCT and VSXC, on the other hand, attempt to improve the performance by fitting the free parameters to give good agreement with experimental data. This means that these functionals often perform (slightly) better than the non-empirical ones for systems that resemble those in the parameterization set. Since the parameterization data are usually molecular systems, this means that they are often preferred for chemical purposes, but often give inferior performance for, for example, periodic systems such as metals. Note also that most common functionals belong to levels 2 and 4, as inclusion of HF exchange historically has preceded the development of functionals using derivatives beyond first order. As one moves along the rungs of the ladder, it is expected (or hoped) that the accuracy will improve, but there is no guarantee that this is the case.

6.6 Performance and Properties of Density Functional Methods

An evaluation of the performance of the plethora of different functionals for a variety of properties is a major undertaking.⁴ We will here just quote two sets of results:

- (1) *Root Mean Square (RMS)* errors of atomization energies, ionization potentials, electron and proton affinities over the data set of 407 compounds selected from the G3 data set against experimental data.⁶⁷ In addition the RMS error for the residual gradient at the experimental equilibrium geometry is taken as a measure of the accuracy of the functionals for predicting equilibrium geometries. It should be noted that the evaluation data are the same data used for *optimizing* the parameters in the HTCT functional, and this functional will therefore naturally display good performance. The results are obtained by using a TZP type basis set.
- (2) *Mean Absolute Deviation (MAD)* of atomization energies over the 223 molecules in the G3 data set against experimental data.⁶⁴ The results were obtained using the 6-311++G(3df,3pd) basis set.

While results with the above basis sets are not converged to the basis set limit, the residual basis set errors are presumably well below the inherent errors in the functionals, and the performance thus reflects the quality of the exchange–correlation functionals (Table 6.2).

Note that the performance ordering of the functionals is not the same for the two sets of results. Only a minor part of this discrepancy can be attributed to the difference in basis sets, the remaining discrepancy is due to differences in the data sets.

The LSDA method performs somewhat better than Hartree–Fock, but all the gradient-corrected methods are clearly far superior. The PW91 and PBE functionals are somewhat poorer than the other GGA functionals, reflecting the fact that these do not contain parameters that have been fitted to give a good performance for these systems. Hybrid methods including exact exchange tend to perform (slightly) better than the corresponding pure functionals (e.g. BLYP/B3LYP and PBE/PBE0), but several of the more recent “pure” functionals such as OLYP and VSXC are comparable to for example B3LYP. Since the inclusion of HF exchange is computationally expensive for implementations relying on plane waves for expanding the orbitals, or for programs taking advantage of various density fitting schemes, this represent a computational

Table 6.2 Comparison of the performance of DFT methods

Functional	RMS (gradient)	RMS (kJ/mol)	MAD (kJ/mol)
HF	35	649	885
LSDA	16	439	510
PW91	15	80	99
PBE	16	87	93
PKBZ	21	75	29
BLYP	19	41	40
PBE0	11	50	28
OLYP	14	40	25
B3LYP	11	40	21
VSXC	11	39	14
HTCT	11	33	30
τ -HCTH	11	31	
τ -HCTH-hybrid	10	26	
TPSS			24
TPSSh			16

advantage more than a fundamental theoretical improvement. In general, it is found that DFT methods often give geometries and vibrational frequencies for stable molecules of the same or better quality than MP2, at a computational cost similar to HF. For systems containing multi-reference character, where MP2 usually fails badly, DFT methods are often found to generate results of a quality comparable to those obtained with coupled cluster methods⁶⁸ (see also Section 11.7.3). Handy and Cohen have argued that the BLYP and B3LYP forms are probably close to the optimum with respect to performance for a functional depending only on the gradients of the density.⁶⁹

A significant advantage is that DFT methods based on unrestricted determinants (analogous to UHF, Section 3.7) for open-shell systems are not very prone to “spin contamination”, i.e. $\langle \mathbf{S}^2 \rangle$ is normally close to $S_z(S_z + 1)$ (see also Sections 4.4 and 11.5.3). This is a consequence of electron correlation being included in the single-determinantal wave function (by means of E_{xc}). Actually, it has been argued that “spin contamination” is not well defined in DFT methods, and that $\langle \mathbf{S}^2 \rangle$ should *not* be equal to $S_z(S_z + 1)$.⁷⁰ The argument is that real systems display “spin polarization”, i.e. there are point in space where ρ_α is larger than ρ_β (assuming that the number of α -electron is larger than the number of β electrons). This effect cannot be achieved by a restricted open-shell type determinant (analogous to ROHF), only by an unrestricted treatment that allows the α - and β orbitals to be different. It is somewhat unclear whether this argument hold for cases with $\langle \mathbf{S}^2 \rangle$ values very different from $S_z(S_z + 1)$, as in for example systems with multiple open-shell fragments.⁷¹ Another consequence of the presence of E_{xc} is that restricted type determinants are much more stable toward symmetry breaking to an unrestricted determinant (Section 3.8.3) than Hartree–Fock wave functions. For ozone (Section 4.4), for example, it is not possible to find a lower energy solution corresponding to UHF for “pure” DFT methods (such as LSDA or BLYP), although

those including exact exchange (such as B3LYP) display a triplet instability. This “inverse” symmetry breaking is in some cases problematic. In radical cations, for example, DFT methods usually refuse to localize the spin and charge, and thereby create unrealistic energy surfaces.

The Lagrange multipliers arising in Hartree–Fock theory from the orthogonality constraints of the orbitals are molecular orbital energies, and the occupied orbital energies correspond to ionization potentials in a frozen orbital approximation via Koopmans’ theorem. The corresponding Lagrange multipliers in DFT do not have the same formal relationship, since Koopmans’ theorem does not hold unless the *exact* exchange–correlation functional is employed. For approximate XC functionals, the Lagrange multipliers can be interpreted as the derivative of the total energy with respect to the occupation number of the orbital, often called the *Janak theorem*⁷² but discussed first by Slater,⁷³ and this is of course also closely related to experimentally measured ionization potential.

$$\frac{\partial E}{\partial n_i} = \epsilon_i \quad (6.53)$$

The Lagrange multipliers may also be considered as approximations to ionization potentials using *relaxed* orbitals, and in practice give quite accurate results for the valence orbitals.⁷⁴ In earlier work the orbital energies resulting from Kohn–Sham calculations were not considered to have any physical relevance, since they often showed poor agreement with ionization potentials, and orbital energy differences correlated poorly with excitation energies. It is now clear that part of the poor agreement was due to the self-interaction error embedded in LDA and GGA methods, while more modern functionals yield much improved results.

Another difference is that the unoccupied orbital energies in Hartree–Fock theory are determined in the field of N electrons and therefore correspond to *adding* an electron, i.e. the electron affinity. The virtual orbitals in density functional theory, on the other hand, are determined in the field of $N - 1$ electrons and therefore correspond to *exciting* an electron, i.e. unoccupied orbitals in DFT tend to be significantly lower in energy than the corresponding HF ones, and the highest occupied molecular orbital–lowest unoccupied molecular orbital (HOMO–LUMO) gaps are therefore much smaller with DFT methods than for HF. This also means that orbital energy differences in DFT are reasonable estimates of excitation energies, in contrast to HF methods where excitation energies involve additional Coulomb and exchange integrals. The LSDA method usually underestimates the HOMO–LUMO gap, leading to the incorrect prediction of metallic behaviour for certain semiconducting materials.

Although it is clear that there are many similarities between wave mechanics HF theory and DFT, there is an important difference. If the *exact* $E_{xc}[\rho]$ was known, DFT would provide the *exact* total energy, *including* electron correlation. DFT methods therefore have the potential of including the computationally difficult part in wave mechanics, the correlation energy, at a computational effort similar to that for determining the uncorrelated HF energy. Although this certainly is the case for approximations to $E_{xc}[\rho]$ (as illustrated above), this is not necessarily true for the *exact* $E_{xc}[\rho]$. It may well be that the exact $E_{xc}[\rho]$ functional is so complicated that the computational effort for solving the KS equations will be similar to that required for solving the

Schrödinger equation (exactly) with a wave mechanics approach. Indeed, unless one believes that the Schrödinger equation contains superfluous information, this is likely to be the case. Since exact solutions are generally not available in either approach, the important question is instead what the computational cost is for generating a solution of a given accuracy. In this respect, DFT methods have very favourable characteristics.

6.7 DFT Problems

Despite the many successes of DFT, there are some areas where the current functionals are known to perform poorly.

- Weak interactions due to dispersion forces (part of van der Waals type interactions) arise from electron correlation in wave function methods, but this is poorly described by current DFT methods.⁷⁵ Rare gas atoms should show a slight attraction, but most functionals display a purely repulsive energy curve, and those that do predict an attraction underestimate the effect and the variation between systems.⁷⁶ Furthermore, none have the correct R^{-6} limiting behaviour in the long distance limit, although very recent developments appear to provide quite accurate results with only a single parameter.⁷⁷ In some approaches, an empirical attraction term is added that improves the performance,⁷⁸ but this is clearly an *ad hoc* repair. Owing to the general overestimation of bond strengths, LSDA does predict an attraction between rare gas atoms, but significantly overestimates the magnitude. Hydrogen bonding, however, is mainly electrostatic and is reasonably well accounted for by many DFT functionals.
- Loosely bound electrons, such as anions arising from systems with relatively low electron affinities, represent a problem for exchange–correlation functionals that do not include self-interaction corrections or correct for the incorrect long-range behaviour of the exchange–correlation potential. Since loosely bound electrons by definition have most of the associated density far from the nuclei, this may cause the self-interaction error to be larger than the actual binding energy, and thus lead erroneously to an unbound electron. In actual calculations using a limited basis set, this may not be obvious, since the outer electron is confined by the most diffuse basis function. A positive HOMO energy, however, is a clear warning sign, and extending the basis set with many diffuse functions in such cases may cause the outer electron to drift away from the atom. This means that only systems with high electron affinities have a well-defined basis set limiting value. Nevertheless, a medium-sized basis set with a single set of diffuse functions will in many cases give a reasonable estimate of the experimental electron affinity.⁷⁹ The basis set confines the outer electron to be in the correct physical space, and the exchange–correlation functional gives a reasonable estimate of the energy of this density. It should be noted that the relatively good performance is in essence due to a correct physical description, rather than a correct theoretical methodology.
- For chemically bonded systems, analysis⁸⁰ similar to the H_2 system in Section 6.4 suggest that bonds involving:
 - two-centre two-electrons (e.g. normal covalent bonds),
 - two-centre four-electrons (e.g. steric repulsion between closed shell systems), and

- three-centre three-electrons (e.g. radical abstraction) should be reasonably described by gradient-corrected methods. Systems involving:
- two-centre one-electron (e.g. radical cations),
- two-centre three-electrons (e.g. radical anions), and
- three-centre four electrons (e.g. atom transfer transition structures)

are, however, predicted to be too stable. The dissociation of charged odd-electron systems is a problem for most DFT methods, with the dissociation energy profile displaying an artificial barrier and an incorrect dissociation energy, often in error by as much as 100 kJ/mol. Transition structures are similarly predicted to be too stable (barriers are underestimated) by functionals that do not include exact exchange. Since Hartree–Fock overestimates activation barriers, hybrid methods involving exact exchange, however, often give reasonable barriers.

- The absence of a wave function makes a direct description of excited states with the same symmetry as the ground state problematic. Excited states must be orthogonal to the ground state, which is easy to enforce if the spatial or spin symmetry differ, but difficult to ensure for excited states having the same spatial and spin symmetry. Excited state properties, however, can be calculated by time-dependent DFT (linear response) methods, since the excited state is never needed explicitly. Such calculations can give for example excitation energies and transition moments, as well as gradients of the excited surface, which allows excited states to be optimized.
- The accuracy of excitation energies is typically ~ 0.5 eV for valence states, but Rydberg states, where the electron is excited into a diffuse orbital, can be in error by several eV. This problem has the same physical reason as the anion problem above, and can be solved by using corrections for the asymptotic behaviour of the exchange–correlation potential.^{81,82} Such *Asymptotic Corrected* (AC) functionals display much improved predictions for response properties.
- The exchange–correlation functional is inherently local, depending only on the density and possibly its derivatives at a given point, and this causes DFT methods to be inherently unsuitable for describing charge transfer systems, where an electron is transferred over a large distance. Such systems are predicted to have excitation energies that are too low by several eV.⁸³
- Relative energies of states with different spin multiplicity are often poorly described. In HF theory, the energy difference between a singlet and triplet state with the same orbital occupancy is given by an exchange integral. In DFT, this must be described by the exchange–correlation functional, which only depends on the electron density. If the two spin states arise from the same electron configuration the two electron densities are very similar, and this makes the results sensitive to the details of the exchange–correlation functional. These problems are especially problematic for transition metal systems, where several low-energy spin states are often possible, and many of these cannot be described by a single determinant. Pure DFT methods favour low spin states while HF favours high spin states, and hybrid methods with a suitable parameterized amount of exact exchange perform better.⁸⁴ These problems can perhaps be improved by adding current density terms to the DFT formalism but this is not yet a commonly used procedure since it requires that the orbitals be allowed to become complex.

- Individual spatial components of a spin multiplet may have different energies, even in the absence of a magnetic field. The boron atom, for example, has the electron configuration $1s^2 2s^2 2p^1$, and the single p -electron can be in either a p_{-1} , p_0 or p_{+1} orbital. These should all have the same energy, but since the density associated with the p_0 orbital is different from that of a $p_{\pm 1}$ orbital, their energies as a result differ by ~ 25 kJ/mol. This is clearly non-physical, but can be significantly improved by introducing current density terms.⁸⁵

6.8 Computational Considerations

The strength of DFT is that only the total density needs to be considered. In order to calculate the kinetic energy with sufficient accuracy, however, orbitals have to be re-introduced. Nevertheless, Kohn–Sham DFT displays a computational cost similar to HF theory, with the possibility of providing more accurate (exact, in principle) results.

Once an exchange–correlation functional has been selected, the computational problem is very similar to that encountered in wave mechanics HF theory: determine a set of orthogonal orbitals that minimizes the energy. Since the $J[\rho]$ (and $E_{xc}[\rho]$) functional depends on the total density, a determination of the orbitals involves an iterative sequence. The orbital orthogonality constraint may be enforced by the Lagrange method (Section 12.5), again in complete analogy with wave mechanics HF methods (eq. (3.34)).

$$L[\rho] = E_{\text{DFT}}[\rho] - \sum_{ij}^{N_{\text{orb}}} \lambda_{ij} (\langle \phi_i | \phi_j \rangle - \delta_{ij}) \quad (6.54)$$

Requiring the variation of L to vanish provides a set of equations involving an effective one-electron operator (\mathbf{h}_{KS}), similar to the Fock operator in wave mechanics (eq. (3.36)).

$$\begin{aligned} \mathbf{h}_{\text{KS}} \phi_i &= \sum_j^{N_{\text{orb}}} \lambda_{ij} \phi_j \\ \mathbf{h}_{\text{KS}} &= \frac{1}{2} \nabla^2 + \mathbf{V}_{\text{eff}} \\ \mathbf{V}_{\text{eff}}(\mathbf{r}) &= \mathbf{V}_{\text{nc}}(\mathbf{r}) + \int \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}' + \mathbf{V}_{\text{xc}}(\mathbf{r}) \end{aligned} \quad (6.55)$$

The effective potential contains the nuclear contribution, the electronic Coulomb repulsion and the exchange–correlation potential, which is given as the derivative of the energy (eq. (6.29)) with respect to the density.

$$\mathbf{V}_{\text{xc}}(\mathbf{r}) = \frac{\delta E_{\text{xc}}[\rho]}{\delta \rho(\mathbf{r})} = \varepsilon_{\text{xc}}[\rho(\mathbf{r})] + \int \rho(\mathbf{r}') \frac{\delta \varepsilon_{\text{xc}}(\mathbf{r}')}{\delta \rho(\mathbf{r})} d\mathbf{r}' \quad (6.56)$$

A unitary transformation that makes the matrix of the Lagrange multiplier diagonal may again be chosen, producing a set of canonical KS orbitals. The resulting pseudo-eigenvalue equations are known as the Kohn–Sham equations.

$$\mathbf{h}_{\text{KS}} \phi_i = \varepsilon_i \phi_i \quad (6.57)$$

The KS orbitals can be determined completely by a numerical procedure, analogously to numerical HF methods. In practice, such procedures are limited to small systems,

and essentially all calculations employ an expansion of the KS orbitals in an atomic basis set.

$$\phi_i = \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} \quad (6.58)$$

The basis functions are often the same as used in wave mechanics for expanding the HF orbitals, although basis functions specifically optimized for DFT have recently been proposed (see Section 5.4.7 for details).

The variational procedure again leads to a matrix equation in the atomic orbital basis that can be written in the following form (compare to eq. (3.51)).

$$\begin{aligned} \mathbf{h}_{\text{KS}} \mathbf{C} &= \mathbf{S} \mathbf{C} \epsilon \\ h_{\alpha\beta} &= \langle \chi_{\alpha} | \mathbf{h}_{\text{KS}} | \chi_{\beta} \rangle \\ S_{\alpha\beta} &= \langle \chi_{\alpha} | \chi_{\beta} \rangle \end{aligned} \quad (6.59)$$

The \mathbf{h}_{KS} matrix is analogous to the Fock matrix in wave mechanics, and the one-electron and Coulomb parts are identical to the corresponding Fock matrix elements. The exchange–correlation part, however, is given in terms of the electron density, and possibly also involves derivatives of the density or orbitals.

$$\int \chi_{\alpha}(\mathbf{r}) \mathbf{V}_{\text{xc}}[\rho(\mathbf{r}), \nabla \rho(\mathbf{r})] \chi_{\beta}(\mathbf{r}) d\mathbf{r} \quad (6.60)$$

Since the \mathbf{V}_{xc} functional depends on the integration variables implicitly via the electron density, these integrals cannot be evaluated analytically but must be generated by a numerical integration.

$$\int \chi_{\alpha}(\mathbf{r}) \mathbf{V}_{\text{xc}}[\rho(\mathbf{r}), \nabla \rho(\mathbf{r})] \chi_{\beta}(\mathbf{r}) d\mathbf{r} \approx \sum_k^G \mathbf{V}_{\text{xc}}[\rho(\mathbf{r}_k), \nabla \rho(\mathbf{r}_k)] \chi_{\alpha}(\mathbf{r}_k) \chi_{\beta}(\mathbf{r}_k) \Delta \mathbf{v}_k \quad (6.61)$$

As the number of grid points G goes to infinity, the approximation becomes exact. In practice, the number of points is selected based on the desired accuracy of the final results, i.e. if the energy is only required with an accuracy of 10^{-3} , the number of integration points can be smaller than if the energy is required with an accuracy of 10^{-5} .⁸⁶ There are also some technical skills involved in selecting the optimum distribution of a given number of points to yield the best accuracy, i.e. the points should be dense where the function \mathbf{V}_{xc} varies most. The grid is usually selected as being spherical around each nucleus, making it dense in the radial direction near the nucleus, and dense in the angular part in the valence space. For typical applications, 1000–10000 points are used for each atom.⁸⁷ It should be noted that only the larger of such grids approach saturation, i.e. in general the energy will depend on the number (and location) of grid points. In order to compare energies for different systems, the same grid must therefore be used. The grid plays the same role for E_{xc} as the basis set for the other terms. Just as it is improper to compare energies calculated with different basis sets, it is not justified to compare DFT energies calculated with different grid sizes. Furthermore, an incomplete grid may lead to “grid superposition errors” analogous to basis set superposition errors (Section 5.10).

With an expansion of the orbitals in basis functions, the number of integrals necessary for solving the KS equations rises as M_{basis}^4 , owing to the Coulomb integrals in the

J functional (and possibly also “exact” exchange in the hybrid methods). The number of grid points for the numerical E_{xc} integration (eq. (6.61)) increases linearly with the system size, and the computational effort for the exchange–correlation term rises as GM_{basis}^2 , i.e. a cubic dependence of the system size. When the Coulomb (and possibly “exact” exchange) term is evaluated directly from integrals over basis functions, DFT methods scale formally as M_{basis}^4 . However, as discussed in Section 3.8.6, the Coulomb (and exchange) part can be calculated with an effort that scales only as M_{basis}^1 for large systems with for example fast multipole methods. The numerical integration required for the exchange and correlation parts may also be reduced to a computation cost that scales linearly with system size, i.e. with modern techniques DFT methods have true linear scaling.⁸⁸ This opens up the possibility of performing accurate calculations on systems containing thousands of atoms, which is likely to have impacts on many areas outside traditional computational chemistry.

Nevertheless, the formal M_{basis}^4 scaling has spawned approaches that reduce the dependence to M_{basis}^3 . This may be achieved by *fitting* the electron density to a linear combination of functions, and using the fitted density in evaluating the J integrals in the Coulomb term.

$$\rho \approx \sum_{\alpha}^{M_{\text{fit}}} a'_{\alpha} \chi'_{\alpha} \quad (6.62)$$

The density fitting functions may be the same as those used in expanding the orbitals, but more often an auxiliary basis that is optimized for density fitting is used. The fitting constants a'_{α} are often chosen such that the Coulomb energy arising from the *difference* between the exact and fitted densities is minimized, subject to the constraint of charge conservation.⁸⁹ The J integrals then become eq. (6.63), which only involves three basis functions, thereby reducing the computational effort to M_{basis}^3 .

$$\int \chi_{\alpha}(1) \chi_{\beta}(1) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \chi'_{\gamma}(2) d\mathbf{r}_1 d\mathbf{r}_2 \quad (6.63)$$

Alternative versions where the Coulomb part of the Kohn–Sham matrix is assembled using plane waves as the auxiliary basis have also been proposed and, properly implemented, these achieve linear scaling even for small systems and for large basis sets.⁹⁰

The use of grid-based techniques for the numerical integration of the exchange–correlation contribution has some disadvantages when derivatives of the energy are desired. For this reason, there is also interest in developing grid-free DFT methods where the exchange–correlation potential is expressed completely in terms of analytical integrals.⁹¹

The computational cost of a DFT calculation depends strongly on the implementation strategy. The use of DFT in the chemical community has to a large extent been introduced by modifying existing programs designed for wave function methods, and in these cases the numerical integration of the exchange–correlation energy adds a small overhead relative to an HF calculation. Programs designed for DFT from the outset, on the other hand, can exploit the reductions arising from density fitting, and can consequently run significantly faster than a wave function HF calculation.⁹² Furthermore, the use of grid-based methods for evaluating the Coulomb and

exchange–correlation contributions means that almost any kind of basis functions can be used, including Slater type orbitals.

Finally, *DFT methods are one-dimensional* just like HF methods, and increasing the size of the basis set allows a better and better description of the KS orbitals. Since the DFT energy depends directly on the electron density, it has an exponential convergence with respect to basis set size, analogously to HF methods, and a polarized triple zeta type basis usually gives results close to the basis set limit.

6.9 Final Considerations

Should DFT methods be considered *ab initio* or semi-empirical? If *ab initio* is taken to mean the absence of fitting parameters, LSDA methods are *ab initio* but gradient-corrected methods may or may not be. The LSDA exchange energy contains no parameters and the correlation functional is known accurately as a tabulated function of the density. The use of a parameterized interpolation formula in practical calculations does not represent fitting in order to improve the performance for atomic and molecular systems. Some gradient-corrected methods (e.g. the B88 exchange and the LYP correlation), however, contain parameters that are fitted to give the best agreement with experimental atomic data, but the number of parameters is significantly smaller than for semi-empirical methods. The semi-empirical PM3 method (Section 3.11.5), for example, has 18 parameters for each atom, while the B88 exchange functional only has one fitting constant, valid for the whole periodic table. Functionals such as VSXC contains a moderate number of parameters (21), while other functionals such as PBE are derived entirely from theory and can consequently be considered “pure” *ab initio*.

If *ab initio* is taken to mean that the method is based on theory, which in principle is able to produce the exact results, DFT methods are *ab initio*. The only caveat is that current methods cannot yield the exact results, even in the limit of a complete basis set, since the functional form of the exact exchange–correlation energy is not known. At present it is easier to systematically improve on a wave function description than adding corrections to the energy functional in DFT. Methods using reduced density matrices are still in their infancy, but promising results have been obtained in recent years.

It is perhaps a little disturbing that seemingly very different functionals give similar-quality results.⁹³ Levy and Perdew⁹⁴ and others⁹⁵ have shown how wave functions of near exact quality (such as CCSD(T)) can be “inverted” by a “constrained search” method to give near exact KS orbitals and corresponding exchange–correlation potentials. Comparisons of such “exact” \mathbf{V}_{xc} potentials with those discussed in the previous subsections have revealed large deviations and erroneous functional behaviour.⁹⁶ Since many of these functionals perform well in practical applications, it is clear that the performance is not particularly sensitive to details in the functional, and that the good performance to some extent is due to error cancellations.

Although gradient-corrected DFT methods have been shown to give impressive results, even for theoretically difficult problems, the lack of a systematic way of extending a series of calculations to approach the exact result is a major drawback of DFT. The results converge toward a certain value as the basis set is increased, but theory does not allow an evaluation of the errors inherent in this limit (such as the systematic overestimation of vibrational frequencies with wave mechanics HF methods). Fur-

thermore, although a progression of methods such as LSDA, BLYP and B3LYP has provided successively lower errors for a suitable set of reference data (such as that used for calibrating the Gaussian-2 model), there is no guarantee that the same progression will provide better and better results for a specific property of a given system. Indeed, LSDA methods may in some cases provide better results, even in the limit of a large basis set, than either of the more “complete” gradient-corrected models. The quality of a given result can therefore only be determined by comparing the performance for similar systems where experimental or high-quality wave mechanics results are available. In this respect, DFT resembles semi-empirical methods. Nevertheless, DFT methods, especially those involving gradient corrections and hybrid methods, are significantly more accurate (and the errors are much more uniform) than those of for example the MNDO family, and DFT is consequently a valuable tool for systems where a (very) high accuracy is not needed.

References

1. P. Hohenberg, W. Kohn, *Phys. Rev.*, **136** (1964), B864.
2. P.-O. Löwdin, *Int. J. Quant. Chem.*, **S19** (1986), 19.
3. R. G. Parr, W. Yang, *Density Functional Theory*, Oxford University Press, 1989; L. J. Bartolotti, K. Flurichick, *Rev. Comp. Chem.*, **7** (1996), 187; A. St-Amant, *Rev. Comp. Chem.*, **7** (1996), 217; T. Ziegler, *Chem. Rev.*, **91** (1991), 651; E. J. Baerends, O. V. Gritsenko, *J. Phys. Chem.*, **101** (1997), 5383.
4. W. Koch, M. C. Holthausen, *A Chemist's Guide to Density Functional Theory*, Wiley-VCH, 2000.
5. W. Kohn, L. J. Sham, *Phys. Rev.*, **140** (1965), A1133.
6. F. Bloch, *Z. Physik*, **57** (1929), 545.
7. P. A. M. Dirac, *Proc. Cambridge Phil. Soc.*, **26** (1930), 376.
8. R. D. Murphy, *Phys. Rev. A*, **24** (1981), 1682.
9. E. K. U. Gross, R. M. Dreizler, *Z. Phys. A*, **302** (1981), 103.
10. P. S. Svendsen, U. von Barth, *Phys. Rev. B*, **54** (1996), 17402.
11. Y. A. Wang, N. Govind, E. A. Carter, *Phys. Rev. B*, **60** (1999), 16350; S. S. Iyengar, M. Ernzerhof, S. N. Maximoff, G. E. Scuseria, *Phys. Rev. A*, **63** (2001), 052508.
12. A. J. Coleman, V. I. Yukalov, *Reduced Density Matrices: Coulson's Challenge*, Springer-Verlag, New York, 2000.
13. C. Garrod, J. Percus, *J. Math. Phys.*, **5** (1964), 1756.
14. D. A. Mazziotti, *Phys. Rev. A*, **65** (2002), 062511; D. Mazziotti, *Acc. Chem. Res.*, **39** (2006), 207.
15. D. A. Mazziotti, *J. Chem. Phys.*, **121** (2004), 10957.
16. S. Goedecker, C. J. Umrigar, *Phys. Rev. Lett.*, **81** (1998), 866; O. Gritsenko, K. Pernal, E. J. Baerends, *J. Chem. Phys.*, **122** (2005), 204102.
17. C. Kollmar, B. A. Hess, *J. Chem. Phys.*, **119** (2003), 4655, C. Kollmar, *J. Chem. Phys.*, **121** (2004), 11581.
18. P. M. W. Gill, D. L. Crittenden, D. P. O'Neill, N. A. Besley, *Phys. Chem. Chem. Phys.*, **8** (2005), 15.
19. O. V. Gritsenko, P. R. T. Schippen, E. J. Baerends, *J. Chem. Phys.*, **107** (1997), 5007.
20. A. D. Becke, *J. Chem. Phys.*, **122** (2005), 064101.
21. E. V. Ludena, J. M. Ugalde, X. Lopez, J. Fernandez-Rico, G. Ramirez, *J. Chem. Phys.*, **120** (2004), 540.
22. R. K. Nesbet, R. Colle, *J. Math. Chem.*, **26** (1999), 233.

23. J. P. Perdew, A. Zunger, *Phys. Rev. B*, **23** (1981), 5048.
24. O. A. Vydrov, G. E. Scuseria, *J. Chem. Phys.*, **121** (2004), 8187.
25. O. A. Vydrov, G. E. Scuseria, J. P. Perdew, A. Ruzsinszky, G. I. Csonka, *J. Chem. Phys.*, **124** (2006), 094108; O. A. Vydrov, G. E. Scuseria, *J. Chem. Phys.*, **124** (2006), 191101.
26. E. J. Baerends, *Phys. Rev. Lett.*, **87** (2001), 133004; A. D. Becke, *J. Chem. Phys.*, **119** (2003), 2972.
27. J. P. Perdew, J. Tao, V. N. Staroverov, G. E. Scuseria, *J. Chem. Phys.*, **120** (2004), 6898.
28. M. Levy, J. P. Perdew, *Phys. Rev. A*, **32** (1985), 2010.
29. E. H. Lieb, S. Oxford, *Int. J. Quant. Chem.*, **19** (1981), 427.
30. R. van Leeuwen, E. J. Baerends, *Phys. Rev. A*, **49** (1994), 2421.
31. J. P. Perdew, R. G. Parr, M. Levy, J. L. Balduz Jr, *Phys. Rev. Lett.*, **49** (1982), 1691.
32. A more complete list of functionals can be found in: G. E. Scuseria, V. N. Staroverov, in *Theory And Applications of Computational Chemistry: the First Forty Years*, C. E. Dykstra, G. Frenking, K. S. Kim, G. E. Scuseria, Eds, Elsevier, 2005.
33. S. Kurth, J. P. Perdew, P. Blaha, *Int. J. Quant. Chem.*, **75** (1999), 889; J. P. Perdew, A. Ruzsinszky, J. Tao, V. N. Staroverov, G. E. Scuseria, G. I. Csonka, *J. Chem. Phys.*, **123** (2005), 062201.
34. J. C. Slater, *Phys. Rev.*, **81** (1951), 385.
35. W. J. Carr Jr, *Phys. Rev.*, **122** (1961), 1437; W. J. Carr Jr, A. A. Maradudin, *Phys. Rev.*, **133** (1964), A371.
36. S. J. Vosko, L. Wilk, M. Nusair, *Can. J. Phys.*, **58** (1980), 1200; J. P. Perdew, Y. Wang, *Phys. Rev. B*, **45** (1992), 13244.
37. S.-K. Ma, K. A. Brueckner, *Phys. Rev.*, **165** (1968), 18.
38. A. D. Becke, *Phys. Rev. A*, **38** (1988), 3098.
39. G. Ortiz, P. Ballone, *Phys. Rev. B*, **43** (1991), 6376.
40. N. C. Handy, A. J. Cohen, *Mol. Phys.*, **99** (2001), 403.
41. F. A. Hamprecht, A. J. Cohen, D. J. Tozer, N. C. Handy, *J. Chem. Phys.*, **109** (1998), 6264; A. D. Boese, N. Doltsinis, N. C. Handy, M. Sprik, *J. Chem. Phys.*, **112** (2000), 1670; A. D. Boese, N. C. Handy, *J. Chem. Phys.*, **114** (2001), 5497.
42. C. Lee, W. Yang, R. G. Parr, *Phys. Rev. B*, **37** (1988), 785; B. Miehlich, A. Savin, H. Stoll, H. Preuss, *Chem. Phys. Lett.*, **157** (1989), 200.
43. J. P. Perdew, Y. Wang, *Phys. Rev. B*, **33** (1986), 8800.
44. J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh, C. Fiolhais, *Phys. Rev. B*, **46** (1992), 6671.
45. J. P. Perdew, K. Burke, M. Ernzerhof, *Phys. Rev. Lett.*, **77** (1996), 3865.
46. C. Adamo, V. Barone, *J. Chem. Phys.*, **108** (1998), 664.
47. B. Hammer, L. B. Hansen, J. K. Nørskov, *Phys. Rev. B*, **59** (1999), 7413.
48. C. Adamo, V. Barone, *J. Chem. Phys.*, **116** (2002), 5933.
49. T. W. Keal, D. Tozer, *J. Chem. Phys.*, **121** (2005), 5654.
50. M. Brack, B. K. Jennings, Y. H. Chu, *Phys. Lett.*, **65B** (1976), 1.
51. A. D. Becke, M. R. Roussel, *Phys. Rev. A*, **39** (1989), 3761.
52. A. D. Becke, *J. Chem. Phys.*, **104** (1996), 1040.
53. A. D. Boese, N. C. Handy, *J. Chem. Phys.*, **116** (2002), 9559.
54. T. V. Voorhis, G. E. Scuseria, *J. Chem. Phys.*, **109** (1998), 400.
55. J. P. Perdew, S. Kurth, A. Zupan, P. Blaha, *Phys. Rev. Lett.*, **82** (1999), 2544.
56. J. Tao, J. P. Perdew, V. N. Staroverov, G. E. Scuseria, *Phys. Rev. Lett.*, **91** (2003), 146401.
57. J. Harris, *Phys. Rev. A*, **29** (1984), 1648.
58. A. D. Becke, *J. Chem. Phys.*, **98** (1993), 1372.
59. A. D. Becke, *J. Chem. Phys.*, **98** (1993), 5648; P. J. Stephens, F. J. Devlin, C. F. Chabalowski, M. J. Frisch, *J. Phys. Chem.*, **98** (1994), 11623.
60. A. D. Becke, *J. Chem. Phys.*, **107** (1997), 8554; H. L. Schmider, A. D. Becke, **108** (1998), 9624.
61. A. D. Boese, N. C. Handy, *J. Chem. Phys.*, **116** (2002), 9559.

62. M. Ernzerhof, G. E. Scuseria, *J. Chem. Phys.*, **110** (1999), 5029.
63. J. P. Perdew, M. Ernzerhof, K. Burke, *J. Chem. Phys.*, **105** (1996), 9982.
64. V. N. Staroverov, G. E. Scuseria, J. Tao, J. P. Perdew, *J. Chem. Phys.*, **119** (2003), 12129; **121** (2004), 11507 (E).
65. J. D. Talman, W. F. Shadwick, *Phys. Rev. A*, **14** (1976), 36; A. F. Bonetti, E. Engel, R. N. Schmid, R. M. Dreizler, *Phys. Rev. Lett.*, **86** (2001), 2241; R. J. Bartlett, V. F. Lotrich, I. Schweihert, *J. Chem. Phys.*, **123** (2005), 062205.
66. G. K.-L. Chan, N. C. Handy, *J. Chem. Phys.*, **112** (2000), 5639.
67. A. D. Boese, J. M. L. Martin, N. C. Handy, *J. Chem. Phys.*, **119** (2003), 3005.
68. N. Oliphant, R. J. Bartlett, *J. Chem. Phys.*, **100** (1994), 6550.
69. N. C. Handy, A. J. Cohen, *J. Chem. Phys.*, **116** (2002), 5411.
70. J. A. Pople, P. M. W. Gill, N. C. Handy, *Int. J. Quant. Chem.*, **56** (1995), 303.
71. E. R. Davidson, A. E. Clark, *J. Phys. Chem. A*, **106** (2002), 7456.
72. J. F. Janak, *Phys. Rev. B*, **18** (1978), 7165.
73. J. C. Slater, *The Self-Consistent Field for Molecules and Solids: Quantum Theory of Molecules and Solids*, Vol. 4, McGraw-Hill, 1974.
74. D. P. Chong, O. V. Gritsenko, E. J. Baerends, *J. Chem. Phys.*, **116** (2002), 1760; V. Gritsenko, E. J. Baerends, *J. Chem. Phys.*, **120** (2004), 8364.
75. E. J. Meijer, M. Sprik, *J. Chem. Phys.*, **105** (1996), 8684; S. M. Cybulski, C. E. Seversen, *J. Chem. Phys.*, **122** (2005), 014117.
76. J. Tao, J. P. Perdew, *J. Chem. Phys.*, **122** (2005), 114102.
77. E. R. Johnson, A. D. Becke, *J. Chem. Phys.*, **123** (2005), 024101.
78. Q. Wu, W. Yang, *J. Chem. Phys.*, **116** (2001), 515.
79. J. C. Rienstra-Kiracofe, G. S. Tschumper, H. F. Schaefer III, S. Nandi, G. B. Ellison, *Chem. Rev.*, **102** (2002), 231.
80. O. V. Gritsenko, B. Ensing, P. R. T. Schipper, E. J. Baerends, *J. Phys. Chem. A*, **104** (2000), 8558.
81. M. Grüning, O. V. Gritsenko, S. J. A. van Gisbergen, E. J. Baerends, *J. Chem. Phys.*, **114** (2001), 652.
82. D. J. Tozer, N. C. Handy, *Mol. Phys.*, **101** (2003), 2669.
83. A. Dreuw, J. L. Weisman, M. Head-Gordon, *J. Chem. Phys.*, **119** (2003), 2943; O. Gritsenko, E. J. Baerends, *J. Chem. Phys.*, **121** (2004), 655.
84. R. J. Deeth, N. Fey, *J. Comp. Chem.*, **25** (2004), 1840; A. Fouqueau, M. E. Casida, L. M. L. Daku, A. Hauser, F. Neese, *J. Chem. Phys.*, **122** (2005), 044110.
85. S. N. Maximoff, M. Ernzerhof, G. E. Scuseria, *J. Chem. Phys.*, **120** (2004), 2104; A. D. Becke, *J. Chem. Phys.*, **117** (2002), 6935.
86. J. M. Perez-Jorda, A. D. Becke, E. San-Fabian, *J. Chem. Phys.*, **100** (1994), 6520.
87. P. M. W. Gill, B. G. Johnson, *Chem. Phys. Lett.*, **209** (1993), 506.
88. R. E. Stratmann, G. E. Scuseria, M. J. Frisch, *Chem. Phys. Lett.*, **257** (1996), 213.
89. B. I. Dunlap, J. W. D. Connolly, J. R. Sabin, *J. Chem. Phys.*, **71** (1979), 3396; U. Birkenheuer, A. B. Gordienko, V. A. Nasluzov, M. K. Fuchs-Rohr, N. Rösch, *Int. J. Quant. Chem.*, **102** (2005), 743.
90. J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, J. Hutter, *Comp. Phys. Comm.*, **167** (2005), 103.
91. K. S. Werpetsinski, M. Cook, *J. Chem. Phys.*, **106** (1997), 7124; Y. C. Zheng, J. E. Almlöf, *J. Mol. Struct.*, **388** (1996), 277; K. R. Glaesemann, M. S. Gordon, *J. Chem. Phys.*, **112** (2000), 10739.
92. C. F. Guerra, J. G. Snijders, G. te Velde, E. J. Baerends, *Theor. Chem. Acc.*, **99** (1998), 391.
93. R. Neumann, R. H. Nobes, N. C. Handy, *Mol. Phys.*, **87** (1996), 1.
94. M. Levy, J. P. Perdew, in *Density Functional Methods in Physics*, R. M. Dreizler, J. da Providencia, Eds, Plenum, 1985.

95. Q. Zhao, R. C. Morrison, R. G. Parr, *Phys. Rev. A.*, **50** (1994), 2138; O. V. Gritsenko, R. van Leeuwen, E. J. Baerends, *J. Chem. Phys.*, **104** (1996), 8535; D. J. Tozer, V. E. Ingamells, N. C. Handy, *J. Chem. Phys.*, **105** (1996), 9200; J. B. Lucks, A. J. Cohen, N. C. Handy, *Phys. Chem. Chem. Phys.*, **4** (2002), 4612.
96. C. J. Umrigar, X. Gonze, *Phys. Rev. A*, **50** (1994), 3827; S. Hirata, S. Ivanov, I. Grabowski, R. J. Bartlett, K. Burke, J. D. Talman, *J. Chem. Phys.*, **115** (2001), 1635.

7 Valence Bond Methods

Essentially all practical calculations for generating solutions to the electronic Schrödinger equation have been performed with molecular orbital methods. The zeroth-order wave function is constructed as a single Slater determinant and the MOs are expanded in a set of atomic orbitals, the basis set. In a subsequent step the wave function may be improved by adding electron correlation with either CI, MP or CC methods. There are two characteristics of such approaches: (1) the one-electron functions, the MOs, are delocalized over the whole molecule, and (2) an accurate treatment of the electron correlation requires many (millions or billions) “excited” Slater determinants. The delocalized nature of the MOs is partly a consequence of choosing the Lagrange multiplier matrix to be diagonal (canonical orbitals, eq. (3.42)), they may in a subsequent step be mixed to form localized orbitals (see Section 9.4) without affecting the total wave function. Such a localization, however, is not unique. Furthermore, delocalized MOs are at variance with the basic concept in chemistry that molecules are composed of structural units (functional groups) which to a very good approximation are constant from molecule to molecule. The MOs for propane and butane, for example, are quite different, although “common” knowledge is that they contain CH_3 and CH_2 units that in terms of structure and reactivity are very similar for the two molecules. A description of the electronic wave function as having electrons in orbitals formed as linear combinations of all (in principle) atomic orbitals is also at variance with the chemical language of molecules being composed of atoms held together by bonds, where the bonds are formed by pairing unpaired electrons contained in atomic orbitals. Finally, when electron correlation is important (as is usually the case), the need to include many Slater determinants obscures the picture of electrons residing in orbitals.

There is an equivalent way of generating solutions to the electronic Schrödinger equation that conceptually is much closer to the experimentalist’s language, known as *Valence Bond* (VB) theory.¹ We will start by illustrating the concepts for the H_2 molecule, and note how it differs from MO methods.

7.1 Classical Valence Bond Theory

A single-determinant MO wave function for the H_2 molecule within a minimum basis consisting of a single s -function on each nucleus is given in eq. (7.1) (see also Section 4.3).

$$\begin{aligned}\Phi_0 &= \begin{vmatrix} \phi_1(1) & \bar{\phi}_1(1) \\ \phi_1(2) & \bar{\phi}_1(2) \end{vmatrix} \\ \phi_1 &= (\chi_A + \chi_B)\alpha \\ \bar{\phi}_1 &= (\chi_A + \chi_B)\beta\end{aligned}\tag{7.1}$$

We have here ignored the normalization constants. The Slater determinant may be expanded in AOs, as shown in eq. (7.2).

$$\begin{aligned}\Phi_0 &= \phi_1\bar{\phi}_1 - \bar{\phi}_1\phi_1 = (\phi_1\phi_1)[\alpha\beta - \beta\alpha] \\ \Phi_0 &= (\chi_A + \chi_B)(\chi_A + \chi_B)[\alpha\beta - \beta\alpha] \\ \Phi_0 &= (\chi_A\chi_A + \chi_B\chi_B + \chi_A\chi_B + \chi_B\chi_A)[\alpha\beta - \beta\alpha]\end{aligned}\tag{7.2}$$

This shows that the HF wave function consists of equal amounts of ionic ($\chi_A\chi_A$ and $\chi_B\chi_B$) and covalent ($\chi_A\chi_B$ and $\chi_B\chi_A$) terms. In the dissociation limit only the covalent terms are correct, but the single-determinant description does not allow the ratio of covalent to ionic terms to vary. In order to provide a correct description, a second determinant is necessary.

$$\begin{aligned}\Phi_1 &= \begin{vmatrix} \phi_2(1) & \bar{\phi}_2(1) \\ \phi_2(2) & \bar{\phi}_2(2) \end{vmatrix} \\ \phi_2 &= (\chi_A - \chi_B)\alpha \\ \bar{\phi}_2 &= (\chi_A - \chi_B)\beta \\ \Phi_1 &= (\chi_A\chi_A + \chi_B\chi_B - \chi_A\chi_B - \chi_B\chi_A)[\alpha\beta - \beta\alpha]\end{aligned}\tag{7.3}$$

By including the doubly excited determinant Φ_1 , built from the antibonding MO, the amounts of the covalent and ionic terms may be varied, and this is determined completely by the variational principle (eq. (4.20)).

$$\begin{aligned}\Psi_{CI} &= a_0\Phi_0 + a_1\Phi_1 \\ \Psi_{CI} &= \{(a_0 - a_1)(\chi_A\chi_B + \chi_B\chi_A) + (a_0 + a_1)(\chi_A\chi_A + \chi_B\chi_B)\}[\alpha\beta - \beta\alpha]\end{aligned}\tag{7.4}$$

This two-configurational CI wave function allows a qualitatively correct description of the H_2 molecule at all distances and in the dissociation limit, where the weights of the two configurations become equal.

The classical VB wave function, on the other hand, is built from the atomic fragments by coupling the unpaired electrons to form a bond. In the H_2 case, the two electrons are coupled into a singlet pair, properly antisymmetrized. The simplest VB description, known as a *Heitler–London* (HL) function, includes only the two covalent terms in the HF wave function.

$$\Phi_{HL}^{cov} = (\chi_A\chi_B + \chi_B\chi_A)[\alpha\beta - \beta\alpha]\tag{7.5}$$

Just as the single-determinant MO wave function may be improved by including excited determinants, the simple VB-HL function may also be improved by adding

terms that correspond to higher energy configurations for the fragments, in this case ionic structures.

$$\Phi_{\text{HL}}^{\text{ion}} = (\chi_A \chi_A + \chi_B \chi_B) [\alpha\beta - \beta\alpha] \quad (7.6)$$

$$\Psi_{\text{HL}} = a_0 \Phi_{\text{HL}}^{\text{cov}} + a_1 \Phi_{\text{HL}}^{\text{ion}} \quad (7.7)$$

The final description, either in terms of a CI wave function written as a linear combination of two determinants build from delocalized MOs (eq. (7.4)), or as a VB wave function written in terms of two VB-HL structures composed of AOs (eq. (7.7)), is identical.

For the H_2 system, the amount of ionic HL structures determined by the variational principle is 44%, close to the MO-HF value of 50%. The need for including large amounts of ionic structures in the VB formalism is due to the fact that pure atomic orbitals are used.

Consider now a covalent VB function built from “atomic” orbitals that are allowed to distort from the pure atomic shape.

$$\begin{aligned} \Phi_{\text{CF}} &= (\phi_A \phi_B + \phi_B \phi_A) [\alpha\beta - \beta\alpha] \\ \phi_A &= \chi_A + c\chi_B \\ \phi_B &= \chi_B + c\chi_A \end{aligned} \quad (7.8)$$

Such a VB function is known as a *Coulson–Fischer* (CF) type. The c constant is fairly small (for H_2 , c is ~ 0.04), but by allowing the VB orbitals to adopt the optimum shape, the need for ionic VB structures is strongly reduced. Note that the two VB orbitals in eq. (7.8) are not orthogonal – the overlap is given by eq. (7.9).

$$\begin{aligned} \langle \phi_A | \phi_B \rangle &= (1 + c^2) \langle \chi_A | \chi_B \rangle + 2c(\langle \chi_A | \chi_A \rangle + \langle \chi_B | \chi_B \rangle) \\ \langle \phi_A | \phi_B \rangle &= (1 + c^2) S_{\text{AB}} + 4c \end{aligned} \quad (7.9)$$

Compared with the overlap of the undistorted atomic orbitals used in the HL wave function, which is just S_{AB} , it is seen that the overlap is increased (c is positive), i.e. the orbitals distort such that they overlap better in order to make a bond. Although the distortion is fairly small (a few percent), this effectively eliminates the need for including ionic VB terms. When c is variationally optimized, the MO-CI, VB-HL and VB-CF wave functions (eqs (7.4), (7.7) and (7.8)) are all completely equivalent. The MO approach incorporates the flexibility in terms of an “excited” determinant, the VB-HL in terms of “ionic” structures, and the VB-CF in terms of “distorted” atomic orbitals.

In the MO-CI language, the correct dissociation of a single bond requires the addition of a second doubly excited determinant to the wave function. The VB-CF wave function, on the other hand, dissociates smoothly to the correct limit, the VB-orbitals simply reverting to their pure atomic shapes, with the overlap disappearing.

7.2 Spin-Coupled Valence Bond Theory

The generalization of a Coulson–Fischer type wave function to the molecular case with an arbitrary-size basis set is known as *Spin-Coupled Valence Bond* (SCVB) theory.²

It is again instructive to compare with the traditional MO approach, taking the CH_4 molecule as an example. The MO single-determinant description (RHF, which is identical to UHF near the equilibrium geometry) of the valence orbitals is in terms of four delocalized orbitals, each occupied by two electrons with opposite spin. The C—H

bonding is described by four different, orthogonal molecular orbitals, each expanded in a set of AOs.

$$\Phi_{\text{valence-MO}}^{\text{CH}_4} = \mathbf{A}[\phi_1\bar{\phi}_1\phi_2\bar{\phi}_2\phi_3\bar{\phi}_3\phi_4\bar{\phi}_4]$$

$$\phi_i = \sum_{\alpha=1}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} \quad (7.10)$$

Here \mathbf{A} is the usual antisymmetrizer (eq. (3.21)) and a bar above a MO indicates that the electron has a β spin function, no bar indicates an α spin function.

The SCVB description, on the other hand, considers the four bonds in CH_4 as arising from coupling of a single electron at each of the four hydrogen atoms with a single unpaired electron at the carbon atom. Since the ground state of the carbon atom is a triplet, corresponding to the electron configuration $1s^2 2s^2 2p^2$, the first step is formation of four equivalent “hybrid” orbitals by mixing three parts p -function with one part s -function, generating four equivalent “ sp^3 -hybrid” orbitals. Each of these singly occupied hybrid orbitals can then couple with a hydrogen atom to form four equivalent C—H bonds. The electron spins are coupled such that the total spin is a singlet, which can be done in several different ways. The coupling of four electrons to a total singlet state, for example, can be done either by coupling two electrons in a pair to a singlet, and then coupling two singlet pairs, or by first coupling two electrons in a pair to a triplet, and subsequently coupling two triplet pairs to an overall singlet.




Figure 7.1 Two possible schemes for coupling four electrons to an overall singlet

The $\Theta_{S,i}^N$ symbol is used to designate the i th combination of spin functions coupling N electrons to give an overall spin of S , and there are f_S^N number of ways of doing this. The value of f_S^N is given by eq. (7.11).

$$f_S^N = \frac{(2S+1)N!}{(\frac{1}{2}N+S+1)!(\frac{1}{2}N-S)!} \quad (7.11)$$

For a singlet wave function ($S = 0$), the number of coupling schemes for N electrons is given in Table 7.1.

Table 7.1 Number of possible spin coupling schemes for achieving an overall singlet state

N	f_0^N
2	1
4	2
6	5
8	14
10	42
12	132
14	429

For the eight valence electrons in CH_4 there are 14 possible spin couplings resulting in an overall singlet state. The full SCVB function may be written (again neglecting normalization) as in eq. (7.12).

$$\Phi_{\text{valence-SCVB}}^{\text{CH}_4} = \sum_{i=1}^{14} \alpha_i \mathbf{A} \{ [\phi_1 \phi_2 \phi_3 \phi_4 \phi_5 \phi_6 \phi_7 \phi_8] \Theta_{0,i}^N \}$$

$$\phi_i = \sum_{\alpha=1}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha}$$
(7.12)

There are now eight different spatial orbitals, ϕ_i , four of which are essentially carbon sp^3 -hybrid orbitals, with the other four being close to atomic hydrogen s-orbitals. The expansion of each of the VB-orbitals in terms of *all* the basis functions located on *all* the nuclei allows the orbitals to distort from the pure atomic shape. The SCVB wave function is variationally optimized, both with respect to the VB-orbital coefficients $c_{\alpha i}$ and the spin coupling coefficients α_i . The result is that a complete set of optimum “distorted” atomic orbitals is determined together with the weight of the different spin couplings. Each spin coupling term (in the so-called *Rumer* basis) is closely related to the concept of a resonance structure used in organic chemistry textbooks. An SCVB calculation of CH_4 gives as a result that one of the spin coupling schemes completely dominates the wave function, namely that corresponding to the electron pair in each of the C—H bonds being singlet coupled. This is the quantum mechanical analogue of the graphical representation of CH_4 shown in Figure 7.2.

Each of the lines represents a singlet-coupled electron pair between two orbitals that strongly overlap to form a bond, and the drawing in Figure 7.2 is the only important “resonance” form.




Figure 7.2 A representation of the dominating spin coupling in CH_4

Consider now the π -system in benzene. The MO approach will generate linear combinations of the atomic p-orbitals, producing six π -orbitals delocalized over the whole molecule with four different orbital energies (two sets of degenerate orbitals).

The stability of benzene can be attributed to the large gap between the HOMO and LUMO orbitals.

A SCVB calculation considering only the coupling of the six π -electrons, gives a somewhat different picture. The VB π -orbitals are strongly localized on each carbon, resembling p-orbitals that are slightly distorted in the direction of the nearest neighbour atoms. It is now found that five spin coupling combinations are important, these




Figure 7.3 Molecular orbital energies in benzene



Figure 7.4 Representations of important spin coupling schemes in benzene

are shown in Figure 7.4, where a bold line indicates two electrons coupled into a singlet pair.

Each of the two first VB structures contributes $\sim 40\%$ to the wave function, and each of the remaining three contributes $\sim 6\%$.³ The stability of benzene in the SCVB picture is due to *resonance* between these VB structures. It is furthermore straightforward to calculate the resonance energy by comparing the full SCVB energy with that calculated from a VB wave function omitting certain spin coupling functions.

The MO wave function for CH_4 may be improved by adding configurations corresponding to excited determinants, i.e. replacing occupied MOs with virtual MOs. Allowing all excitations in the minimal basis valence space and performing the full optimization corresponds to an [8,8]-CASSCF wave function (Section 4.6). Similarly, the SCVB wave function in eq. (7.12) may be improved by adding ionic VB structures such as CH_3^-/H^+ and CH_3^+/H^- , and this corresponds to exciting an electron from one of the singly occupied VB orbitals into another VB orbital, thereby making it doubly occupied. The importance of these excited/ionic terms can again be determined by the variational principle. If all such ionic terms are included, the fully optimized SCVB+CI wave function is for all practical purposes identical to that obtained by the MO-CASSCF approach (the only difference is a possible slight difference in the description of the carbon 1s-core orbital).

Both types of wave function provide essentially the same total energy, and thus include the same amount of electron correlation. The MO-CASSCF wave function attributes the electron correlation to interaction of 1764 configurations, the Hartree–Fock reference and 1763 excited configurations, with each of the 1763 configurations providing only a small amount of the correlation energy. The SCVB wave function (which includes only one resonance structure), however, contains 90+ % of the correlation energy, and only a few percent is attributed to “excited” structures. The ability of SCVB wave functions to include electron correlation is due to the fact that the VB orbitals are strongly localized and, since they are occupied by only one electron, they have the built-in feature of electrons avoiding each other. In a sense, an SCVB wave function is *the* best wave function that can be constructed in terms of *prod-*

ucts of *spatial* orbitals. By allowing the orbitals to become non-orthogonal, the large majority (80–90%) of what is called electron correlation in an MO approach can be included in a *single*-determinant wave function composed of spatial orbitals, multiplied by proper spin coupling functions.

There are a number of technical complications associated with optimizing the SCVB wave function due to the non-orthogonal orbitals. The MO-CI or MO-CASSCF approaches simplify considerably owing to the orthogonality of the MOs, and thereby also of the Slater determinants. Computationally, the optimization of an SCVB wave function, where N electrons are coupled in all possible ways, is similar to that required for constructing an $[N,N]$ -CASSCF wave function. This effectively limits the size of SCVB wave functions to coupling of 12–16 electrons. The actual optimization of the wave function is usually done by a second-order expansion of the energy in terms of orbital and spin coupling coefficients, and employing a Newton–Raphson type scheme, analogously to MCSCF methods (Section 4.6). The non-orthogonal orbitals have the disadvantage that it is difficult to add dynamical correlation on top of an SCVB wave function by perturbation or coupled cluster theory, although (non-orthogonal) CI methods are straightforward. SCVB+CI approaches may also be used to describe excited states, analogously to MO-CI methods.

It should be emphasized again that the results obtained from an $[N,N]$ -CASSCF and a corresponding N -electron SCVB wave function (or SCVB+CI and MRCI) are virtually identical. The difference is in the way the results can be analyzed. Molecules in the SCVB picture are composed of atoms held together by bonds, where bonds are formed by (singlet) coupling of the electron spins between (two) overlapping orbitals. These orbitals are strongly localized, usually on a single atom, and are basically atomic orbitals slightly distorted by the presence of the other atoms in the molecule. The VB description of a bond as the result of two overlapping orbitals is in contrast to the MO approach where a bond between two atoms arises as a sum over (small) contributions from many delocalized molecular orbitals. Furthermore, the weight of the different ways spin couplings in an SCVB wave function carries a direct analogy with chemical concepts such as “resonance” structures.

The SCVB method is a valuable tool for providing *insight* into the problem. This is to a certain extent also possible from an MO type wave function by localizing the orbitals or by analyzing the natural orbitals (see Sections 9.4 and 9.5 for details). However, there is no unique method for producing localized orbitals, and different methods may give different orbitals. Natural orbitals are analogous to canonical orbitals delocalized over the whole molecule. The SCVB orbitals, in contrast, are uniquely determined by the variational procedure, and there is no freedom to further transforming them by making linear combinations without destroying the variational property.

The primary feature of SCVB is the use of non-orthogonal orbitals, which allows a much more compact representation of the wave function. An MO-CI wave function of a certain quality may involve many thousands of Slater determinants, while a similar-quality VB wave function may be written as only a handful of “resonating” VB structures. Furthermore, the VB orbitals, and spin couplings, of a C—H bond in say propane and butane are very similar, in contrast to the vastly different MO descriptions of the two systems. The VB picture is thus much closer to the traditional descriptive language used with molecules composed of functional groups. The widespread availability of

programs for performing CASSCF calculations, and the fact that CASSCF calculations are computationally more efficient owing to the orthogonality of the MOs, have prompted developments of schemes for transforming CASSCF wave functions to VB structures, denoted CASVB.³ A corresponding procedure using orthogonal orbitals (which introduce large weights of ionic structures) has also been reported.⁴

7.3 Generalized Valence Bond Theory

The SCVB wave function allows all possible spin couplings to take place and has no restrictions on the form of the orbitals. The *Generalized Valence Bond* (GVB) method can be considered as a reduced version of the full problem where only certain subsets of spin couplings are allowed.⁵ For a typical case of a singlet system, the GVB method has two (non-orthogonal) orbitals assigned to each bond, and each pair of electrons in a bond are required to couple to a singlet pair. The coupling of such singlet pairs will then give the overall singlet spin state. This is known as *Perfect Pairing* (PP), and is one of the many possible spin coupling schemes, and such two-electron two-orbital pairs are called *geminal pairs*. Just as an orbital is a wave function for one electron, a geminal is a wave function for two electrons. In order to reduce the computational problem, the *Strong Orthogonality* (SO) condition is normally imposed on the GVB wave function. This means that orbitals belonging to different pairs are required to be orthogonal. While the perfect pairing coupling typically is the largest contribution to the full SCVB wave function, the strong orthogonality constraint is often a quite poor approximation, and may lead to artefacts. For diazomethane, for example, the SCVB wave function is dominated (91%) by the PP coupling, leading to the conclusion that the molecule has essentially normal C=N and N=N π -bonds, perpendicular to the plane defined by the CH₂ moiety.⁶ Taking into account also the in-plane bonding, this suggests that diazomethane is best described with a triple bond between the two nitrogens, thereby making the central nitrogen “hypervalent”, as illustrated in Figure 7.5.




Figure 7.5 A representation of the SCVB wave function for diazomethane

There are strong overlaps between the VB orbitals, the *smallest* overlap (between the carbon and terminal nitrogen) is ~ 0.4 , and that between the two orbitals on the central nitrogen is ~ 0.9 . The GVB-SOPP approach, however, forces these geminal pairs to be orthogonal, leading to the conclusion that the electronic structure of diazomethane has a very strong diradical nature, as illustrated in Figure 7.6.




Figure 7.6 A representation of the GVB wave function for diazomethane

References

1. S. Shaik, P. C. Hiberty, *Rev. Comp. Chem.*, **20** (2004), 1.
2. D. L. Cooper, J. Gerratt, M. Raimondi, *Chem. Rev.*, **91** (1991), 929; J. Gerratt, D. L. Cooper, P. B. Karadakov, M. Raimondi, *Chem. Soc. Rev.*, **26** (1997), 87.
3. D. L. Cooper, T. Thorsteinsson, J. Gerratt, *Int. J. Quant. Chem.*, **65** (1997), 439.
4. K. Hirao, H. Nakano, K. Nakayama, M. Dupuis, *J. Chem. Phys.*, **105** (1996), 9227.
5. W. A. Goddard III, L. B. Harding, *Ann. Rev. Phys. Chem.*, **29** (1978), 363.
6. D. L. Cooper, J. Gerratt, M. Raimondi, S. C. Wright, *Chem. Phys. Lett.*, **138** (1987), 296.

8 Relativistic Methods

The central theme in relativity is that the speed of light, c , is constant in all inertia frames (coordinate systems that move with respect to each other). Augmented with the requirement that physical laws should be identical in such frames, this has as a consequence that time and space coordinates become “equivalent”. A relativistic description of a particle thus requires four coordinates, three space and one time coordinate.¹ The latter is usually multiplied by c to have units identical to the space variables.

A change between different coordinate systems can be described by a *Lorentz* transformation, which may mix space and time coordinates. The postulate that physical laws should be identical in all coordinate systems is equivalent to the requirement that equations describing the physics must be invariant (unchanged) to a Lorentz transformation. Considering the time-dependent Schrödinger equation (8.1), it is clear that it is not Lorentz invariant since the derivative with respect to space coordinates is of second order, but the time derivative is only first order. The fundamental structure of the Schrödinger equation is therefore not relativistically correct.

$$\left[-\frac{1}{2m} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \right) + \mathbf{V} \right] \Psi = i \frac{\partial \Psi}{\partial t} \quad (8.1)$$

For use below, we have elected here to explicitly write the electron mass as m , although it is equal to one in atomic units.

One of the consequences of the constant speed of light is that the mass of a particle, which moves at a substantial fraction of c , increases over the rest mass m_0 .

$$m = m_0 \left(\sqrt{1 - \frac{v^2}{c^2}} \right)^{-1} \quad (8.2)$$

The energy of a 1s-electron in a hydrogen-like system (one nucleus and one electron) is $-Z^2/2$, and classically this is equal to minus the kinetic energy, $1/2mv^2$, owing to the virial theorem ($E = -T = 1/2V$). In atomic units ($m = 1$) the classical velocity of a 1s-electron is thus Z . The speed of light in atomic units is 137.036, and it is clear that relativistic effects cannot be neglected for the core electrons in heavy nuclei. For atoms with large Z , the 1s-electrons are relativistic and thus heavier, which has the effect that

the 1s-orbital shrinks in size, by the same factor as the mass increases (eq. (8.2)). In order to maintain orthogonality, the higher s-orbitals also contract. This provides a more effective screening of the nuclear charge for the higher angular momentum orbitals, which consequently increase in size. For p-orbitals the spin-orbit interaction, which mixes s- and p-orbitals, counteracts the inflation. The net effect is that p-orbitals are relatively unaffected in size, while d- and f-orbitals become larger and more diffuse.

In terms of total energy, the relativistic correction becomes comparable to the correlation energy already for $Z \sim 10$, while it becomes comparable to the exchange energy for $Z \sim 50$. Since the majority of the relativistic effects are concentrated in the core orbitals, there is a large error cancellation for molecular properties. Relativistic effects for geometries and energetics are normally negligible for the first three rows in the periodic table (up to Kr, $Z = 36$, corresponding to a “mass correction” of 1.04), the fourth row represents an intermediate case, while relativistic corrections are necessary for the fifth and sixth rows, and for lanthanide/actinide metals. For effects involving electron spin (e.g. spin-orbit coupling), which are purely relativistic in origin, there is no non-relativistic counterpart, and the “relativistic correction” is of course everything.

Although an in-depth treatment of relativistic effects is outside the scope of this book, it may be instructive to point out some of the features and problems in a relativistic quantum description of atoms and molecules. Furthermore, we will require some operators derived from a relativistic treatment for calculating molecular properties in Chapter 10.

8.1 The Dirac Equation

For a free electron, Dirac proposed that the (time-dependent) Schrödinger equation should be replaced by eq. (8.3).

$$[c\boldsymbol{\alpha} \cdot \mathbf{p} + \beta mc^2]\Psi = i\frac{\partial \Psi}{\partial t} \quad (8.3)$$

Here $\boldsymbol{\alpha}$ and β are 4×4 matrices, $\boldsymbol{\alpha}$ is written in terms of the three Pauli 2×2 spin matrices $\boldsymbol{\sigma}$, and β in term of a 2×2 unit matrix \mathbf{I} .

$$\begin{aligned} \boldsymbol{\alpha}_{x,y,z} &= \begin{pmatrix} 0 & \boldsymbol{\sigma}_{x,y,z} \\ \boldsymbol{\sigma}_{x,y,z} & 0 \end{pmatrix} & \beta &= \begin{pmatrix} \mathbf{I} & 0 \\ 0 & \mathbf{I} \end{pmatrix} \\ \boldsymbol{\sigma}_x &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} & \boldsymbol{\sigma}_y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} & \boldsymbol{\sigma}_z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} & \mathbf{I} &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \end{aligned} \quad (8.4)$$

Except for a factor of $1/2$, the $\boldsymbol{\sigma}_{x,y,z}$ matrices can be viewed as representations of the \mathbf{s}_x , \mathbf{s}_y and \mathbf{s}_z spin operators, respectively, when the α and β spin functions are taken as (1,0) and (0,1) vectors.

$$\begin{aligned} \mathbf{s}_z &= \frac{1}{2} \boldsymbol{\sigma}_z \\ \mathbf{s}_z \begin{pmatrix} 1 \\ 0 \end{pmatrix} &= \frac{1}{2} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ \mathbf{s}_z \begin{pmatrix} 0 \\ 1 \end{pmatrix} &= -\frac{1}{2} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \end{aligned} \quad (8.5)$$

The α function is an eigenfunction of the s_z operator with an eigenvalue of $1/2$, and the β function similarly has an eigenvalue of $-1/2$.

The Dirac equation is of same order in all variables (space and time), since the momentum operator \mathbf{p} ($= -i\nabla$) involves a first-order differentiation with respect to the space variables. It should be noted that the free electron rest energy in eq. (8.3) is mc^2 , equal to 0.511 MeV, while this situation is defined as zero in the non-relativistic case. The zero point of the energy scale is therefore shifted by 5.11×10^5 eV, a large amount compared with the binding energy of 13.6 eV for a hydrogen atom. The two (relativistic and non-relativistic) energy scales may be aligned by subtracting the electron rest energy, which corresponds to replacing the β matrix in eq. (8.3) by β' .

$$\beta' = \begin{pmatrix} 0 & 0 \\ 0 & 2\mathbf{I} \end{pmatrix} \quad (8.6)$$

The Dirac equation corresponds to satisfying the requirements of *special relativity* in connection with the quantum behaviour of the electron. Special relativity considers only systems that move with a constant velocity with respect to each other, which can hardly be considered a good approximation for the movement of an electron around a nucleus. A relativistic treatment of accelerated systems is described by *general relativity*, which is a gravitational theory. For atomic systems, however, the gravitational interaction between electrons and nuclei (or between electrons) is insignificant compared with the electrostatic interaction. Furthermore, a consistent theory describing the quantum aspects of gravitation has not yet been developed.

The Dirac equation is four-dimensional, and the relativistic wave function consequently contains *four components*. Two of the degrees of freedom are accounted for by assigning an intrinsic magnetic moment (spin), while the other two are interpreted as two different particles, electron and positron. The positronic solutions show up as a continuum of “negative” energy states, having energies below $-2mc^2$, as illustrated in Figure 8.1. Note that the spacing between bound states has been exaggerated, as the binding energy is of the order of eV while $2mc^2$ is of the order of MeV.

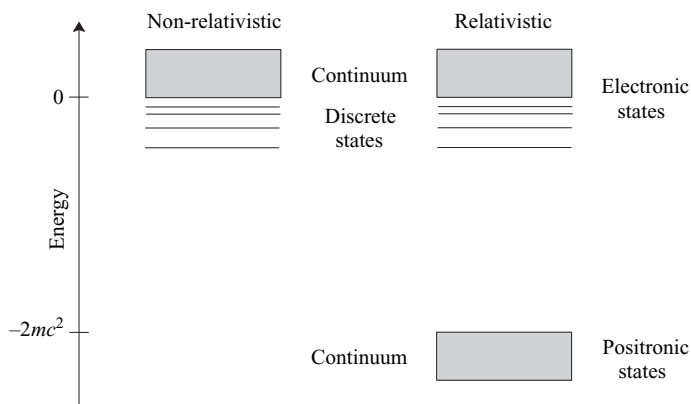


Figure 8.1 Non-relativistic and relativistic solutions

It is conventional to write the relativistic wave function as in eq. (8.7).

$$\Psi = \begin{pmatrix} \Psi_{L\alpha} \\ \Psi_{L\beta} \\ \Psi_{S\alpha} \\ \Psi_{S\beta} \end{pmatrix} \quad (8.7)$$

Here Ψ_L and Ψ_S are the *large* and *small* components of the wave function, and α and β indicate the usual spin functions. Note that the spatial parts of $\Psi_{L\alpha}/\Psi_{L\beta}$, and $\Psi_{S\alpha}/\Psi_{S\beta}$, are not necessarily identical. For electrons, the large component reduces to the solutions of the Schrödinger equation when $c \rightarrow \infty$ (the non-relativistic limit), and the small component disappears. The small component of the electronic wave function corresponds to a coupling with the positronic states.

8.2 Connections Between the Dirac and Schrödinger Equations

8.2.1 Including electric potentials

In the presence of an electric potential \mathbf{V} (e.g. from nuclei), the time-independent Dirac equation may be written as in eq. (8.8), where we have again explicitly indicated the electron mass.

$$[c\boldsymbol{\alpha} \cdot \mathbf{p} + \beta' mc^2 + \mathbf{V}]\Psi = E\Psi \quad (8.8)$$

Since $\boldsymbol{\alpha}$ and β' are block matrices in terms of $\boldsymbol{\sigma}$ and \mathbf{I} , eq. (8.8) can be factored out in two equations.

$$\begin{aligned} c(\boldsymbol{\sigma} \cdot \mathbf{p})\Psi_S + \mathbf{V}\Psi_L &= E\Psi_L \\ c(\boldsymbol{\sigma} \cdot \mathbf{p})\Psi_L + (-2mc^2 + \mathbf{V})\Psi_S &= E\Psi_S \end{aligned} \quad (8.9)$$

Here Ψ_L and Ψ_S are (large and small) two-component wave functions that include the α and β spin functions. The latter equation can be solved for Ψ_S .

$$\Psi_S = (E + 2mc^2 - \mathbf{V})^{-1} c(\boldsymbol{\sigma} \cdot \mathbf{p})\Psi_L \quad (8.10)$$

The inverse quantity can be factorized as in eq. (8.11).

$$\begin{aligned} (E + 2mc^2 - \mathbf{V})^{-1} &= (2mc^2)^{-1} \left(1 + \frac{E - \mathbf{V}}{2mc^2} \right)^{-1} = (2mc^2)^{-1} \mathbf{K} \\ \mathbf{K} &= \left(1 + \frac{E - \mathbf{V}}{2mc^2} \right)^{-1} \end{aligned} \quad (8.11)$$

Eq. (8.10) may then be written as in eq. (8.12).

$$\Psi_S = \mathbf{K} \frac{\boldsymbol{\sigma} \cdot \mathbf{p}}{2mc} \Psi_L \quad (8.12)$$

The top equation in (8.9) then becomes eq. (8.13).

$$\left[\frac{1}{2m} (\boldsymbol{\sigma} \cdot \mathbf{p}) \mathbf{K} (\boldsymbol{\sigma} \cdot \mathbf{p}) + (\mathbf{V} - E) \right] \Psi_L = 0 \quad (8.13)$$

In the non-relativistic limit ($c \rightarrow \infty$) the \mathbf{K} factor is 1, and the first term becomes $(\boldsymbol{\sigma} \cdot \mathbf{p})(\boldsymbol{\sigma} \cdot \mathbf{p})$. Using the vector identity $(\boldsymbol{\sigma} \cdot \mathbf{p})(\boldsymbol{\sigma} \cdot \mathbf{p}) = \mathbf{p} \cdot \mathbf{p} + i\boldsymbol{\sigma}(\mathbf{p} \times \mathbf{p})$, this gives the non-relativistic kinetic energy $\mathbf{p}^2/2m$, since the vector product of any vector with itself is zero ($\mathbf{p} \times \mathbf{p} = 0$). The equation for the large component therefore reduces to the Schrödinger equation.

$$\left[\frac{\mathbf{p}^2}{2m} + \mathbf{V} \right] \Psi_L = E \Psi_L \quad (8.14)$$

The electron spin is still present in eq. (8.14), since Ψ_L is a two-component wave function, but this can trivially be separated out since the operators do not contain any spin dependence.

In the non-relativistic limit the small component of the wave function is given by eq. (8.15).

$$\Psi_s = \frac{\boldsymbol{\sigma} \cdot \mathbf{p}}{2mc} \Psi_L \quad (8.15)$$

For a hydrogenic wave function ($\Psi_L \approx e^{-Zr}$), this gives eq. (8.16) in atomic units (setting $m = 1$).

$$\Psi_s \approx \frac{Z}{2c} \Psi_L \quad (8.16)$$

For a hydrogen atom the small component accounts for only $\sim 0.4\%$ of the total wave function and $10^{-3}\%$ of the electron density, but for a uranium 1s-electron it is a third of the wave function and $\sim 10\%$ of the density.

We may obtain relativistic corrections by expanding the \mathbf{K} factor in eq. (8.11).

$$\mathbf{K} = \left(1 + \frac{E - \mathbf{V}}{2mc^2} \right)^{-1} \approx 1 - \frac{E - \mathbf{V}}{2mc^2} + \dots \quad (8.17)$$

This is only valid when $E - \mathbf{V} \ll 2mc^2$, however all atoms have a region close to the nucleus where this is not fulfilled (since $\mathbf{V} \rightarrow -\infty$ for $\mathbf{r} \rightarrow 0$). Inserting (8.17) in (8.13), assuming a Coulomb potential $-Z/r$ (i.e. \mathbf{V} is the attraction to a nucleus), gives after renormalization of the (large component) wave function and some rearrangement the terms shown in eq. (8.18).

$$\left[\frac{\mathbf{p}^2}{2m} + \mathbf{V} - \frac{\mathbf{p}^4}{8m^3c^2} + \frac{Z\mathbf{s} \cdot \mathbf{I}}{2m^2c^2r^3} + \frac{Z\pi\delta(\mathbf{r})}{2m^2c^2} \right] \Psi_L = E \Psi_L \quad (8.18)$$

Eq. (8.18) is called the *Pauli equation*. The first two terms are the usual non-relativistic kinetic and potential energy operators, the \mathbf{p}^4 term is called the *mass-velocity* correction, and is due to the dependence of the electron mass on the velocity. The next is the *spin-orbit* term (\mathbf{s} is the electron spin and \mathbf{I} is the angular momentum operator $\mathbf{r} \times \mathbf{p}$), which corresponds to an interaction of the electron spin with the magnetic field generated by the movement of the electron. The last term involving the δ function is the *Darwin* correction, which corresponds to a correction that can be interpreted as the electron making a high-frequency oscillation around its mean position, sometimes referred to as *Zitterbewegung*. The mass-velocity and Darwin corrections are often collectively called the *scalar* relativistic

corrections. Since they have opposite signs, they do to a certain extent cancel each other.

Owing to the divergence of the \mathbf{K} expansion near the nuclei, the mass-velocity and Darwin corrections can only be used as first-order corrections. Inclusion of such operators in a variational sense will result in a collapse of the wave function. An alternative method is to partition eq. (8.11) as in eq. (8.19), which avoids the divergence near the nucleus.

$$(E + 2mc^2 - \mathbf{V})^{-1} = (2mc^2 - \mathbf{V})^{-1} \left(1 + \frac{E}{2mc^2 - \mathbf{V}} \right)^{-1} = (2mc^2 - \mathbf{V})^{-1} \mathbf{K}' \quad (8.19)$$

$$\mathbf{K}' = \left(1 + \frac{E}{2mc^2 - \mathbf{V}} \right)^{-1}$$

In contrast to eq. (8.17), the factor $E/(2mc^2 - \mathbf{V})$ is always much smaller than 1, and \mathbf{K}' may be expanded in powers of $E/(2mc^2 - \mathbf{V})$, analogously to eq. (8.17). Keeping only the zeroth-order term (i.e. setting $\mathbf{K}' = 1$) gives the *Zeroth-Order Regular Approximation* (ZORA) method, eq. (8.20).²

$$\left[\frac{c^2 \mathbf{p}^2}{2mc^2 - \mathbf{V}} + \frac{2c^2}{(2mc^2 - \mathbf{V})^2} - \frac{\mathbf{Zs} \cdot \mathbf{I}}{r^3} + \mathbf{V} \right] \Psi_L = E \Psi_L \quad (8.20)$$

Note that in this case the spin-orbit coupling is already included in zeroth order. Including the first-order term from an expansion of \mathbf{K}' defines the *First-Order Regular Approximation* (FORA) method. A disadvantage of these methods is that they are not gauge invariant.³

8.2.2 Including both electric and magnetic potentials

The presence of a magnetic field can be included in the so-called minimal coupling by addition of a *vector potential* \mathbf{A} to the momentum operator \mathbf{p} , forming a *generalized* momentum operator $\boldsymbol{\pi}$, which for an electron (charge of -1) is given by eq. (8.21).

$$\boldsymbol{\pi} = \mathbf{p} + \mathbf{A} \quad (8.21)$$

The magnetic field is defined as the curl of the vector potential.

$$\mathbf{B} = \nabla \times \mathbf{A} \quad (8.22)$$

For an external magnetic field, it is conventional to write the vector potential as in eq. (8.23).

$$\mathbf{A}(\mathbf{r}) = \frac{1}{2} \mathbf{B} \times (\mathbf{r} - \mathbf{R}_G) \quad (8.23)$$

Here \mathbf{R}_G is the *gauge origin*, i.e. the “zero” point for the vector potential. The gauge origin is often taken as the centre of mass for the system, but this is by no means unique. The results from an exact calculation will be independent of \mathbf{R}_G but, for approximate calculations, this is not guaranteed, and the results may thus depend on where the gauge origin is chosen. Such gauge-dependent properties are clearly undesirable, since different values can be generated by selecting different (arbitrary) gauge origins.

With the generalized momentum operator $\boldsymbol{\pi}$ replacing \mathbf{p} , the time-independent Dirac equation may be separated analogously to the procedure in Section 8.2.1 to give the equivalent of eq. (8.13).

$$\left[\frac{1}{2m} (\boldsymbol{\sigma} \cdot \boldsymbol{\pi}) \mathbf{K} (\boldsymbol{\sigma} \cdot \boldsymbol{\pi}) + (\mathbf{V} - E) \right] \Psi_L = 0 \quad (8.24)$$

Taking the non-relativistic limit corresponding to $\mathbf{K} = 1$ gives $(\boldsymbol{\sigma} \cdot \boldsymbol{\pi})(\boldsymbol{\sigma} \cdot \boldsymbol{\pi})$ for the first term. Using again the vector identity $(\boldsymbol{\sigma} \cdot \boldsymbol{\pi})(\boldsymbol{\sigma} \cdot \boldsymbol{\pi}) = \boldsymbol{\pi} \cdot \boldsymbol{\pi} + i\boldsymbol{\sigma}(\boldsymbol{\pi} \times \boldsymbol{\pi})$, this may be written as in eq. (8.25).

$$\left[\frac{1}{2m} (\boldsymbol{\pi} \cdot \boldsymbol{\pi} + i\boldsymbol{\sigma} \cdot (\boldsymbol{\pi} \times \boldsymbol{\pi})) + \mathbf{V} \right] \Psi_L = E \Psi_L \quad (8.25)$$

In contrast to the situation without a magnetic field, the latter vector product no longer disappears. The $\boldsymbol{\pi} \times \boldsymbol{\pi}$ term can be expanded by inserting the definition of $\boldsymbol{\pi}$ from eq. (8.21).

$$\begin{aligned} \boldsymbol{\pi} \times \boldsymbol{\pi} &= (\mathbf{p} + \mathbf{A}) \times (\mathbf{p} + \mathbf{A}) \\ &= \mathbf{p} \times \mathbf{p} + \mathbf{p} \times \mathbf{A} + \mathbf{A} \times \mathbf{p} + \mathbf{A} \times \mathbf{A} \end{aligned} \quad (8.26)$$

The first and last terms are zero (since $\mathbf{a} \times \mathbf{a} = 0$). With $\mathbf{p} = -i\nabla$ the other two terms yield eq. (8.27).

$$\begin{aligned} (\mathbf{p} \times \mathbf{A} + \mathbf{A} \times \mathbf{p})\Psi &= -i\nabla \times (\mathbf{A}\Psi) - i\mathbf{A} \times (\nabla\Psi) \\ &= -i(\nabla \times \mathbf{A})\Psi - i(\nabla\Psi) \times \mathbf{A} - i\mathbf{A} \times (\nabla\Psi) \end{aligned} \quad (8.27)$$

The two last terms cancel (since $\mathbf{a} \times \mathbf{b} = -\mathbf{b} \times \mathbf{a}$), and the curl of the vector potential is the magnetic field, eq. (8.22). The final result is given in eq. (8.28).

$$\left[\frac{\boldsymbol{\pi}^2}{2m} + \mathbf{V} + \frac{\boldsymbol{\sigma} \cdot \mathbf{B}}{2m} \right] \Psi_L = E \Psi_L \quad (8.28)$$

The $\boldsymbol{\sigma} \cdot \mathbf{B}$ term is called the (spin) *Zeeman* interaction, and represents the interaction of an (external) magnetic field with the intrinsic magnetic moment associated with the electron. As noted in eq. (8.5), $\boldsymbol{\sigma}$ represents the spin operator (except for a factor of $1/2$), and the $\boldsymbol{\sigma} \cdot \mathbf{B}/2m$ interaction can (in atomic units) also be written as $\mathbf{s} \cdot \mathbf{B}$, with \mathbf{s} being the electron spin operator. In a more refined treatment, by including quantum field corrections, it turns out that the electron magnetic moment is not exactly equal to the spin. It is conventional to write the interaction as $g_e \mu_B \mathbf{s} \cdot \mathbf{B}$ where the Bohr magneton $\mu_B (= e\hbar/2m)$ has a value of $1/2$ in atomic units and the electronic g-factor g_e is approximately equal to 2.0023 (the deviation from the value of 2 (exactly) is due to quantum field fluctuations).

Although electron spin is often said to arise from relativistic effects, the above shows that spin naturally arises in the non-relativistic limit of the Dirac equation. It may also be argued that electron spin is actually present in the non-relativistic case, as the kinetic energy operator $\mathbf{p}^2/2m$ is mathematically equivalent to $(\boldsymbol{\sigma} \cdot \mathbf{p})^2/2m$. If the kinetic energy is written as $(\boldsymbol{\sigma} \cdot \mathbf{p})^2/2m$ in the Schrödinger Hamiltonian, then electron spin is present in the non-relativistic case, although it would only have consequences in the presence of a magnetic field.

The Dirac equation automatically includes effects due to electron spin, while this must be introduced in a more or less *ad hoc* fashion in the Schrödinger equation (the Pauli principle). Furthermore, once the spin-orbit interaction is included, the total electron spin is no longer a “good” quantum number: an orbital no longer contains an integer number of α and β spin functions. The proper quantum number in relativistic theory is therefore the total angular momentum obtained by vector addition of the orbital and spin moments.

Turning now to the π^2 term in eq. (8.28), it can with the use of eq. (8.21) be expanded into eq. (8.29).

$$\pi^2 = (\mathbf{p} + \mathbf{A})^2 = \mathbf{p}^2 + \mathbf{p} \cdot \mathbf{A} + \mathbf{A} \cdot \mathbf{p} + \mathbf{A}^2 \quad (8.29)$$

The \mathbf{p}^2 gives the usual (non-relativistic) kinetic energy operator. Since $\mathbf{p} = -i\nabla$, the $\mathbf{p} \cdot \mathbf{A}$ term gives eq. (8.30).

$$(\mathbf{p} \cdot \mathbf{A})\Psi = -i(\nabla \cdot \mathbf{A})\Psi = -i\mathbf{A} \cdot (\nabla\Psi) - i\Psi(\nabla \cdot \mathbf{A}) \quad (8.30)$$

The *Coulomb gauge* is defined by $\nabla \cdot \mathbf{A} = 0$, and in this gauge we have $\mathbf{p} \cdot \mathbf{A} = \mathbf{A} \cdot \mathbf{p}$. The two terms involving \mathbf{A} in eq. (8.29) can be evaluated by inserting the expression for the vector potential (8.23).

$$\begin{aligned} \mathbf{A} \cdot \mathbf{p} &= \left(\frac{1}{2}\mathbf{B} \times (\mathbf{r} - \mathbf{R}_G)\right) \cdot \mathbf{p} \\ &= \frac{1}{2}\mathbf{B} \cdot (\mathbf{r} - \mathbf{R}_G) \times \mathbf{p} \\ &= \frac{1}{2}\mathbf{B} \cdot \mathbf{L}_G \\ \mathbf{A}^2 &= \left(\frac{1}{2}\mathbf{B} \times (\mathbf{r} - \mathbf{R}_G)\right) \cdot \left(\frac{1}{2}\mathbf{B} \times (\mathbf{r} - \mathbf{R}_G)\right) \\ &= \frac{1}{4}(\mathbf{B}^2 \cdot (\mathbf{r} - \mathbf{R}_G))^2 - (\mathbf{B} \cdot (\mathbf{r} - \mathbf{R}_G))^2 \end{aligned} \quad (8.31)$$

Here the vector identities $\mathbf{a} \times \mathbf{b} \cdot \mathbf{c} = \mathbf{a} \cdot \mathbf{b} \times \mathbf{c}$ and $(\mathbf{a} \times \mathbf{b}) \cdot (\mathbf{c} \times \mathbf{d}) = (\mathbf{a} \cdot \mathbf{c})(\mathbf{b} \cdot \mathbf{d}) - (\mathbf{a} \cdot \mathbf{d})(\mathbf{b} \cdot \mathbf{c})$ have been used. In addition to the Zeeman term for electron spin (eq. (8.28)), the presence of a magnetic field introduces two new terms, being linear and quadratic in the field. The linear operator represents an (orbital) Zeeman type interaction of the magnetic field with the magnetic moment generated by the movement of the electron, as described by the angular momentum operator \mathbf{L}_G , while the quadratic term gives rise to a component of the magnetizability in a perturbation treatment, as discussed in Section 10.7.6.

8.3 Many-Particle Systems

A fully relativistic treatment of more than one particle would have to start from a full QED treatment of the system (Chapter 1), and perform a perturbation expansion in terms of the radiation frequency. There is no universally accepted way of doing this, and a full relativistic many-body equation has not yet been developed. For many-particle systems it is assumed that each electron can be described by a Dirac operator ($c\boldsymbol{\alpha} \cdot \boldsymbol{\pi} + \beta mc^2$) and the many-electron operator is a sum of such terms, in analogy with the kinetic energy in non-relativistic theory. Furthermore, potential energy operators are added to form a total operator equivalent to the Hamiltonian operator in non-relativistic theory. Since this approach gives results that agree with experiments, the assumptions appear justified.

The Dirac operator incorporates relativistic effects for the kinetic energy. In order to describe atomic and molecular systems, the potential energy operator must also be modified. In non-relativistic theory, the potential energy is given by the Coulomb operator.

$$\mathbf{V}(\mathbf{r}_{12}) = \frac{q_1 q_2}{r_{12}} \quad (8.32)$$

According to this equation, the interaction between two charged particles depends only on the distance between them, *but not on time*. This cannot be correct when relativity is considered, as it implies that the attraction/repulsion between two particles occurs *instantly* over the distance r_{12} , violating the fundamental relativistic principle that nothing can move faster than the speed of light. The interaction between distant particles must be “later” than between particles that are close, and the potential is consequently “retarded” (delayed). The relativistic interaction requires a description, *Quantum ElectroDynamics* (QED), which involves exchange of photons between charged particles. The photons travel at the speed of light and carry the information equivalent to the classical Coulomb interaction. The relativistic potential energy operator becomes complicated and cannot be written in closed form. For actual calculations, it may be expanded in a Taylor series in $1/c$ and, for chemical purposes, it is normally only necessary to include terms up to $1/c^2$. In this approximation, the potential energy operator for the electron–electron repulsion is given by eq. (8.33).

$$\mathbf{V}_{ee}(\mathbf{r}_{12}) = \frac{1}{r_{12}} - \frac{1}{r_{12}} \left[\boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2 + \frac{(\boldsymbol{\alpha}_1 \times \mathbf{r}_{12})(\boldsymbol{\alpha}_2 \times \mathbf{r}_{12})}{r_{12}^2} \right] \quad (8.33)$$

Note that the subscript on the $\boldsymbol{\alpha}$ -matrices refers to the particle, and $\boldsymbol{\alpha}$ here includes all of the α_x , α_y and α_z components in eq. (8.4). The first correction term in the square bracket is called the *Gaunt* interaction, and the whole term in the square bracket is the *Breit* interaction. The Dirac matrices appear since they represent the velocity operators in a relativistic description. The Gaunt term is a magnetic interaction (spin) while the other term represents a retardation effect. Equation (8.33) is more often written in the form shown in eq. (8.34).

$$\mathbf{V}_{ee}^{\text{Coulomb-Breit}}(\mathbf{r}_{12}) = \frac{1}{r_{12}} - \frac{1}{2r_{12}} \left[\boldsymbol{\alpha}_1 \cdot \boldsymbol{\alpha}_2 + \frac{(\boldsymbol{\alpha}_1 \cdot \mathbf{r}_{12})(\boldsymbol{\alpha}_2 \cdot \mathbf{r}_{12})}{r_{12}^2} \right] \quad (8.34)$$

Relativistic corrections to the nuclear–electron attraction (\mathbf{V}_{ne}) are of order $1/c^3$ (owing to the much smaller velocity of the nuclei) and are normally neglected.

An expansion in powers of $1/c$ (or, equivalently, in powers of the fine-structure constant $\alpha = 1/c$ in atomic units) is a standard approach for deriving relativistic correction terms. Taking into account electron (\mathbf{s}) and nuclear spins (\mathbf{I}), and indicating explicitly an external electric potential by means of the field ($\mathbf{F} = -\nabla\phi$, or $-\nabla\phi - \partial\mathbf{A}/\partial t$ if time dependent), an expansion up to order $1/c^2$ of the Dirac Hamiltonian including the Coulomb–Breit potential gives the following set of operators,⁴ where the QED correction to the electron spin has been introduced by means of the $g_e\mu_B$ factor. Note that many of these operators arise from the minimal coupling of the magnetic field via the generalized momentum operator, as discussed in more detail in Section 10.10.7.

One electron operators:

$$\begin{aligned}
 \mathbf{H}_e^{\text{Zeeman}} &= g_e \mu_B \sum_{i=1}^{N_{\text{elec}}} \left[\mathbf{s}_i \cdot \mathbf{B}_i - \frac{1}{2mc^2} (\mathbf{s}_i \cdot \mathbf{B}_i) \boldsymbol{\pi}_i^2 \right] \\
 \mathbf{H}_e^{\text{mv}} &= -\frac{1}{8m^3 c^2} \sum_{i=1}^{N_{\text{elec}}} \boldsymbol{\pi}_i^4 \\
 \mathbf{H}_e^{\text{SO}} &= -\frac{g_e \mu_B}{4mc^2} \sum_{i=1}^{N_{\text{elec}}} [\mathbf{s}_i \cdot \boldsymbol{\pi}_i \times \mathbf{F}_i - \mathbf{s}_i \cdot \mathbf{F}_i \times \boldsymbol{\pi}_i] \\
 \mathbf{H}_e^{\text{Darwin}} &= -\frac{1}{8m^2 c^2} \sum_{i=1}^{N_{\text{elec}}} \nabla \cdot \mathbf{F}_i
 \end{aligned} \tag{8.35}$$

Here \mathbf{F}_i and \mathbf{B}_i indicate the (electric and magnetic) fields at the position of particle i . $\mathbf{H}_e^{\text{Zeeman}}$ has the $\mathbf{s} \cdot \mathbf{B}$ term from eq. (8.28) and a relativistic correction, and \mathbf{H}_e^{mv} is the mass–velocity correction, as is also present in eq. (8.18). \mathbf{H}_e^{SO} and $\mathbf{H}_e^{\text{Darwin}}$ are spin–orbit and Darwin type correction with respect to an external electric field. It should be noted that the generalized momentum operator contains magnetic fields via the vector potential $\boldsymbol{\pi} = \mathbf{p} + \mathbf{A}$, and eq. (8.35) therefore implicitly includes higher order effects.

Two electron operators:

$$\begin{aligned}
 \mathbf{H}_{ee}^{\text{SO}} &= -\frac{g_e \mu_B}{2mc^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j \neq i}^{N_{\text{elec}}} \frac{\mathbf{s}_i \cdot (\mathbf{r}_{ij} \times \mathbf{p}_i)}{r_{ij}^3} \\
 \mathbf{H}_{ee}^{\text{SOO}} &= -\frac{g_e \mu_B}{mc^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j \neq i}^{N_{\text{elec}}} \frac{\mathbf{s}_i \cdot (\mathbf{r}_{ij} \times \mathbf{p}_j)}{r_{ij}^3} \\
 \mathbf{H}_{ee}^{\text{SS}} &= \frac{g_e^2 \mu_B^2}{2c^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j \neq i}^{N_{\text{elec}}} \left(\frac{\mathbf{s}_i \cdot \mathbf{s}_j}{r_{ij}^3} - 3 \frac{(\mathbf{s}_i \cdot \mathbf{r}_{ij})(\mathbf{r}_{ij} \cdot \mathbf{s}_j)}{r_{ij}^5} - \frac{8\pi}{3} (\mathbf{s}_i \cdot \mathbf{s}_j) \delta(\mathbf{r}_{ij}) \right) \\
 \mathbf{H}_{ee}^{\text{OO}} &= -\frac{1}{4m^2 c^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j \neq i}^{N_{\text{elec}}} \left(\frac{\boldsymbol{\pi}_i \cdot \boldsymbol{\pi}_j}{r_{ij}} + \frac{(\boldsymbol{\pi}_i \cdot \mathbf{r}_{ij})(\mathbf{r}_{ij} \cdot \boldsymbol{\pi}_j)}{r_{ij}^3} \right) \\
 \mathbf{H}_{ee}^{\text{Darwin}} &= -\frac{\pi}{2m^2 c^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{j \neq i}^{N_{\text{elec}}} \delta(\mathbf{r}_{ij})
 \end{aligned} \tag{8.36}$$

The sums run over all values of i and j , excluding the $i = j$ term, and there is consequently a factor of $1/2$ included to avoid overcounting. $\mathbf{H}_{ee}^{\text{SO}}$ is a *spin–orbit* operator, describing the interaction of the electron spin with the magnetic field generated by its own movement, as given by the angular momentum operator $\mathbf{r}_{ij} \times \mathbf{p}_i$. $\mathbf{H}_{ee}^{\text{SOO}}$ is a *spin–other-orbit* operator, describing the interaction of an electron spin with the magnetic field generated by the movement of the other electrons, as given by the angular momentum operator $\mathbf{r}_{ij} \times \mathbf{p}_j$. $\mathbf{H}_{ee}^{\text{SS}}$ and $\mathbf{H}_{ee}^{\text{OO}}$ are *spin–spin* and *orbit–orbit* terms, accounting for additional magnetic interactions, where the orbit–orbit term comes from the Breit correction to \mathbf{V}_{ee} (eq. (8.34)). The (two-electron) Darwin interaction $\mathbf{H}_{ee}^{\text{Darwin}}$ contains a δ function, which arise from the divergence of the field ($\nabla \cdot \mathbf{F}$) from the (electron–electron) potential energy operator, i.e. $\nabla \cdot (\nabla(1/\mathbf{r})) = -4\pi\delta(\mathbf{r})$. The spin–spin interaction $\mathbf{H}_{ee}^{\text{SS}}$ also has a δ function, which comes from taking the curl of the vector potential associated with the magnetic dipole corresponding to the electron spin. A mathematical reformulation leads to a term involving the divergence of the \mathbf{r}/r^3

operator, giving $\nabla \cdot (\mathbf{r}/r^3) = (4\pi/3)\delta(\mathbf{r})$. Such terms are often called contact interactions, since they depend on the two particles being at the same position ($\mathbf{r} = 0$). In the spin–spin case, it is normally called the *Fermi Contact* (FC) term.

Operators involving one nucleus and one electron:

$$\begin{aligned} \mathbf{H}_{\text{ne}}^{\text{SO}} &= \frac{g_e \mu_B}{2mc^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{A=1}^{N_{\text{nuclei}}} Z_A \frac{\mathbf{s}_i \cdot (\mathbf{r}_{iA} \times \boldsymbol{\pi}_i)}{r_{iA}^3} \\ \mathbf{H}_{\text{ne}}^{\text{PSO}} &= \frac{\mu_N}{mc^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{A=1}^{N_{\text{nuclei}}} g_A \frac{\mathbf{I}_A \cdot (\mathbf{r}_{iA} \times \mathbf{p}_i)}{r_{iA}^3} \\ \mathbf{H}_{\text{ne}}^{\text{SS}} &= -\frac{g_e \mu_B \mu_N}{c^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{A=1}^{N_{\text{nuclei}}} g_A \left(\frac{\mathbf{s}_i \cdot \mathbf{I}_A}{r_{iA}^3} - 3 \frac{(\mathbf{s}_i \cdot \mathbf{r}_{iA})(\mathbf{r}_{iA} \cdot \mathbf{I}_A)}{r_{iA}^5} - \frac{8\pi}{3} (\mathbf{s}_i \cdot \mathbf{I}_A) \delta(\mathbf{r}_{iA}) \right) \\ \mathbf{H}_{\text{ne}}^{\text{Darwin}} &= \frac{\pi}{2m^2 c^2} \sum_{i=1}^{N_{\text{elec}}} \sum_{A=1}^{N_{\text{nuclei}}} Z_A \delta(\mathbf{r}_{iA}) \end{aligned} \quad (8.37)$$

The $\mathbf{H}_{\text{ne}}^{\text{SO}}$ operator is the one-electron part of the spin–orbit interaction, while the $\mathbf{H}_{\text{ee}}^{\text{SO}}$ and $\mathbf{H}_{\text{ee}}^{\text{SOO}}$ operators in eq. (8.36) define the two-electron part. The one-electron term dominates and the two-electron contribution is often neglected or accounted for approximately by introducing an effective nuclear charge in $\mathbf{H}_{\text{ne}}^{\text{SO}}$ (corresponding to a screening of the nucleus by the electrons). The effect of the spin–orbit operators is to mix states having different total spin, as for example singlet and triplet states.

The equivalent of the spin–other-orbit operator in eq. (8.36) splits into two contributions, one involving the interaction of the electron spin with the magnetic field generated by the movement of the nuclei, and one describing the interaction of the nuclear spin with the magnetic field generated by the movement of the electrons. Only the latter survives within the Born–Oppenheimer approximation, and it is normally denoted the *Paramagnetic Spin–Orbit* (PSO) operator. The spin–spin term is analogous to that in eq. (8.36), while the term describing the orbit–orbit interaction disappears owing to the Born–Oppenheimer approximation. The spin–orbit and (one-electron) Darwin terms are the same as given in eq. (8.18), except for the quantum field correction factor of $g_e \mu_B$.

All of the terms in eqs (8.35)–(8.37) may be used as perturbation operators in connection with non-relativistic theory,⁵ as discussed in more detail in Chapter 10. It should be noted, however, that some of the operators are inherently divergent and should not be used beyond a first-order perturbation correction.

8.4 Four-Component Calculations

Although relativistic effects can be included by perturbative operators describing corrections to the non-relativistic wave function, this rapidly becomes cumbersome if higher order corrections are required, and it is then perhaps more satisfying to include relativistic effects by solving the Dirac equation directly. The simplest approximative wave function is a single determinant constructed from four-component one-electron functions, called *spinors*, having large and small components multiplied with the two spin functions. The spinors are the relativistic equivalents of the spin-orbitals in non-relativistic theory. With such a wave function, the relativistic equation

corresponding to the Hartree–Fock equation is the *Dirac–Fock* equation, which in its time-independent form (setting $\boldsymbol{\pi} = \mathbf{p}$ and $m = 1$ in eq. (8.8)) can be written as in eq. (8.38).

$$[c\boldsymbol{\alpha} \cdot \mathbf{p} + \beta'c^2 + \mathbf{V}]\Psi = E\Psi \quad (8.38)$$

The requirement that the wave function should be stationary with respect to a variation in the orbitals, results in an equation that is formally the same as in non-relativistic theory, $\mathbf{FC} = \mathbf{SC}\epsilon$ (eq. (3.51)). However, the presence of solutions for the positronic states means that the desired solution is no longer the global minimum (Figure 8.1), and care must be taken that the procedure does not lead to variational collapse. The choice of basis set is an essential component in preventing this. Since practical calculations necessarily use basis sets that are far from complete, the large and small component basis sets must be properly balanced. The large component corresponds to the normal non-relativistic wave function, and has similar basis set requirements. The small component basis set is chosen to obey the *kinetic balance* condition, which follows from (8.15).

$$\chi^{\text{small}} = \frac{\boldsymbol{\sigma} \cdot \mathbf{p}}{2c} \chi^{\text{large}} \quad (8.39)$$

The use of kinetic balance ensures that the relativistic solution smoothly reduces to the non-relativistic wave function as c is increased. The presence of the momentum operator in eq. (8.39) means that the small component basis set must contain functions that are derivatives of the large component basis set, making the former roughly twice the size of the latter. This means that there are ~ 8 times as many *large–small* two-electron integrals and ~ 16 times as many *small–small* integrals, than there are *large–large* type integrals. A relativistic calculation thus requires roughly 25 times as many two-electron integrals compared with a non-relativistic calculation.

When the Dirac operator is invoked, the point charge model of the nucleus also becomes problematic. For a non-relativistic hydrogen atom, the orbitals have a *cusp* (discontinuous derivative) at the nucleus. However, the relativistic solutions have a *singularity*. A singularity is much harder to represent in an approximate treatment (such as an expansion in a Gaussian basis) than a cusp. Consequently, a (more realistic) finite-size nucleus is often used in relativistic methods. A finite nucleus model removes the singularity of the orbitals, which now assume a Gaussian type behaviour within the nucleus. Neither experiments nor theory, however, provide a good model for how the positive charge is distributed *within* the nucleus. The wave function and energy will of course depend on the exact form used for describing the nuclear charge distribution. A popular choice is either a uniformly charged sphere, where the radius is proportional to the nuclear mass to the $1/3$ power, or a Gaussian charge distribution (which facilitates the calculation of the additional integrals) with the exponent depending on the nuclear mass. Note that this implies that the energy (and derived properties) depends on the specific isotope, not just the atomic charge, i.e. the results for say ^{37}Cl will be (slightly) different from ^{35}Cl . The difference between a finite and a point charge nuclear model is large in terms of total energy (~ 1 au), however, the exact shape for the finite nucleus is not important. For valence properties, any “reasonable” model gives essentially the same results.

8.5 Relativistic Effects

The differences due to relativity can be described as:

- (1) Differences in the dynamics due to the velocity-dependent mass of the electron. This alters the size of the orbitals: s- and p-orbitals contract while d- and f-orbitals expand.
- (2) New (magnetic) interactions in the Hamiltonian operator due to electron spin. The spin-orbit coupling, for example, destroys the picture of an orbital having a definite spin.
- (3) Introduction of “negative” energy (positron) states. The coupling between the electronic and positronic states introduces a “small” component in the electronic wave function. This leads to a change in the shape of the orbitals: relativistic orbitals, for example, do not have nodes.
- (4) Modification of the potential operator due to the finite speed of light. In the lowest order approximation, this corresponds to addition of the Breit operator to the Coulomb interaction.

Results from fully relativistic calculations are scarce, and there is no clear consensus on which effects are the most important. The Breit (Gaunt) term is believed to be small and many relativistic calculations neglect this term, or include it as a perturbational term evaluated from the converged wave function. For geometries, the relativistic contraction of the s-orbitals normally means that bond lengths become shorter.

Working with a full four-component wave function and the Dirac–Fock operator is significantly more complicated than solving the Roothaan–Hall equations. The spin dependence can no longer be separated out, and the basis set for the small component of the wave function must contain derivatives of the corresponding large component basis. This means that the basis set becomes three to four times as large as in the non-relativistic case for a comparable accuracy. Furthermore, the presence of magnetic terms (spin) in the Hamiltonian operator means that the wave function contains both real and imaginary parts, yielding a factor of two in complexity. In practice, a (single-determinant) Dirac–Fock–Coulomb calculation is about two orders of magnitude more expensive than the corresponding non-relativistic Hartree–Fock case, although implementation of integral screening techniques is likely to reduce this factor.⁶ Since heavy atom systems by definition contain many electrons, even small systems (in terms of the number of atoms) are demanding. A relativistic calculation for a single radon atom with a DZP quality basis, for example, is computationally equivalent to a non-relativistic calculation of a $\text{C}_{13}\text{H}_{28}$ alkane, for a comparable quality in term of basis set limitations. To further complicate matters, there are many more systems that cannot be adequately described by a single-determinant wave function in a relativistic treatment owing to the spin–orbit coupling, and therefore require MCSCF type wave functions.

Since working with the full four-component wave function is so demanding, various approximative methods have been developed where the small component of the wave function is “eliminated” to a certain order in $1/c$ or approximated (such as the *Foldy–Wouthuysen*⁷ or *Douglas–Kroll* transformations,⁸ thereby reducing the four-component wave function to only two components. A description of such methods is beyond the scope of this book.

Table 8.1 Properties of the sixth group dihydrides

System	Non-relativistic			Relativistic correction				
	Total energy (au)	R_{eq} (Å)	θ_{eq} (°)	ΔE_{atom} (kJ/mol)	Total energy (au)	R_{eq} (Å)	θ_{eq} (°)	ΔE_{atom} (kJ/mol)
H ₂ O	-76.054	0.9391	107.75	643.8	-0.055	-0.00003	-0.07	-1.6
H ₂ S	-398.641	1.3429	94.23	514.1	-1.107	-0.00015	-0.09	-4.5
H ₂ Se	-2400.977	1.4530	93.14	459.4	-28.628	-0.00260	-0.27	-13.3
H ₂ Te	-6612.797	1.6557	92.57	392.5	-182.072	-0.00720	-0.58	-37.7
H ₂ Po	-20676.709	1.7539	92.21	350.2	-1555.822	-0.01060	-1.62	-126.8




Figure 8.2 Converging the computational results by increasing the basis set, the amount of electron correlation and description of the relativistic effects

Table 8.1 illustrates the magnitude of relativistic effects for dihydrides of the sixth main group in the periodic table, where the relativistic calculations are of the Dirac–Fock–Coulomb type (i.e. a single-determinant wave function and neglecting the Breit interaction).⁹ The relativistic correction to the total energy is significant: even for a first row species such as H_2O is the difference 0.055 au (145 kJ/mol). It increases rapidly down the periodic table, and reaches $\sim 7\%$ of the total energy for H_2Po , but the equilibrium distances and angles change only marginally. Similarly, the atomization energy (for breaking both X–H bonds completely) is remarkably insensitive to the large changes in the total energies. This is of course due to a high degree of cancellation of errors, the major relativistic correction is associated with the inner-shell electrons of the heavy atom, with the correction being almost constant for the atom and the molecule. For the lighter elements the effect on the atomization energies is almost solely due to the spin–orbit interaction in the triplet X atom (e.g. $\text{H}_2\text{O} \rightarrow {}^3\text{O} + 2{}^2\text{H}$) which is not present in the singlet H_2X molecule.

Similar results have been obtained for the fourth group tetrahydrides, CH_4 , SiH_4 , SnH_4 , GeH_4 and PbH_4 , where the Gaunt term has been shown to give corrections typically an order of magnitude less than the other relativistic changes.¹⁰ The general conclusion is that relativistic effects for geometries and energetics can normally be neglected for molecules containing only first and second row elements. This is also true for third row elements, unless a high accuracy is required. Although the geometry and atomization energy changes for H_2S and H_2Se in Table 8.1 may be considered significant, it should be noted that the errors due to incomplete basis sets and neglect of electron correlation are much larger than the relativistic corrections. The experimental geometries for H_2S and H_2Se , for example, are 1.3356 Å and 92.12°, and 1.4600 Å and 90.57°, respectively. While the relativistic contraction of the H–Se bond is 0.0026 Å, the basis set and electron correlation error is 0.0070 Å. Relativistic effects typically

become comparable to those from electron correlation at atomic numbers ~40–50. For molecules involving atoms beyond the fourth row in the periodic table, however, relativistic effects cannot be neglected for quantitative work. It should be noted that an approximate inclusion of the scalar relativistic effects, most notably the change in orbital size, can be modelled by replacing the inner electrons with a relativistic pseudopotential, as discussed in Section 5.9.

Relativistic methods can be extended to include electron correlation by methods analogous to the non-relativistic cases, e.g. CI, MCSCF, MP and CC. Such methods are currently at the development stage.¹¹ Once relativistic effects are considered, one may thus expand the two-dimensional Figure 4.2 with a third axis describing how accurate the relativistic effects are treated, for example measured in terms of one-, two- or four-component wave functions.

References

1. R. E. Moss, *Advanced Molecular Quantum Mechanics*, Chapman and Hall, 1973; P. Pyykko, *Chem. Rev.*, **88** (1988), 563; J. Almlöf, O. Gropen, *Rev. Comp. Chem.*, **8** (1996), 203; K. Balasubramanian, *Relativistic Effects in Chemistry*, Wiley, 1997.
2. E. van Lenthe, E. J. Baerends, J. G. Snijder, *J. Chem. Phys.*, **99** (1993), 4597; J. G. Snijder, A. J. Sadlej, *Chem. Phys. Lett.*, **252** (1996), 51.
3. R. van Leeuwen, E. van Lenthe, E. J. Baerends, J. G. Snijder, *J. Chem. Phys.*, **101** (1994), 1272.
4. R. McWeeny, *Methods of Molecular Quantum Mechanics*, Academic Press, 1992; S. A. Perera, R. J. Bartlett, *Adv. Quant. Chem.*, **48** (2005), 435.
5. S. Coriani, T. Helgaker, P. Jørgensen, W. Klopper, *J. Chem. Phys.*, **121** (2004), 6591.
6. T. Saue, K. Faegri, T. Helgaker, O. Gropen, *Mol. Phys.*, **91** (1997), 937.
7. L. L. Foldy, S. A. Wouthuysen, *Phys. Rev.*, **78** (1950), 29.
8. M. Douglas, N. M. Kroll, *Ann. Phys. NY*, **82** (1974), 89.
9. L. Pisani, E. Clementi, *J. Chem. Phys.*, **101** (1994), 3079.
10. O. Visser, L. Visscher, P. J. C. Aerts, W. C. Nieuwpoort, *Theor. Chem. Acta*, **81** (1992), 405.
11. L. Visscher, T. J. Lee, K. G. Dyall, *J. Chem. Phys.*, **105** (1996), 8769; L. Visscher, *J. Comp. Chem.*, **23** (2002), 759.

9 Wave Function Analysis

The previous chapters have focused on various methods for obtaining more or less accurate solutions to the Schrödinger equation. The natural “byproduct” of determining the electronic wave function is the energy. However, there are many other properties that may be derived. Although the quantum mechanical description of a molecule is in terms of positive nuclei surrounded by a cloud of negative electrons, chemistry is still formulated as “atoms” held together by “bonds”. This raises questions such as: given a wave function, how can we define an atom and its associated electron population, or how do we determine whether two atoms are bonded?

Atomic charge is an example of a property often used for discussing/rationalizing structural and reactivity differences.¹ There are three commonly used methods for assigning a charge to a given atom:

- (1) Partitioning the wave function in terms of the basis functions.
- (2) Fitting schemes.
- (3) Partitioning the electron density into atomic domains.

9.1 Population Analysis Based on Basis Functions

The electron density ρ (probability of finding an electron) at a certain position \mathbf{r} from a single molecular orbital containing one electron is given as the square of the MO ϕ .

$$\rho_i(\mathbf{r}) = \phi_i^2(\mathbf{r}) \quad (9.1)$$

Assuming that the MO is expanded in a set of normalized, but non-orthogonal, basis functions χ , this can be written as in eq. (9.2) (see also eq. (3.49)).

$$\begin{aligned} \phi_i &= \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} \\ \phi_i^2 &= \sum_{\alpha\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta i} \chi_{\alpha} \chi_{\beta} \end{aligned} \quad (9.2)$$

Integrating and summing over all occupied MOs gives the total number of electrons, N_{elec} .

$$\sum_i^{N_{\text{occ}}} \int \phi_i^2 d\mathbf{r} = \sum_i^{N_{\text{occ}}} \sum_{\alpha\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta i} \int \chi_{\alpha} \chi_{\beta} d\mathbf{r} = \sum_i^{N_{\text{occ}}} \sum_{\alpha\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta i} S_{\alpha\beta} = N_{\text{elec}} \quad (9.3)$$

We may generalize this by introducing an occupation number (number of electrons), n , for each MO. For a single-determinant wave function, this will be either 0, 1 or 2, while it may be a fractional number for a correlated wave function (Section 9.5).

$$\sum_i^{N_{\text{orb}}} n_i \int \phi_i^2 d\mathbf{r} = \sum_{\alpha\beta}^{M_{\text{basis}}} \left(\sum_i^{N_{\text{orb}}} n_i c_{\alpha i} c_{\beta i} \right) S_{\alpha\beta} = \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} S_{\alpha\beta} = N_{\text{elec}} \quad (9.4)$$

The sum of the product of MO coefficients and the occupation numbers is the density matrix defined in eq. (3.52), and the sum over the product of the density and overlap matrices elements is the number of electrons.

The *Mulliken Population Analysis* uses the $\mathbf{D} \cdot \mathbf{S}$ matrix for distributing the electrons into atomic contributions² ($\mathbf{D} \cdot \mathbf{S}$ is the entrywise product matrix, Section 16.1, i.e. the products of elements, not elements of the product matrix). A diagonal element $D_{\alpha\alpha} S_{\alpha\alpha}$ is the number of electrons in the α AO, and an off-diagonal element $D_{\alpha\beta} S_{\alpha\beta}$ is (half) the number of electrons shared by AOs α and β (there is an equivalent $D_{\beta\alpha} S_{\beta\alpha}$ element). The contributions from all AOs located on a given atom A may be summed up to give the number of electrons associated with atom A. This requires a decision on how a contribution involving basis functions on different atoms should be divided. The simplest, and the one used in the Mulliken scheme, is to partition the contribution equally between the two atoms. The Mulliken electron population is thereby defined as in eq. (9.5).

$$\rho_A = \sum_{\alpha \in A}^{M_{\text{basis}}} \sum_{\beta}^{M_{\text{basis}}} D_{\alpha\beta} S_{\alpha\beta} \quad (9.5)$$

The *gross charge* on atom A is the sum of the nuclear and electronic contributions.

$$Q_A = Z_A - \rho_A \quad (9.6)$$

The Mulliken method corresponds to a partitioning of the $\mathbf{D} \cdot \mathbf{S}$ matrix product, another commonly used method is the *Löwdin* partitioning, which uses the $\mathbf{S}^{1/2} \cdot \mathbf{D} \cdot \mathbf{S}^{1/2}$ matrix for analysis.³ These are mathematically related as shown in eq. (9.7).

$$\begin{aligned} \sum \mathbf{D} \cdot \mathbf{S} &= N_{\text{elec}} \\ \sum \mathbf{S}^{1/2} \cdot (\mathbf{D} \cdot \mathbf{S}) \cdot \mathbf{S}^{-1/2} &= \mathbf{S}^{1/2} N_{\text{elec}} \mathbf{S}^{-1/2} \\ \sum \mathbf{S}^{1/2} \cdot \mathbf{D} \cdot \mathbf{S}^{-1/2} &= N_{\text{elec}} \end{aligned} \quad (9.7)$$

The Löwdin method is equivalent to a population analysis of the density matrix in the orthogonalized basis set (Section 13.2) formed by transforming the original set of functions by $\mathbf{S}^{-1/2}$.

$$\chi' = \mathbf{S}^{-1/2} \chi \quad (9.8)$$

The Mulliken and Löwdin methods are just particular examples of a whole family of population analysis using $\mathbf{S}^n \cdot \mathbf{D} \cdot \mathbf{S}^{1-n}$ matrices.⁴ The Mulliken and Löwdin methods

give different atomic charges but mathematically there is nothing to indicate which of these partitionings gives the “best” result. It should be noted, however, that the Löwdin method is not rotationally invariant if the basis set contains Cartesian polarizations functions rather than spherical functions.⁵ The lack of rotational invariance means that symmetry equivalent atoms may end up having different charges.

There are some common problems with all population analyses based on partitioning the wave function in terms of basis functions:

- (1) The diagonal elements may be larger than two. This implies more than two electrons in an orbital, violating the Pauli principle.
- (2) The off-diagonal elements may become negative. This implies a negative number of electrons between two basis functions, which clearly is physically impossible.
- (3) There is no objective reason for dividing the off-diagonal contributions equally between the two orbitals. It may be argued that the most “electronegative” (which then needs to be defined) atom (orbital) should receive most of the shared electrons.
- (4) Basis functions centred on atom A may have a small exponent, such that they effectively describe the wave function far from atom A. Nevertheless, the electron density is counted as only belonging to A.
- (5) The dipole, quadrupole, etc., moments are in general not conserved, i.e. a set of population atomic charges does not reproduce the original multipole moments.

The Mulliken scheme suffers from all of the above, while the Löwdin method solves problems (1), (2) and (3). In the orthogonalized basis, all off-diagonal elements are 0, and the diagonal elements are restricted to values between 0 and 2.

Problem (4) is especially troublesome, as a few examples for the water molecule will demonstrate. A reasonable description of the wave function can be obtained by an HF single determinant with a DZP basis set. An equally good wave function (in terms of energy) may be constructed by having a very large number of basis functions centred on oxygen, and none on the hydrogens. The latter will, according to the above population analysis, have a +1 charge on hydrogen, and a –2 charge on oxygen. Worse, another equally good wave function may be constructed by having a large number of basis functions only on the hydrogens. This will give charges of –4 for each of the hydrogens and +8 for the oxygen. Or the basis functions can be taken to be non-nuclear-centred, in which case the electrons are not associated with any nuclei at all, i.e. atomic charges of +1 and +8!

The fundamental problem is that basis functions often describe electron density near nucleus other than the one they are centred on. An *s*-type Gaussian function on oxygen with an exponent of 0.15, for example, has a maximum in the radial distribution ($r^2\phi^2$) that peaks at 0.97 Å, i.e. at the distance where the hydrogen nuclei are located. Atomic charges calculated from a Mulliken or Löwdin analysis will therefore not converge to a constant value as the size of the basis set is increased. Enlarging the basis set involves the addition of more and more diffuse basis functions, often leading to unpredictable changes in the atomic charges. This is a case where a “better” theoretical procedure is actually counterproductive. Basis function derived population analyses are therefore most useful for comparing trends in electron distributions, when small- or medium-sized basis sets (which only contain relatively tight functions) are used.

The density matrix can also be used for generating information about bond strengths. A quantitative measure is given by the *Bond Order* (BO). It was originally defined from bond distances as shown in eq. (9.9).⁶

$$\text{BO} = e^{-(r-r_0)/a} \quad (9.9)$$

If the bond orders for ethane, ethylene and acetylene are defined to be 1, 2 and 3, respectively, the a constant is found to have a value of approximately 0.3 Å. For bond orders less than 1 (i.e. breaking and forming single bonds), it appears that a value of 0.6 Å is a more appropriate proportionality constant. A “Mulliken” style measure of the bond strength between atoms A and B can be defined from the density matrix as eq. (9.10) (note that this involves elements of the product of the **D** and **S** matrices).⁷

$$\text{BO}_{AB} = \sum_{\alpha \in A}^{M_{\text{basis}}} \sum_{\beta \in B}^{M_{\text{basis}}} (\mathbf{DS})_{\alpha\beta} (\mathbf{DS})_{\beta\alpha} \quad (9.10)$$

This will again be basis set dependent, but not nearly as sensitive as atomic charges. The concept can be generalized to higher order quantities, i.e. three-, four-, five-, etc., centre bond indices, which are derived from products of **DS** elements.⁸

Population analysis with semi-empirical methods requires a special comment. These methods normally employ the ZDO approximation, i.e. the overlap **S** is a unit matrix. The population analysis can therefore be performed directly on the density matrix. In some cases, however, a Mulliken population analysis is performed with **D**·**S**, which requires an explicit calculation of the **S** matrix.

9.2 Population Analysis Based on the Electrostatic Potential

One area where the concept of atomic charges is deeply rooted is in force field methods (Chapter 2). A significant part of the non-bonded interaction between polar molecules is described in terms of electrostatic interactions between fragments having an internal asymmetry in the electron distribution. The fundamental interaction is between the *ElectroStatic Potential* (ESP), also called the *Molecular Electrostatic Potential* (MEP), generated by one molecule (or fraction thereof) and the charged particles of another. The ESP at position **r** is given as a sum of contributions from the nuclei and the electronic wave function.

$$\phi_{\text{ESP}}(\mathbf{r}) = \sum_A^{\text{nuclei}} \frac{Z_A}{|\mathbf{r} - \mathbf{R}_A|} - \int \frac{|\Psi(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}' \quad (9.11)$$

The first part of the potential is trivially calculated from the nuclear charges and positions, but the electronic contribution requires knowledge of the wave function. The latter is not available in force field methods, and the simplest way of modelling the electrostatic potential is to assign partial charges to each atom (Section 2.2.6). Atomic charges may be treated as regular force field parameters, and assigned values based on fitting to experimental data, such as dipole, quadrupole, octopole, etc., moments, but there are rarely enough data to allow a unique assignment.

A common way of deriving partial atomic charges in force fields is to choose a set of parameters that in a least squares sense generates the best fit to the actual electrostatic potential as calculated from an electronic wave function.⁹ The electrostatic

potential stretches far beyond the molecular dimension (the Coulomb interaction falls off as R^{-1} (charge) or as R^{-3} (dipole)), but the most important region is just beyond the van der Waals distance. The potential is sampled by placing a suitable grid of points around each nucleus with distances from just outside the van der Waals radius to about twice that distance, with a typical sampling having a few hundred points for each atom. The atomic charges are determined as those parameters that reproduce the electrostatic potential as closely as possible at these points, subject to the constraint that the sum is equal to the total molecular charge. In some cases, the atomic charges may also be constrained to reproduce for example the dipole moment. Additional constraints such as forcing the total charge of a subgroup (such as a methyl group or an amino acid) to be zero are also often employed as this improves the parameter transferability and computational issues related to calculating the electrostatic energy. The various schemes for deriving atomic charges differ in the number and location of points used in the fitting, and whether additional constraints beyond preservation of charge are added, and may produce slightly different results. In many cases, the fitted set of charges is uniformly increased by 10–20% to model the fact that polarization in condensed phases will increase the effective dipole moment relative to the isolated molecule case (Section 2.2.7), or the charges are derived by fitting to an ESP which naturally overestimates the charge polarization (for example HF/6-31G(d), Section 4.3).

The electrostatic potential depends directly on the wave function and therefore converges as the size of the basis set and amount of electron correlation is increased. Since the potential depends directly on the electron density ($\rho = |\Psi|^2$), it is in general found to be fairly insensitive to the level of sophistication, i.e. an HF calculation with a DZP type basis set already gives quite good results. One might thus anticipate that atomic charges based on fitting to the electrostatic potential would lead to well-defined values. This, however, is not the case. Besides the already mentioned dependence on the sampling points, another problem is that a straightforward fitting tends to give conformationally dependent charges.¹⁰ The three hydrogens in a freely rotating methyl group, for example, may end up having significantly different charges, or two conformations may give two widely different sets of fitted parameters. This is a problem in connection with force field methods that rely on the fundamental assumption that parameters are transferable between similar fragments, and consequently atoms that are easily interchanged (e.g. by bond rotation) should have identical parameters. Conformationally dependent charges can be modelled in force field methods by fluctuating charge or polarization models (Section 2.2.7), but this leads to significantly more complicated force fields, and consequently loss of computational efficiency.

One way of eliminating the problem with conformationally dependent charges is to add additional constraints, for example forcing the three hydrogens in a methyl group to have identical charges¹¹ or averaging over different conformations.¹² A more fundamental problem is that the fitting procedure becomes statistically underdetermined for large systems, although the severity of this depends on how the fitting is done.¹³ The difference between the true electrostatic potential and that generated by a set of atomic charges on say 80% of the atoms is not significantly reduced by having fitting parameters on all atoms. The electrostatic potential experienced outside the molecule is mainly determined by the atoms near the surface, and consequently the charges on atoms buried within a molecule cannot be assigned with any great confidence. Even

for a medium-sized molecule, it may only be statistically valid to assign charges to perhaps half the nuclei. Having a full set of atomic charges thus forms a redundant set: many different sets of charges may be chosen, all of which are capable of reproducing the true electrostatic potential to almost the same accuracy. Although a very large number of sampling points (several thousand) may be chosen to be fitted by relatively few (perhaps 20–30) parameters, the fact that the sampling points are highly correlated makes the problem underdetermined. In practical applications, additional penalty terms are therefore often added, to ensure that only those atoms that contribute significantly in a statistical sense acquire non-zero charges.¹⁴

Another problem with atomic charges determined by fitting is related to the absolute accuracy. Although inclusion of charges on all atoms does not significantly improve the results over that determined from a reduced set of parameters, the *absolute* deviation between the true and fitted electrostatic potentials can be quite large. Interaction energies as calculated by an atom-centred charge model in a force field may be off by several kJ/mol per atom in certain regions of space just outside the molecular surface, an error of one or two orders of magnitude larger than the van der Waals interaction. In order to improve the description of the electrostatic interaction, additional non-nuclear-centred charges may be added,¹⁵ or dipole, quadrupole, etc., moments may be added at nuclear or bond positions.¹⁶ These descriptions are essentially equivalent since a dipole may be generated as two oppositely charged monopoles, a quadrupole as four monopoles, etc.



Figure 9.1 Generation of dipole and quadrupole moments by charges

The *Distributed Multipole Analysis* (DMA) developed by A. Stone uses the fact that the electrostatic potential arising from the charge overlap between two basis functions can be written in terms of a multipole expansion around a point between the two nuclei.¹⁷ These moments can be calculated *directly* from the density matrix and the basis functions, and are not a result of a fitting procedure. The multipole expansion is furthermore finite, the highest non-vanishing term is given as the sum of the angular momenta for the two basis functions, e.g. the product of two *p*-functions gives at most rise to a quadrupole moment. For Gaussian orbitals the expansion point is given in eq. (9.12), where \mathbf{R}_A and \mathbf{R}_B are the positions of the two nuclei, and α and β are the exponents of the basis functions (this follows since the product of two Gaussians is a single Gaussian located between the two original, eq. (3.60)).

$$\mathbf{R}_C = \frac{\alpha \mathbf{R}_A + \beta \mathbf{R}_B}{\alpha + \beta} \quad (9.12)$$

If such distributed multipoles are assigned for each pair of basis functions, the electrostatic potential as seen from outside the charge distribution is reproduced exactly. This, however, would mean that $\sim M_{\text{basis}}^2$ different sites are required. In practice, only the nuclei and possibly the midpoints of all bonds are selected as multipole points, and

all the pair expansion points are moved to the nearest multipole point. By moving the origin, the termination after a finite number of terms is destroyed, and an infinite sum over all moments must be used for an exact representation. Since most of the pair expansion points are rather close to either a nucleus or the centre of a bond, the higher order moments are usually quite small. Furthermore, since the majority of the electron density can be represented with just s - and p -functions for elements belonging to the first or second row of the periodic table, it follows that a representation in terms of charges, dipoles and quadrupoles located on all nuclei and bond centres gives a quite accurate representation of the electrostatic potential. If only nuclear-centred multipoles are used, an expansion up to quadrupoles will typically generate an electrostatic potential of the same quality as a model based on fitted atomic charges. A disadvantage of the DMA approach is that the calculated multipole moments are quite sensitive to the employed basis set, in analogy with other analysis based directly on the basis functions used for representing the wave function. Alternatively, multipoles may be fitted to either the electrostatic potential¹⁸ or the DMA multipoles.¹⁹ Such fitted multipole methods typically reduce the required moments by one or two, i.e. fitted charges and dipoles can reproduce DMA results including up to quadrupoles or octopoles.

9.3 Population Analysis Based on the Electron Density

The examples in Section 9.1 illustrate that it would be desirable to base a population analysis on properties of the wave function or electron density itself, and not on the basis set chosen for representing the wave function. The electron density is the square of the wave function integrated over $N_{\text{elec}} - 1$ coordinates (it does not matter which coordinates since the electrons are indistinguishable).

$$\rho(\mathbf{r}_1) = \int |\Psi(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3 \dots \mathbf{r}_{N_{\text{elec}}})|^2 d\mathbf{r}_2 d\mathbf{r}_3 \dots d\mathbf{r}_{N_{\text{elec}}} \quad (9.13)$$

Some textbooks state that it is impossible to define a unique atomic charge since there is no quantum mechanical operator associated with charge. This is not true: the electronic charge operator is simply the negative of the number operator (the charge from an electron is -1). The problem is in the definition of an “atom” within a molecule. If the total molecular volume could somehow be divided into subsections, each belonging to one specific nucleus, then the electron density could be integrated to give the number of electrons present in each of these atomic basins Ω , and the (net) atomic charge Q is then obtained by adding the nuclear charge Z .

$$Q_A = Z_A - \int_{\Omega_A} \rho(\mathbf{r}) d\mathbf{r} \quad (9.14)$$

The division into atomic basins requires a *choice* to be made as to whether a certain point in space belongs to one nucleus or another, and several different schemes have been proposed.

9.3.1 Atoms In Molecules

Perhaps the most rigorous way of dividing a molecular volume into atomic subspaces is the *Atoms In Molecules* (AIM) method of R. Bader.²⁰ The electron density is a

function of three spatial coordinates, and it may be analyzed in terms of its topology (maxima, minima and saddle points). In the large majority of cases it is found that the only maxima in the electron density occur at the nuclei (or very close to them), which is reasonable since they are the only sources of positive charge. The nuclei thus act as *attractors* of the electron density. At each point in space the gradient of the electron density points in the direction of the strongest (local) attractor. This forms a rigorous way of dividing the physical space into atomic subspaces: starting from a given point in space a series of infinitesimal steps may be taken in the gradient direction until an attractor is encountered. The collection of all such points forms the atomic basin associated with the attractor (nucleus). If the negative of the electron density is considered, the attractors are local minima, and a basin is then defined as points which end up at the local minimum by a steepest descent minimization (Section 12.2.1). In the other direction (away from other nuclei) the gradient goes asymptotically to zero, and the atomic basin stretches into infinity in this direction. The border between two three-dimensional atomic basins is a two-dimensional surface, as illustrated in Figure 9.2.




Figure 9.2 Dividing surface between two atomic basins
Reprinted with permission from The Americal Chemical Society.²¹

The carbon and hydrogen atomic basins in cyclopropane are shown in Figure 9.3.

Once the molecular volume has been divided up, the electron density may be integrated within each of the atomic basins to give atomic charges and dipole, quadrupole, etc., moments. As the dividing surface is rigorously defined in terms of the electron density, these quantities will converge to specific values as the quality of the wave function is increased. Furthermore, as only the electron density is involved, the results are fairly insensitive to the theoretical level used for generating the wave function. If the net charges are taken as nuclear centred (analogous to partial charges for force field methods), they do not reproduce the molecular dipole, quadrupole, etc., moments, nor do they yield a good representation of the molecular electrostatic potential, and they are therefore not suitable for transferring to a force field environment for modelling purposes. If, however, the dipole moments of the atomic basins are also considered, the total molecular dipole moment is reproduced, and similarly for higher order moments. The dipole moment of CO, for example, is close to zero (0.122 debye), despite calculated AIM charges of ± 1.1 . The large dipole moment generated by the charge




Figure 9.3 Hydrogen and carbon AIM basins for cyclopropane; dots indicate bond and ring critical points²²

transfer is almost exactly cancelled by compensating atomic dipoles. The AIM method is often criticized for generating too large atomic charges for polar bonds, but it should be recognized that this is largely due to the neglect of higher order moments.

A more fundamental problem in the AIM approach is the presence of non-nuclear attractors in certain metallic systems, such as lithium and sodium clusters.²³ While these are of interest by themselves, they spoil the picture of electrons associated with nuclei, forming atoms within molecules. It should be noted that non-nuclear attractors can also be found for other systems such as ethyne when a low level of theory is used for calculating the electron density.

For a point on a dividing surface between two atomic basins the gradient of the density must necessarily be tangential to the surface. Following the gradient path for such a point leads to a stationary point on the surface where the total derivative is zero, marked with a dot in Figure 9.2. The basin attractor is also a stationary point on the electron density surface. The second derivative of the electron density, the Hessian, is a function of the three (Cartesian) coordinates, i.e. it is a 3×3 matrix. At stationary points, it may be diagonalized and the number of negative eigenvalues determined. The basin attractor is an overall maximum, it has three negative eigenvalues. Other stationary points are usually found between nuclei that are “bonded”. Such points have a minimum in the electron density in the direction of the nuclei, and a maximum in the perpendicular directions, i.e. there is one positive and two negative eigenvalues in the Hessian. These are known as *bond critical points*. If the negative of the electron density is considered instead, the attractors are minima (all positive eigenvalues in the second derivative matrix) and the bond critical points are analogous to transition structures (one negative eigenvalue). Comparing with potential energy surfaces (Section 13.1), the (negative) electron density surface may be analyzed in terms of “reaction paths” connecting “transition structures” with minima. Such paths trace the maximum electron density connecting the two nuclei, and may be taken as the molecular “bond”. It should be noted that bond critical points are not necessarily located on a straight line connecting two nuclei: small strained rings such as cyclopropane, for example, have bond paths that are significantly curved, as illustrated in Figure 9.3. Indeed, the degree of bending tends to correlate with the strain energy.

The value of the electron density at the bond critical point correlates with the strength of the bond, the bond order. As mentioned above, there are certain systems such as metal clusters that have non-nuclear-centred attractors. The corresponding bond critical points have electron densities at least an order of magnitude smaller than “normal” single bonds, and the value of the density at the local maximum is only slightly larger than at the bond critical point. The non-nuclear-centred attractors are thus only weakly defined, and may be considered as a special kind of metal bonding, where a “sea” of electrons with weak local maxima surrounds the positive nuclei, which are strong local maxima. In certain cases, bond critical points may also be found between atoms that are not bonded, but experience a strong steric repulsion, corresponding to situations where two atoms are forced to be closer than the sum of their van der Waals radii. Such systems usually have values of the electron density at the bond critical point that are at least an order of magnitude smaller than ordinary “bonded” atoms.²⁴

There are two other types of critical points, having either one or no negative eigenvalues in the density Hessian. The former are usually found in the centre of a ring (illustrated in Figure 9.3 for cyclopropane), and are consequently denoted as a *ring critical point*. The latter are typically found at the centre of a cage (e.g. cubane), and are denoted as a *cage critical point*. They correspond to local minima in the electron density in two or three directions.

The second derivative of the electron density, the Laplacian $\nabla^2\rho$, provides information on where electron density is depleted or increased. At a bond critical point the sign of the Laplacian has been used for characterizing the nature of the bond, i.e. a negative value indicates a covalent bond, while a positive value indicates an ionic bond or a van der Waals interaction.

The division of the molecular volume into atomic basins follows from a deeper analysis based on the principle of stationary action. The shapes of the atomic basins, and the associated electron densities, in a given functional group are very similar in different molecules.²⁵ The *local* properties of the wave function are therefore transferable to a very good approximation, which rationalizes the basis for organic chemistry, i.e. functional groups react similarly in different molecules. It may be shown that any observable molecular property may be written as a sum of corresponding atomic contributions.

$$\langle \mathbf{A} \rangle = \sum_i^{\text{atomic basins}} \langle \mathbf{A} \rangle_i \quad (9.15)$$

The total energy, for example, may be written as a sum of atomic energies, and these atomic energies are again almost constant for the same structural units in different molecules. The atomic basins are probably the closest quantum mechanical analogy to the chemical concepts of atoms within a molecule. The good degree of transferability furthermore provides a rationale for defining atom types in force field methods.

9.3.2 Voronoi, Hirshfeld and Stewart atomic charges

The AIM approach partitions the physical space into atomic basins based on a topological analysis of the electron density itself, but several other methods have been proposed for dividing the space into atomic contributions.

Voronoi charges are based on dividing the physical space according to a distance criterion, i.e. a given point in space belongs to the nearest nucleus. This is reminiscent of the Mulliken equal partitioning, except that it is the *physical space* between two nuclei that is divided equally to each side, not the *Hilbert space* defined by the basis functions. The atomic basins are bounded by planes perpendicular to the interatomic bonds, and are called *Voronoi polyhedra* or *Voronoi cells*. Voronoi charges tend to be rather large. A modified approach where these dividing planes are moved away from the bond midpoint by a distance related to the relative atom sizes, defined by their van der Waals radii, has also been proposed, and this gives significantly smaller charges.²⁶

Hirshfeld (or *stockholder*) charges are based on using atomic densities for partitioning the molecular electron density.²⁷ The *promolecular* density is defined as the sum of atomic densities placed at the nuclear geometries in the molecule. The *actual* molecular electron density at each point in space is then partitioned by weighting factors according to the promolecular contributions.

$$\begin{aligned} \rho_{\text{promolecule}}(\mathbf{r}) &= \sum_A^{M_{\text{atoms}}} \rho_A^{\text{atomic density}}(\mathbf{r}) \\ w_A(\mathbf{r}) &= \frac{\rho_A^{\text{atomic density}}(\mathbf{r})}{\rho_{\text{promolecule}}(\mathbf{r})} \\ Q_A &= Z_A - \int w_A(\mathbf{r}) \rho_{\text{molecule}}(\mathbf{r}) d(\mathbf{r}) \end{aligned} \quad (9.16)$$

Hirshfeld charges may be considered as a soft-boundary version of the Voronoi charges. An ambiguity in the Hirshfeld method is the source of the atomic densities.

The normal approach is to use spherically averaged ground state densities for neutral atoms but, in some cases, other valence configurations may be considered.²⁸ Furthermore, the level of theory (method and basis set) for calculating the atomic density is a potential variable.

Stewart atoms are defined as the spherical densities centred at the nuclei that in a least squares sense fit the molecular density as well as possible, and the resulting densities can be integrated to yield atomic charges and higher order electric moments.²⁹ The Stewart atomic densities often have small *negative* contributions far from the nuclei, and the resulting charges are often large and counterintuitive, but give good representations of the molecular electrostatic potential.

9.3.3 Generalized atomic polar tensor charges

The derivative of the dipole moment with respect to the nuclear coordinates determines intensities of IR absorptions (Section 10.1.5). A central quantity in this respect is the *Atomic Polar Tensor* (APT), which for a given atom is defined in eq. (9.17).

$$\mathbf{V}^{\text{APT}} = \begin{pmatrix} \frac{\partial \mu_x}{\partial x} & \frac{\partial \mu_x}{\partial y} & \frac{\partial \mu_x}{\partial z} \\ \frac{\partial \mu_y}{\partial x} & \frac{\partial \mu_y}{\partial y} & \frac{\partial \mu_y}{\partial z} \\ \frac{\partial \mu_z}{\partial x} & \frac{\partial \mu_z}{\partial y} & \frac{\partial \mu_z}{\partial z} \end{pmatrix} \quad (9.17)$$

Such a matrix is not independent of the coordinate system, but the trace is. J. Cioslowski has proposed a definition of atomic charges as one-third of the trace over the APT, denoted *Generalized Atomic Polar Tensor* (GAPT) charges.³⁰ The charge on atom A is defined in eq. (9.18).

$$Q_A^{\text{GAPT}} = \frac{1}{3} \left(\frac{\partial \mu_x}{\partial x_A} + \frac{\partial \mu_y}{\partial y_A} + \frac{\partial \mu_z}{\partial z_A} \right) \quad (9.18)$$

Since the dipole moment itself is the first derivative of the energy with respect to an external electric field (Section 10.1.1), a calculation of GAPT charges requires the second derivative of the energy. This is a computationally expensive method for generating atomic charges, but if vibrational frequencies are calculated anyway, GAPT charges may be determined with very little additional effort. Dipole derivatives determine intensities of IR absorptions and GAPT charges are therefore directly related to experimentally observable quantities. The GAPT charges are computationally expensive to generate, and are sensitive to the amount of electron correlation in the wave function, which has limited the general use of GAPT charges.

9.4 Localized Orbitals

A Hartree–Fock wave function can be written as a single Slater determinant, composed of a set of orthonormal MOs (eqs (9.19) and (3.20)).

$$\Phi = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(1) & \phi_2(1) & \cdots & \phi_N(1) \\ \phi_1(2) & \phi_2(2) & \cdots & \phi_N(2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(N) & \phi_2(N) & \cdots & \phi_N(N) \end{vmatrix} \quad (9.19)$$

For computational purposes, it is convenient to work with canonical MOs, i.e. those that make the matrix of Lagrange multipliers diagonal, and that are eigenfunctions of the Fock operator at convergence (eq. (3.42)). This corresponds to a specific choice of a unitary (orthogonal) transformation of the occupied MOs. Once the SCF procedure has converged, however, other sets of orbitals may be chosen by forming linear combinations of the canonical MOs. The total wave function, and thus all observable properties, is independent of such a rotation of the MOs.

$$\begin{aligned} \phi' &= \mathbf{U}\phi \\ \phi'_i &= \sum_{j=1}^{N_{\text{orb}}} u_{ij} \phi_j \end{aligned} \quad (9.20)$$

The traditional view of molecular bonds is that they are due to an increased probability of finding electrons between two nuclei, as compared with a sum of the contributions of the pure atomic orbitals. The canonical MOs are delocalized over the whole molecule and do not readily reflect this, since the density between two nuclei is the result of many small contributions from many (all) the MOs. There is furthermore little similarity between MOs for systems which by chemical measures should be similar, such as a series of alkanes. The canonical MOs therefore do not reflect the concept of functional groups, nor do they readily allow identification of the bonding properties of the system.

The goal of *Localized Molecular Orbitals* (LMO) is to define MOs that are spatially confined to a relatively small volume, and therefore clearly display which atoms are bonded and furthermore have the property of being approximately constant between structurally similar units in different molecules. A set of LMOs may be defined by optimizing the expectation value of a two-electron operator Ω .³¹

$$\langle \Omega \rangle = \sum_{i=1}^{N_{\text{orb}}} \langle \phi'_i \phi'_i | \Omega | \phi'_i \phi'_i \rangle \quad (9.21)$$

The expectation value depends on the u_{ij} parameters in eq. (9.20), i.e. this is again a function optimization problem (Chapter 12). In practice, however, the localization is often done by performing a series of 2×2 orbital rotations (Section 16.2). It should be noted that the unitary transformation of the orbitals preserves the orthogonality, i.e. the resulting LMOs are also orthogonal.

Since all observable properties depend only on the total electron density, and not the individual MOs, there is no unique choice for Ω .

The *Boys* localization scheme uses the square of the distance between two electrons as the operator, and minimizes the expectation value.³²

$$\langle \Omega \rangle_{\text{Boys}} = \sum_{i=1}^{N_{\text{orb}}} \langle \phi'_i \phi'_i | (\mathbf{r}_1 - \mathbf{r}_2)^2 | \phi'_i \phi'_i \rangle \quad (9.22)$$

This corresponds to determining a set of LMOs that minimizes the spatial extent, i.e. they are as compact as possible. For extended (periodic) systems described by plane wave basis functions, the equivalent of the Boys LMOs is called *Wannier* orbitals.³³ Feng *et al.*³⁴ have shown that the Boys LMOs can be made even more compact by 10–25% by allowing the localized orbitals to be non-orthogonal, but this requires a general optimization procedure, rather than a simple unitary transformation.

The *Edmiston–Ruedenberg* localization scheme uses the inverse of the distance between two electrons as the operator, and maximizes the expectation value.³⁵

$$\langle \Omega \rangle_{\text{ER}} = \sum_{i=1}^{N_{\text{orb}}} \left\langle \phi'_i \phi'_i \left| \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \right| \phi'_i \phi'_i \right\rangle \quad (9.23)$$

This corresponds to determining a set of LMOs that maximizes the self-repulsion energy.

The *von Niessen* localization scheme uses the δ function of the distance between two electrons as the operator, and maximizes the expectation value.³⁶

$$\langle \Omega \rangle_{\text{vN}} = \sum_{i=1}^{N_{\text{orb}}} \langle \phi'_i \phi'_i | \delta(\mathbf{r}_1 - \mathbf{r}_2) | \phi'_i \phi'_i \rangle \quad (9.24)$$

This corresponds to determining a set of LMOs that maximizes the “self-charge”.

The *Pipek–Mezey* localization scheme corresponds to maximizing the sum of the Mulliken atomic charges.³⁷ The contribution to atom A is given in eq. (9.25).

$$\rho_A = \sum_{i=1}^{N_{\text{orb}}} \sum_{\alpha \in A}^{M_{\text{basis}}} \sum_{\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta i} S_{\alpha \beta} \quad (9.25)$$

The function to be maximized is given in eq. (9.26).

$$\langle \Omega \rangle_{\text{PM}} = \sum_{A=1}^{\text{Atoms}} |\rho_A|^2 \quad (9.26)$$

There is little experience with the von Niessen method but, for most molecules, the other three schemes tend to give very similar LMOs. The main exception is systems containing both σ - and π -bonds, such as ethylene. The Pipek – Mezey procedure will preserve the σ/π -separation, while the Edmiston–Ruedenberg and Boys schemes produce bent “banana” bonds. Similarly, for planar molecules that contain lone pairs (such as water or formaldehyde), the Pipek–Mezey method will produce one in-plane σ -type lone pair and one out-of-plane π -type lone pair, while the Edmiston–Ruedenberg and Boys schemes produce two equivalent “rabbit ear” lone pairs. The canonical MOs and the Boys and Pipek – Mezey LMOs for ethylene are shown in Figure 9.4 for the valence orbitals.

9.4.1 Computational considerations

It may be shown that minimization of $\langle \Omega \rangle_{\text{Boys}}$ is equivalent to maximizing the distance between centroids of the orbitals, defined by the following functional.

$$\langle \Omega' \rangle_{\text{Boys}} = \sum_{i>j}^{N_{\text{orb}}} (\langle \phi'_i | \mathbf{r} | \phi'_i \rangle - \langle \phi'_j | \mathbf{r} | \phi'_j \rangle)^2 \quad (9.27)$$




Figure 9.4 Canonical and localized molecular valence orbitals for ethylene

It is also equivalent to maximizing the distance from the (arbitrary) origin of the coordinate system, i.e. maximizing the following functional.

$$\langle \Omega' \rangle_{\text{Boys}} = \sum_i^{N_{\text{orb}}} \langle \phi'_i | \mathbf{r} | \phi'_i \rangle^2 \quad (9.28)$$

The dipole integrals in the molecular basis may be obtained from the corresponding AO integrals.

$$\langle \phi_i | \mathbf{r} | \phi_j \rangle = \sum_{\alpha}^{M_{\text{basis}}} c_{\alpha i} \left(\sum_{\beta}^{M_{\text{basis}}} c_{\beta j} \langle \chi_{\alpha} | \mathbf{r} | \chi_{\beta} \rangle \right) \quad (9.29)$$

This is a process that increases as the cube of the basis set size, and the optimization of the $\langle \Omega' \rangle_{\text{Boys}}$ function is therefore an M_{basis}^3 method. The Edmiston–Ruedenberg

localization in the above formulation requires standard two-electron integrals over MOs, analogous to those used in electron correlation methods (eq. (4.11)), and it therefore involves a computational effort that increases as M_{basis}^5 . Since only integrals involving occupied MOs are needed, the transformation is not particularly time-consuming for reasonably sized systems,³⁸ but it will ultimately require a significant effort for large systems. Recent work by Head-Gordon and coworkers³⁹ has shown that the problem can be reformulated as a series of iterative one-index transformations, which reduces the formal scaling to M_{basis}^3 . The von Niessen method may be shown to involve a computational effort that increases as M_{basis}^5 , while the Pipek–Mezey charge localization only involves overlap integrals between basis functions, and consequently has an M_{basis}^3 computational dependence.

Although the localization by an energy criterion (Edmiston–Ruedenberg) may be considered more “fundamental” than one based on distance (Boys) or atomic charge (Pipek–Mezey), the difference in computational effort means that the Boys or Pipek–Mezey procedures are often used in practice, especially since there is normally little difference in the shape of the final LMOs.

Localized molecular orbitals are generally found to reflect the usual picture of bonding, i.e. they are localized between two nuclei, or in some cases, such as diborane, extended over three nuclei. Although they indicate which atoms are bonded, they do not directly give any information about the strength of the bonds. Furthermore, localizing a set of MOs basically corresponds to determining orbitals containing electron pairs. In structures with delocalized electrons (e.g. transition structures) it may be difficult to achieve a proper localization of the MOs, and molecules with several resonance structures, such as benzene, may have more than one set of (equivalent) LMOs. It should be noted that LMOs are often used as starting points for local electron correlation methods, as discussed in Section 4.13.

9.5 Natural Orbitals

The electron density calculated from a wave function is given as the square of the function, $|\Psi|^2 = \Psi^* \Psi$. The reduced density matrix of order k , γ_k , is defined by eq. (9.30).⁴⁰

$$\gamma_k(\mathbf{r}_1, \dots, \mathbf{r}_k, \mathbf{r}'_1, \dots, \mathbf{r}'_k) = \left(\frac{N_{\text{elec}}}{k} \right) \int \Psi^*(\mathbf{r}'_1, \dots, \mathbf{r}'_k, \mathbf{r}_{k+1} \dots \mathbf{r}_{N_{\text{elec}}}) \int \Psi(\mathbf{r}_1, \dots, \mathbf{r}_k, \mathbf{r}_{k+1} \dots \mathbf{r}_{N_{\text{elec}}}) d\mathbf{r}_{k+1} \dots d\mathbf{r}_{N_{\text{elec}}} \quad (9.30)$$

Note that the coordinates for Ψ^* and Ψ are different. Of special importance in electronic structure theory are the first- and second-order reduced density matrices, $\gamma_1(\mathbf{r}_1, \mathbf{r}'_1)$ and $\gamma_2(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}'_1, \mathbf{r}'_2)$, since the Hamiltonian operator only contains one- and two-electron operators. Integrating the first-order density matrix over coordinate 1 yields the number of electrons, N_{elec} , while the integral of the second-order density matrix over coordinates 1 and 2 is $N_{\text{elec}}(N_{\text{elec}} - 1)$, i.e. the number of electron pairs. The first-order density matrix may be diagonalized, and the corresponding eigenvectors and eigenvalues are called *Natural Orbitals* (NO) and *Occupation Numbers*. The corresponding eigenfunctions for the second-order density matrix are called *Natural Geminals*. For a single-determinant RHF wave function, the first-order density matrix is identical to the density matrix used in the formation of the Fock matrix (eq. (3.52)), and the natural

orbitals have occupation numbers of either 0 or 2 (exactly). Since there are $1/2 N_{\text{elec}}$ orbitals with degenerate eigenvalues of 2, the HF natural orbitals are not uniquely defined and they may be any linear combination of the canonical orbitals. For a multi-determinant wave function (MCSCF, CI, MP or CC) the occupation numbers may assume fractional values between 0 and 2. UHF wave functions (when different from RHF) will in general also give fractional occupations.

The original definition of natural orbitals was in terms of the density matrix from a full CI wave function, i.e. the best possible for a given basis set.⁴¹ In that case, the natural orbitals have the significance that they provide the fastest convergence. In order to obtain the lowest energy for a CI expansion using only a limited set of orbitals, the natural orbitals with the largest occupation numbers should be used.

When natural orbitals are determined from a wave function that only includes a limited amount of electron correlation (i.e. not full CI), the convergence property is not rigorously guaranteed but, since most practical methods recover 80–90% of the total electron correlation, the occupation numbers provide a good guideline for how important a given orbital is. This is the reason why natural orbitals are often used for evaluating which orbitals should be included in an MCSCF wave function (Section 4.6).

9.6 Natural Atomic Orbital and Natural Bond Orbital Analysis

The concept of natural orbitals may be used for distributing electrons into atomic and molecular orbitals, and thereby deriving atomic charges and molecular bonds. The idea in the *Natural Atomic Orbital* (NAO) and *Natural Bond Orbital* (NBO) analysis developed by Weinhold and coworkers⁴² is to use the one-electron density matrix for defining the shape of the atomic orbitals in the molecular environment, and to derive molecular bonds from electron density between atoms.

Let us assume that the basis functions have been arranged such that all orbitals located on centre A are before those on centre B, which are before those on centre C, etc.

$$\chi_1^A, \chi_2^A, \chi_3^A, \dots, \chi_k^B, \chi_{k+1}^B, \chi_{k+2}^B, \dots, \chi_n^C, \chi_{n+1}^C, \chi_{n+2}^C, \dots \quad (9.31)$$

The density matrix can be written in terms of blocks of basis functions belonging to a specific centre, as shown in eq. (9.32).

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}^{AA} & \mathbf{D}^{AB} & \mathbf{D}^{AC} & \vdots \\ \mathbf{D}^{AB} & \mathbf{D}^{BB} & \mathbf{D}^{BC} & \vdots \\ \mathbf{D}^{AC} & \mathbf{D}^{BC} & \mathbf{D}^{CC} & \vdots \\ \dots & \dots & \dots & \ddots \end{pmatrix} \quad (9.32)$$

The natural atomic orbitals for atom A in the molecular environment may be defined as those that diagonalize the \mathbf{D}^{AA} block, NAOs for atom B as those that diagonalize the \mathbf{D}^{BB} block, etc. These NAOs will in general not be orthogonal, and the orbital occupation numbers will therefore not sum to the total number of electrons. To achieve a well-defined division of the electrons, the orbitals should be orthogonalized.

The NAOs will normally resemble the pure atomic orbitals (as calculated for an isolated atom), and may be divided into a “natural minimal basis” (corresponding to the

occupied atomic orbitals for the isolated atom), and a remaining set of natural “Rydberg” orbitals based on the magnitude of the occupation numbers. The minimal set of NAOs will normally be strongly occupied (i.e. having occupation numbers significantly different from zero), while the Rydberg NAO usually will be weakly occupied (i.e. having occupation numbers close to zero). There are as many NAOs as the size of the atomic basis set, and the number of Rydberg NAOs thus increases as the basis set is enlarged. It is therefore desirable that the orthogonalization procedure preserves the form of the strongly occupied orbitals as much as possible, which is achieved by using an occupancy-weighted orthogonalizing matrix. If all orbital occupancies are exactly 2 or 0, the orthogonalization is identical to the Löwdin method (eq. (9.8)). The procedure is as follows:

- (1) Each of the atomic blocks in the density matrix is diagonalized to produce a set of non-orthogonal NAOs, often denoted “pre-NAOs”.
- (2) The strongly occupied pre-NAOs for each centre are made orthogonal to all the strongly occupied pre-NAOs on the other centres by an occupancy-weighted procedure.
- (3) The weakly occupied pre-NAOs on each centre are made orthogonal to the strongly occupied NAOs on the same centre by a standard Gram – Schmidt orthogonalization.
- (4) The weakly occupied NAOs are made orthogonal to all the weakly occupied NAOs on the other centres by an occupancy-weighted procedure.




Figure 9.5 Illustration of the orthogonalization order in the NAO analysis

The final set of orthogonal orbitals are simply denoted NAOs, and the diagonal elements of the density matrix in this basis are the orbital populations. Summing all contributions from orbitals belonging to a specific centre produces the atomic charge. It is usually found that the natural minimal NAOs contribute 99+% of the electron density, and they form a very compact representation of the wave function in terms of atomic orbitals. The further advantage of the NAOs is that they are defined from the density matrix, guaranteeing that the electron occupation is between 0 and 2, and that they converge to well-defined values as the size of the basis set is increased. Furthermore, the analysis may also be performed for correlated wave functions. The disad-

vantage is that the NAOs may still extend quite far from the atom upon which they are derived, and analogously to the Mulliken approach, these NAOs may describe electron density that is near another nucleus but are counted as belonging to the nucleus upon which they are centred.

Once the density matrix has been transformed to the NAO basis, bonds between atoms may be identified from the off-diagonal blocks. The procedure involves the following steps:

- (1) NAOs for an atomic block in the density matrix that have occupation numbers very close to 2 (say > 1.999) are identified as core orbitals. Their contributions to the density matrix are removed.
- (2) NAOs for an atomic block in the density matrix that have large occupancy numbers (say > 1.90) are identified as lone pair orbitals. Their contributions to the density matrix are also removed.
- (3) Each pair of atoms (AB, AC, BC, . . .) are now considered, and the two-by-two sub-blocks of the density matrix (with the core and lone pair contributions removed) are diagonalized. Natural bond orbitals are identified as eigenvectors that have large eigenvalues (occupation numbers larger than say > 1.90).
- (4) If an insufficient number of NBOs are generated by the above procedure (sum of occupation numbers for core, lone pair and bond orbitals significantly less than the number of electrons), the criteria for accepting an NBO may be gradually lowered until a sufficiently large fraction of the electrons has been assigned to bonds. Alternatively, a search may be initiated for three-centre bonds. The contributions to the density matrix from all diatomic bonds are removed, and all three-by-three sub-blocks are diagonalized. Such three-centre bonds are quite rare, boron systems being the most notable exception.

Once NBOs have been identified, they may be written as linear combinations of the NAOs, forming a localized picture of which “atomic” orbitals are involved in the bonding.

9.7 Computational Considerations

Population analysis based on basis functions (such as Mulliken or Löwdin) require insignificant computational time. The NAO analysis involves only matrix diagonalization of small subsets of the density matrix, and also requires a negligible amount of computer time, although it is more involved than a Mulliken or Löwdin analysis. The determination of ESP fitted charges requires an evaluation of the potential at many (often several thousand) points in space, and a subsequent solution of a matrix equation for minimizing the least squares expression. For large systems, this is no longer completely trivial in terms of computer time. The AIM population analysis requires a complete topological analysis of the electron density surface, and a subsequent numerical integration of the atomic basins. For medium-sized systems and medium-quality wave functions, such an analysis may be more time-consuming than determining the wave function itself. Voronoi and Hirshfeld charges similarly require a numerical integration of the electron density, and the determination of Stewart atoms has proven to be computationally quite difficult. GAPT charges require calculation of the second derivative of the wave function, which is computationally demanding, especially for

large molecules and/or correlated wave functions. There is little doubt that these computational considerations partly explain the popularity of especially the Mulliken population analysis, despite its well-known shortcomings. For analysis purposes, the NAO procedure is an attractive method, but for modelling purposes (i.e. force field charges) ESP charges are clearly the logical choice.

9.8 Examples

Tables 9.1 and 9.2 give some examples of atomic charges and bond orders calculated by various methods as a function of the basis set at the HF level of theory. It is evident that the Mulliken and Löwdin methods do not converge as the basis set is increased, and the values in general behave unpredictably. In particular, the presence of diffuse functions leads to absurd behaviours, as the aug-cc-pVXZ basis sets illustrate for CH₄. Note also that for sufficiently large basis sets, the charge on oxygen in H₂O can be calculated to be *less* than that on carbon in CH₄! The ESP fitted charges, as well as those derived by the NAO and AIM procedures, attain well-defined values as the basis set is enlarged, and they are rather insensitive to the presence of diffuse functions. The charges assigned by these three methods, however, differ significantly, e.g. the carbon in CH₄ may be assigned charges between +0.20 (AIM) and −0.74 (NAO).

Table 9.1 Atomic charges for carbon in CH₄ ($R_{\text{CH}} = 1.092 \text{ \AA}$)

Basis	Mulliken	Löwdin	ESP fit	NAO	AIM
STO-3G	−0.25	−0.14	−0.37	−0.20	+0.22
3-21G	−0.79	−0.38	−0.44	−0.90	−0.05
6-31G(d,p)	−0.47	−0.43	−0.36	−0.88	+0.22
6-311G(2df,2pd)	+0.08	+0.12	−0.36	−0.70	+0.17
cc-pVDZ	−0.15	−0.31	−0.30	−0.80	+0.30
cc-pVTZ	−0.36	−0.02	−0.34	−0.72	+0.14
cc-pVQZ	−0.27	+0.08	−0.35	−0.74	+0.25
cc-pV5Z	−0.07	+0.17	−0.35	−0.74	+0.20
aug-cc-pVDZ	+0.55	+0.05	−0.34	−0.78	+0.31
aug-cc-pVTZ	−1.18	+0.30	−0.36	−0.72	+0.13
aug-cc-pVQZ	−0.97	+0.62	−0.35	−0.75	+0.25
aug-cc-pV5Z	−0.58	+0.70	−0.35	−0.74	+0.20

The Voronoi, Hirshfeld and Stewart charges with the 6-31G(d,p) basis set are +1.85, −0.10 and −0.58, respectively.²⁶

Table 9.2 Atomic charges for oxygen in H₂O ($R_{\text{OH}} = 0.957 \text{ \AA}$, $\theta_{\text{HOH}} = 104.5^\circ$)

Basis	Mulliken	Löwdin	ESP fit	NAO	AIM
STO-3G	−0.37	−0.25	−0.62	−0.40	−0.82
3-21G	−0.73	−0.46	−0.87	−0.87	−0.91
6-31G(d,p)	−0.67	−0.45	−0.79	−0.97	−1.23
6-311G(2df,2pd)	−0.41	+0.26	−0.73	−0.90	−1.25
cc-pVDZ	−0.31	−0.22	−0.75	−0.91	−1.26
cc-pVTZ	−0.48	+0.18	−0.74	−0.91	−1.25
cc-pVQZ	−0.53	+0.41	−0.73	−0.92	−1.24
cc-pV5Z	−0.56	+0.55	−0.73	−0.91	−1.24
aug-cc-pVDZ	−0.30	−0.01	−0.73	−0.96	−1.25
aug-cc-pVTZ	−0.44	+0.39	−0.72	−0.92	−1.27
aug-cc-pVQZ	−0.58	+0.88	−0.72	−0.92	−1.24
aug-cc-pV5Z	−0.80	+0.93	−0.72	−0.91	−1.24

The Voronoi, Hirshfeld and Stewart charges with the 6-31G(d,p) basis set are −0.72, −0.32 and −0.79, respectively.

Table 9.3 Bond orders for CH₄ and H₂O

Basis	CH ₄		H ₂ O	
	DS	AIM	DS	AIM
STO-3G	0.99	1.00	0.95	0.84
3-21G	0.93	1.00	0.83	0.81
6-31G(d,p)	0.98	1.00	0.88	0.63
cc-pVDZ	1.00	0.99	1.02	0.61
cc-pVTZ	0.98		0.99	
cc-pVQZ	0.99		1.00	
cc-pV5Z	0.99		0.99	
aug-cc-pVDZ	0.98	0.99	1.01	0.62
aug-cc-pVTZ	0.91		1.00	
aug-cc-pVQZ	0.94		0.97	
aug-cc-pV5Z	0.98		0.96	

The bond order denoted with **DS** is calculated from eq. (9.10)

References

1. S. M. Bachrach, *Rev. Comp. Chem.*, **5** (1994), 171.
2. R. S. Mulliken, *J. Chem. Phys.*, **36** (1962), 3428.
3. P.-O. Löwdin, *Adv. Quant. Chem.*, **5** (1970), 185.
4. T. Kar, L. Behera, A. B. Sannigrahi, *Chem. Phys. Lett.*, **163** (1989), 157.
5. I. Mayer, *Chem. Phys. Lett.*, **393** (2004), 209.
6. L. Pauling, *J. Am. Chem. Soc.*, **69** (1947), 542.
7. I. Mayer, *Chem. Phys. Lett.*, **97** (1983), 270.
8. P. Bultinck, R. Ponc, S. Van Damme, *J. Phys. Org. Chem.*, **18** (2005), 706.
9. D. E. Williams, *Rev. Comp. Chem.*, **2** (1991), 219.
10. M. M. Francl, L. E. Chirlian, *Rev. Comp. Chem.*, **14** (2000), 1.

11. C. I. Bayly, P. Cieplak, W. D. Cornell, P. A. Kollman, *J. Phys. Chem.*, **97** (1993), 10269.
12. C. A. Reynolds, J. W. Essex, W. G. Richards, *J. Am. Chem. Soc.*, **114** (1992), 9075.
13. M. M. Francl, C. Carey, L. E. Chilian, D. M. Gange, *J. Comp. Chem.*, **17** (1996), 367; E. Sigfridsson, U. Ryde, *J. Comp. Chem.*, **19** (1998), 377.
14. C. I. Bayly, P. Cieplak, W. D. Cornell, P. A. Kollman, *J. Phys. Chem.*, **97** (1993), 10269.
15. C. Aleman, M. Orozco, F. J. Luque, *Chem. Phys.*, **189** (1994), 573.
16. U. Koch, E. Egert, *J. Comp. Chem.*, **16** (1995), 937.
17. A. J. Stone, M. Alderton, *Mol. Phys.*, **56** (1985), 1047.
18. E. V. Tsiper, K. Burke, *J. Chem. Phys.*, **120** (2004), 1153.
19. G. G. Ferency, P. J. Winn, C. A. Reynolds, *J. Phys. Chem. A*, **101** (1997), 5446.
20. R. F. W. Bader, *Atoms in Molecules*, Clarendon Press, Oxford, 1990; R. F. W. Bader, *Chem. Rev.*, **91** (1991), 893; P. L. A. Popelier, *Atoms in Molecules: An Introduction*, Pearson Education, 1999.
21. N. Singh, R. J. Loader, P. J. O'Malley, P. L. A. Popelier, *J. Phys. Chem. A*, **110** (2006), 6498. Reproduced by permission of ACS.
22. Illustration courtesy of M. Rafat and P. L. A. Popelier.
23. C. Mei, K. E. Edgecombe, V. H. Smith Jr, A. Heilingbrunner, *Int. J. Quant. Chem.*, **48** (1993), 287.
24. C. F. Matta, N. Castillo, R. J. Boyd, *J. Phys. Chem. A*, **109** (2005), 3669.
25. P. L. A. Popelier, F. M. Aicken, *J. Am. Chem. Soc.*, **125** (2003), 1284; P. L. A. Popelier, F. M. Aicken, *Chem. Eur. J.*, **9** (2003), 1207; P. L. A. Popelier, F. M. Aicken, *Chem. Phys. Chem.*, **4** (2003), 824.
26. B. Rousseau, A. Peeters, C. Van Alsenoy, *J. Mol. Struct. Theochem.*, **538** (2001), 235.
27. F. L. Hirshfeld, *Theor. Chim. Acta*, **44** (1977), 129.
28. C. F. Guerra, J.-W. Handgraaf, E. J. Baerends, F. M. Bickelhaupt, *J. Comp. Chem.*, **25** (2003), 189.
29. A. T. B. Gilbert, P. M. W. Gill, S. W. Taylor, *J. Chem. Phys.*, **120** (2004), 7887.
30. J. Cioslowski, *J. Am. Chem. Soc.*, **111** (1989), 8333.
31. J. Pipek, P. G. Mezey, *J. Chem. Phys.*, **90** (1989), 4916; D. A. Kleier, T. A. Halgren, J. H. Hall Jr, W. N. Lipscomb, *J. Chem. Phys.*, **61** (1974), 3905.
32. S. F. Boys, *Rev. Mod. Phys.*, **32** (1960), 296.
33. N. Marzari, D. Vanderbilt, *Phys. Rev. B*, **56** (1997), 12847.
34. H. Feng, J. Bian, J. Li, W. Yang, *J. Chem. Phys.*, **120** (2004), 9458.
35. C. Edmiston, K. Ruedenberg, *J. Chem. Phys.*, **43** (1965), S97.
36. W. von Niessen, *J. Chem. Phys.*, **56** (1972), 4290.
37. J. Pipek, P. G. Mezey, *J. Chem. Phys.*, **90** (1989), 4916.
38. R. C. Raffaneti, K. Ruedenberg, C. L. Jansen, H. F. Schaefer, *Theor. Chim. Acta*, **86** (1992), 149.
39. J. E. Subotnik, Y. Shao, W. Liang, M. Head-Gordon, *J. Chem. Phys.*, **121** (2004), 9220.
40. R. G. Parr, W. Yang, *Density Functional Theory*, Oxford University Press, 1989.
41. P.-O. Löwdin, *Phys. Rev.*, **97** (1955), 1474.
42. A. E. Reed, L. A. Curtiss, F. Weinhold, *Chem. Rev.*, **88** (1988), 899.

10 Molecular Properties

The focus in Chapters 3 and 4 is on determining the wave function and its energy at a given geometry in the absence of external fields (electric or magnetic). While relative energies are certainly of interest, there are many other molecular properties that can be calculated by electronic structure methods. Most properties may be defined as the response of a wave function, an energy or an expectation value of an operator to a perturbation, where the perturbation may be any kind of operator not present in the Hamiltonian used for solving the Schrödinger equation. It may for example be terms arising in a relativistic treatment (e.g. spin–orbit interactions), which can be added as perturbations in non-relativistic theory. It may also be external fields (electric or magnetic) or an internal perturbation, such as a nuclear or electron spin. If we furthermore include “perturbations” such as adding or removing an electron, electron affinities and ionization potentials are also included in this definition. There are a few remaining properties that cannot easily be characterized as a response to a perturbation, most notably transition moments, which determine absorption intensities. These depend on matrix elements between two different wave functions.

We will here consider four types of perturbations:

- External electric field (**F**)
- External magnetic field (**B**)
- Nuclear magnetic moment (nuclear spin, **I**)
- Change in the nuclear geometry (**R**)

The first two, electric and magnetic fields, may either be *time independent*, which lead to *static* properties, or *time dependent*, leading to *dynamic* properties. Time-dependent fields are usually associated with electromagnetic radiation characterized by a frequency, and static properties may be considered as the limiting case of dynamic properties when the frequency goes to zero. We will focus on the static case here, and again concentrate on properties of a single molecule for a fixed geometry. A direct comparison with (gas phase) experimental macroscopic quantities may be done by proper averaging over for example vibrational and rotational states. We will furthermore concentrate on the electronic contribution to properties; the corresponding nuclear

contribution (if present) is normally trivial to calculate as it is independent of the wave function.¹

The nuclear magnetic moment may be considered as an artificial perturbation, since it is an inherent part of a given nucleus (isotope). In most applications, however, the nucleus is modelled as a point particle with an electric charge, and the magnetic moment needs only to be included in the Hamiltonian if the interest is in magnetic interactions involving the nucleus. These interactions are furthermore small and can be treated as a perturbative correction. One may analogously also neglect terms in the Hamiltonian involving electron spins, but one cannot neglect the electron spin in the wave function. The fermion character of the electrons leads to the requirement of wave function antisymmetry, which must be accounted for right from the outset for any theory. With a spin-free Hamiltonian, the spin dependence in the wave function can be integrated out. When properties related to magnetic interactions with the electron spin are desired, the spin-dependent terms in the Hamiltonian can be re-introduced and treated as perturbations.

There are three main methods for calculating the effect of a perturbation:

- Derivative techniques.
- Perturbation theory based on the energy.
- Perturbation theory based on expectation values of properties, often called response or propagator methods.

The derivative formulation is perhaps the easiest to understand. In this case, the energy is expanded in a Taylor series in the perturbation strength λ .

$$E(\lambda) = E(0) + \frac{\partial E}{\partial \lambda} \lambda + \frac{1}{2} \frac{\partial^2 E}{\partial \lambda^2} \lambda^2 + \frac{1}{6} \frac{\partial^3 E}{\partial \lambda^3} \lambda^3 + \dots \quad (10.1)$$

The n th-order property is the n th-order derivative of the energy, $\partial^n E / \partial \lambda^n$. Note that the perturbation is usually a vector, and the first derivative is therefore also a vector, the second derivative a matrix, the third derivative a (third-order) tensor, etc.

10.1 Examples of Molecular Properties

10.1.1 External electric field

The interaction of an electronic charge distribution $\rho(\mathbf{r})$ with an electric potential $\phi(\mathbf{r})$ gives an energy contribution.

$$E = \int \rho(\mathbf{r}) \phi(\mathbf{r}) d\mathbf{r} \quad (10.2)$$

Since the electric field ($\mathbf{F} = -\partial\phi/\partial\mathbf{r}$) is normally fairly uniform at the molecular level, it is useful to write E as a multipole expansion.

$$E = q\phi - \boldsymbol{\mu}\mathbf{F} - \frac{1}{2}\mathbf{Q}\mathbf{F}' - \dots \quad (10.3)$$

Here q is the net charge (monopole), $\boldsymbol{\mu}$ is the (electric) dipole moment, \mathbf{Q} is the quadrupole moment, and \mathbf{F} and \mathbf{F}' are the field and field gradient ($\partial\mathbf{F}/\partial\mathbf{r}$), respectively. The dipole moment and electric field are vectors, and the $\boldsymbol{\mu}\mathbf{F}$ term should be inter-

puted as the dot product ($\boldsymbol{\mu}\mathbf{F} = \mu_x F_x + \mu_y F_y + \mu_z F_z$). The quadrupole moment and field gradient are 3×3 matrices, and $\mathbf{Q}\mathbf{F}'$ denotes the sum of all product terms. For an external field it is rarely necessary to go beyond the quadrupole term, but for molecular interactions the octupole moment may also be important (it is for example the first non-vanishing moment for spherical molecules such as CH_4).

In the absence of an external field, the unperturbed dipole and quadrupole moments may be calculated from the electronic wave function as simple expectation values.

$$\begin{aligned}\boldsymbol{\mu}_0 &= -\langle \Psi | \mathbf{r} | \Psi \rangle \\ \mathbf{Q}_0 &= \langle \Psi | \mathbf{r} \mathbf{r}^t | \Psi \rangle\end{aligned}\quad (10.4)$$

The minus sign for the dipole moment arises from the negative charge on the electron. The superscript t denotes a transposition of the \mathbf{r} -vector, i.e. converting it from a column to a row vector. The $\mathbf{r}\mathbf{r}^t$ notation therefore indicates the outer product of \mathbf{r} with itself, and the quadrupole moment is thus a 3×3 matrix, where the Q_{xy} component is calculated as the expectation value of xy .

The presence of a field influences the wave function and leads to induced dipole, quadrupole, etc., moments. For the dipole moment this may be written as in eq. (10.5).

$$\boldsymbol{\mu} = \boldsymbol{\mu}_0 + \boldsymbol{\alpha}\mathbf{F} + \frac{1}{2}\boldsymbol{\beta}\mathbf{F}^2 + \frac{1}{6}\boldsymbol{\gamma}\mathbf{F}^3 \dots \quad (10.5)$$

Here $\boldsymbol{\mu}_0$ is the permanent dipole moment, $\boldsymbol{\alpha}$ is the (dipole) *polarizability*, $\boldsymbol{\beta}$ is the (first) *hyperpolarizability*, $\boldsymbol{\gamma}$ is the second hyperpolarizability, etc. The quadrupole moment may similarly be expanded in the field by means of a quadrupole polarizability, hyperpolarizability, etc.

For a homogeneous field (i.e. the field gradient and higher derivatives are zero), the total energy of a neutral molecule may be written as a Taylor expansion, where all the derivatives are evaluated at $\mathbf{F} = 0$. There will be a derivative for each individual component of the field, which rapidly leads to a large number of indices and summations in a proper mathematical formulation. In order to avoid this notational cluttering, we will adopt a slightly non-standard notation where the field is indicated by a vector notation, implying that derivatives should be taken along all the individual field components.

$$E(\mathbf{F}) = E(0) + \left. \frac{\partial E}{\partial \mathbf{F}} \right|_{\mathbf{F}=0} \mathbf{F} + \frac{1}{2} \left. \frac{\partial^2 E}{\partial \mathbf{F}^2} \right|_{\mathbf{F}=0} \mathbf{F}^2 + \frac{1}{6} \left. \frac{\partial^3 E}{\partial \mathbf{F}^3} \right|_{\mathbf{F}=0} \mathbf{F}^3 + \frac{1}{24} \left. \frac{\partial^4 E}{\partial \mathbf{F}^4} \right|_{\mathbf{F}=0} \mathbf{F}^4 + \dots \quad (10.6)$$

According to eq. (10.3) we also have that $\partial E / \partial \mathbf{F} = -\boldsymbol{\mu}$, where $\boldsymbol{\mu}$ is given by the expression in eq. (10.5). Differentiation of eq. (10.6) with respect to \mathbf{F} gives eq. (10.7).

$$\boldsymbol{\mu} = -\left. \frac{\partial E}{\partial \mathbf{F}} \right|_{\mathbf{F}=0} - \left. \frac{\partial^2 E}{\partial \mathbf{F}^2} \right|_{\mathbf{F}=0} \mathbf{F} - \frac{1}{2} \left. \frac{\partial^3 E}{\partial \mathbf{F}^3} \right|_{\mathbf{F}=0} \mathbf{F}^2 - \frac{1}{6} \left. \frac{\partial^4 E}{\partial \mathbf{F}^4} \right|_{\mathbf{F}=0} \mathbf{F}^3 + \dots \quad (10.7)$$

Comparing eqs (10.5) and (10.7) shows that the first derivative is the (permanent) dipole moment $\boldsymbol{\mu}_0$, the second derivative is the polarizability $\boldsymbol{\alpha}$, the third derivative is the hyperpolarizability $\boldsymbol{\beta}$, etc.

$$\boldsymbol{\mu}_0 = -\left. \frac{\partial E}{\partial \mathbf{F}} \right|_{\mathbf{F}=0} ; \boldsymbol{\alpha} = -\left. \frac{\partial^2 E}{\partial \mathbf{F}^2} \right|_{\mathbf{F}=0} ; \boldsymbol{\beta} = -\left. \frac{\partial^3 E}{\partial \mathbf{F}^3} \right|_{\mathbf{F}=0} ; \boldsymbol{\gamma} = -\left. \frac{\partial^4 E}{\partial \mathbf{F}^4} \right|_{\mathbf{F}=0} \quad (10.8)$$

10.1.2 External magnetic field

The interaction with a magnetic field may similarly be written in terms of magnetic dipole, quadrupole, etc., moments (there is no magnetic monopole, corresponding to electric charge). Since the magnetic interaction is substantially smaller in magnitude than the electric, only the dipole term is normally considered.

$$E = -\mathbf{m}\mathbf{B} - \frac{1}{2}\xi\mathbf{B}^2 - \dots \quad (10.9)$$

The dipole moment \mathbf{m}_0 for an unperturbed system depends on the total angular momentum, which may be written in terms of the orbital angular momentum operator \mathbf{L}_G and the total electron spin \mathbf{S} .

$$\begin{aligned} \mathbf{m}_0 &= -\frac{1}{2}\langle\Psi|\mathbf{L}_G + g_e\mathbf{S}|\Psi\rangle \\ \mathbf{L}_G &= (\mathbf{r} - \mathbf{R}_G) \times \mathbf{p} \end{aligned} \quad (10.10)$$

Here \mathbf{R}_G is the gauge origin (discussed in Section 10.7), and the electronic g_e -factor is a constant approximately equal to 2.0023. The orbital part of the permanent magnetic dipole moment will be zero for all non-degenerate wave functions (i.e. belonging to A, B or Σ representations), since the \mathbf{L}_G operator is purely imaginary ($\mathbf{p} = -i\nabla$) and the wave function in such cases is real. Similarly, only open-shell states (doublet, triplet, etc.) have the spin part of the magnetic dipole moment different from zero. Since the large majority of stable molecules are closed shell singlets, it follows that permanent magnetic dipole moments are quite rare. The presence of a field, however, may induce a magnetic dipole moment, with the quantity corresponding to the electric polarizability being the *magnetizability* ξ (the corresponding macroscopic quantity is called the magnetic *susceptibility* χ).

$$\mathbf{m} = \mathbf{m}_0 + \xi\mathbf{B} + \dots \quad (10.11)$$

The energy can again be expanded in a Taylor series.

$$E(\mathbf{B}) = E(\mathbf{0}) + \left.\frac{\partial E}{\partial \mathbf{B}}\right|_{\mathbf{B}=\mathbf{0}} \mathbf{B} + \frac{1}{2} \left.\frac{\partial^2 E}{\partial \mathbf{B}^2}\right|_{\mathbf{B}=\mathbf{0}} \mathbf{B}^2 + \dots \quad (10.12)$$

As for the electric field, this leads to the definition of the dipole and magnetizability as first and second derivatives of the total energy with respect to the magnetic field.

$$\mathbf{m}_0 = -\left.\frac{\partial E}{\partial \mathbf{B}}\right|_{\mathbf{B}=\mathbf{0}}; \quad \xi = -\left.\frac{\partial^2 E}{\partial \mathbf{B}^2}\right|_{\mathbf{B}=\mathbf{0}} \quad (10.13)$$

10.1.3 Internal magnetic moments

The perturbation can also be an internal magnetic moment \mathbf{I} , arising from a nuclear spin (the $g_A\mu_N$ factor for converting from spin to magnetic moment has been neglected here).

$$\begin{aligned} E(\mathbf{I}_1, \mathbf{I}_2, \dots) &= E(\mathbf{0}) + \frac{\partial E}{\partial \mathbf{I}_1} \mathbf{I}_1 + \frac{\partial E}{\partial \mathbf{I}_2} \mathbf{I}_2 + \frac{1}{2} \frac{\partial^2 E}{\partial \mathbf{I}_1 \partial \mathbf{I}_2} \mathbf{I}_1 \mathbf{I}_2 + \dots \\ E(\mathbf{I}_1, \mathbf{I}_2, \dots) &= E(\mathbf{0}) - \mathbf{A}_1 \mathbf{I}_1 - \mathbf{A}_2 \mathbf{I}_2 - h \mathbf{J}_{12} \mathbf{I}_1 \mathbf{I}_2 + \dots \end{aligned} \quad (10.14)$$

The first derivative is the nuclear–electron hyperfine coupling constant **A**, while the second derivative with respect to two different nuclear spins is the NMR coupling constant **J** (Planck’s constant appears owing to the convention of reporting coupling constants in Hertz, and the factor of $1/2$ disappears since we implicitly only consider distinct pairs of nuclei). The corresponding interaction between two electron spins determines the zero field splitting of the individual components of a triplet (or higher multiplet) state in the absence of a magnetic field.

10.1.4 Geometry change

The change in energy for moving a nucleus can also be written as a Taylor expansion.

$$E(\mathbf{R}) = E(\mathbf{R}_0) + \frac{\partial E}{\partial \mathbf{R}}(\mathbf{R} - \mathbf{R}_0) + \frac{1}{2} \frac{\partial^2 E}{\partial \mathbf{R}^2}(\mathbf{R} - \mathbf{R}_0)^2 + \frac{1}{6} \frac{\partial^3 E}{\partial \mathbf{R}^3}(\mathbf{R} - \mathbf{R}_0)^3 + \dots \quad (10.15)$$

$$E(\mathbf{R}) = E(\mathbf{R}_0) + \mathbf{g}(\mathbf{R} - \mathbf{R}_0) + \frac{1}{2} \mathbf{H}(\mathbf{R} - \mathbf{R}_0)^2 + \frac{1}{6} \mathbf{K}(\mathbf{R} - \mathbf{R}_0)^3 + \dots$$

The first derivative is the gradient **g**, the second derivative is the force constant (Hessian) **H**, the third derivative is the anharmonicity **K**, etc. If the \mathbf{R}_0 geometry is a stationary point (**g** = **0**) the force constant matrix may be used for evaluating harmonic vibrational frequencies and normal coordinates, **q**, as discussed in Section 16.2.2. If higher order terms are included in the expansion, it is possible also to determine anharmonic frequencies and phenomena such as Fermi resonance.

10.1.5 Mixed derivatives

Mixed derivatives refer to cross terms if the energy is expanded in more than one perturbation. There are many such mixed derivatives that translate into molecular properties, a few of which are given below.

The change in the dipole moment with respect to a geometry displacement along a normal coordinate **q** is related to the intensity of an IR absorption. In the so-called double harmonic approximation (terminating the expansion at first order in the electric field and geometry), the intensity is (except for some constants) given by eq. (10.16).

$$\text{IR intensity} \propto \left(\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{q}} \right)^2 \propto \left(\frac{\partial^2 E}{\partial \mathbf{R} \partial \mathbf{F}} \right)^2 \quad (10.16)$$

Only fundamental bands can have an intensity different from zero in the double harmonic approximation. Including higher order terms in the expansion allows the calculation of intensities of overtone bands, as well as adding contributions to the fundamental bands.

The intensity of a Raman band in the harmonic approximation is given by the derivative of the polarizability with respect to a normal coordinate.

$$\text{Raman intensity} \propto \left(\frac{\partial \boldsymbol{\alpha}}{\partial \mathbf{q}} \right)^2 \propto \left(\frac{\partial^3 E}{\partial \mathbf{R} \partial \mathbf{F}^2} \right)^2 \quad (10.17)$$

The mixed derivative of an external and a nuclear magnetic field (nuclear spin) is the NMR shielding tensor σ .

$$\text{NMR shielding} \propto \left(\frac{\partial^2 E}{\partial \mathbf{B} \partial \mathbf{I}} \right) \quad (10.18)$$

The corresponding quantity related to the electron spin is the ESR \mathbf{g} -tensor.

Table 10.1 gives some examples of properties that may be calculated from derivatives of a certain order with respect to the above four perturbations.

$$\text{Property} \propto \frac{\partial^{n_F+n_B+n_I+n_R} E}{\partial \mathbf{F}^{n_F} \partial \mathbf{B}^{n_B} \partial \mathbf{I}^{n_I} \partial \mathbf{R}^{n_R}} \quad (10.19)$$

All of these properties can be calculated at various levels of sophistication (electron correlation and basis sets). It should be noted that dynamic properties, where one or more of the external electric and/or magnetic fields are time dependent, may involve one or several different frequencies. Time-dependent properties are discussed in Section 10.9.

Table 10.1 Examples of properties that may be calculated as derivatives of the energy

n_F	n_B	n_I	n_R	Property
0	0	0	0	Energy
1	0	0	0	Electric dipole moment
0	1	0	0	Magnetic dipole moment
0	0	1	0	Hyperfine coupling constant
0	0	0	1	Molecular (nuclear) gradient
2	0	0	0	Electric polarizability
0	2	0	0	Magnetizability
0	0	2	0	Nuclear spin–spin coupling
0	0	0	2	Harmonic vibrational frequencies
1	0	0	1	Infrared absorption intensities
1	1	0	0	Optical rotation, circular dichroism
0	1	1	0	Nuclear magnetic shielding
3	0	0	0	(first) Electric hyperpolarizability
0	3	0	0	(first) Hypermagnetizability
0	0	0	3	(cubic) Anharmonic corrections to vibrational frequencies
2	0	0	1	Raman intensities
3	0	0	1	Hyper-Raman effects
2	1	0	0	Magnetic circular dichroism (Faraday effect)
1	0	0	2	Infrared intensities for overtone and combination bands
4	0	0	0	(second) Electric hyperpolarizability
0	4	0	0	(second) Hypermagnetizability
0	0	0	4	(quartic) Anharmonic corrections to vibrational frequencies
2	0	0	2	Raman intensities for overtone and combination bands
2	2	0	0	Cotton–Mutton effect

10.2 Perturbation Methods

Let us first look at some general features. The presence of a perturbation will give rise to extra terms in the Hamiltonian, and we will in the following need to consider operators that are both linear and quadratic in the perturbation.

$$\mathbf{H} = \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 \quad (10.20)$$

\mathbf{H}_0 is the normal electronic Hamiltonian operator, and the perturbations are described by the operators \mathbf{P}_1 and \mathbf{P}_2 , with λ determining the strength. Based on an expansion in *exact* wave functions, Rayleigh–Schrödinger perturbation theory (Section 4.8) gives the first- and second-order energy corrections.

$$\begin{aligned} W_1 &= \lambda \langle \Psi_0 | \mathbf{P}_1 | \Psi_0 \rangle \\ W_2 &= \lambda^2 \left[\langle \Psi_0 | \mathbf{P}_2 | \Psi_0 \rangle + \sum_{i \neq 0} \frac{\langle \Psi_0 | \mathbf{P}_1 | \Psi_i \rangle \langle \Psi_i | \mathbf{P}_1 | \Psi_0 \rangle}{E_0 - E_i} \right] \end{aligned} \quad (10.21)$$

The first-order term is identical to eq. (4.37), while the second-order equation corresponds to eq. (4.39) with an additional term involving the expectation value of \mathbf{P}_2 over the unperturbed wave function. The first-order energy correction is identified with the first-order property, the second-order correction is the second-order property, etc. Although these expressions only hold for exact wave functions, they may also be used for approximative wave functions. The methodology of how the general expressions can be reduced to formulas involving molecular integrals is analogous to that used in Section 3.3.

The first-order term is simply the expectation value of the perturbation operator over the unperturbed wave function, and is easy to calculate. The second-order property, however, involves a sum over all excited states. In some cases, mainly associated with semi-empirical methods, second-order properties are evaluated directly from eq. (10.21), known as *Sum Over States* (SOS) methods. Since this involves a determination of all excited states, it is very inefficient for *ab initio* methods. A computationally efficient way of calculating such properties is by means of response methods (Section 10.9).

10.3 Derivative Techniques

Derivative techniques consider the energy in the presence of the perturbation, perform an analytical differentiation of the energy n times to derive a formula for the n th-order property, and let the perturbation strength go to zero.

Let us write the energy as in eq. (10.22).

$$E(\lambda) = \langle \Psi(\lambda) | \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 | \Psi(\lambda) \rangle \quad (10.22)$$

This is strictly true for HF, MCSCF and CI wave functions, and can be generalized to MP and CC methods as shown in Section 10.4. The perturbation-dependent terms in the operator are written explicitly, while the wave function dependence is implicit, via the parameterization (orbital and state coefficients) and possibly also the basis functions.

The first derivative of the energy can be written as in eq. (10.23).

$$\begin{aligned} \frac{\partial E}{\partial \lambda} = & \left\langle \frac{\partial \Psi}{\partial \lambda} \middle| \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 \middle| \Psi \right\rangle + \langle \Psi | \mathbf{P}_1 + 2\lambda \mathbf{P}_2 | \Psi \rangle + \\ & \langle \Psi | \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 \middle| \frac{\partial \Psi}{\partial \lambda} \rangle \end{aligned} \quad (10.23)$$

For real wave functions the first and third terms are identical. Letting the perturbation strength go to zero yields eq. (10.24).

$$\left. \frac{\partial E}{\partial \lambda} \right|_{\lambda=0} = \langle \Psi_0 | \mathbf{P}_1 | \Psi_0 \rangle + 2 \left\langle \frac{\partial \Psi_0}{\partial \lambda} \middle| \mathbf{H}_0 | \Psi_0 \right\rangle \quad (10.24)$$

The wave function depends on the perturbation indirectly, via parameters in the wave function (\mathbf{C}), and possibly also the basis functions (χ). The wave function parameters may be orbital coefficients (HF), state coefficients (CI, MP, CC) or both (MCSCF).

$$\frac{\partial \Psi}{\partial \lambda} = \frac{\partial \Psi}{\partial \mathbf{C}} \frac{\partial \mathbf{C}}{\partial \lambda} + \frac{\partial \Psi}{\partial \chi} \frac{\partial \chi}{\partial \lambda} \quad (10.25)$$

Assuming for the moment that the basis functions are independent of the perturbation ($\partial \chi / \partial \lambda = 0$), the derivative (10.24) may be written as in eq. (10.26).

$$\left. \frac{\partial E}{\partial \lambda} \right|_{\lambda=0} = \langle \Psi_0 | \mathbf{P}_1 | \Psi_0 \rangle + 2 \frac{\partial \mathbf{C}}{\partial \lambda} \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \middle| \mathbf{H}_0 | \Psi_0 \right\rangle \quad (10.26)$$

If the wave function is variationally optimized with respect to *all* parameters (HF or MCSCF, but *not* CI), the last term disappears since the energy is stationary with respect to a variation of the MO/state coefficients (\mathbf{H}_0 , \mathbf{P}_1 and \mathbf{P}_2 do not depend on the parameters \mathbf{C}).

$$\begin{aligned} \frac{\partial E}{\partial \mathbf{C}} = \frac{\partial}{\partial \mathbf{C}} \langle \Psi | \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 | \Psi \rangle &= 2 \left\langle \frac{\partial \Psi}{\partial \mathbf{C}} \middle| \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 | \Psi \right\rangle \\ \left. \frac{\partial E}{\partial \mathbf{C}} \right|_{\lambda=0} &= 2 \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \middle| \mathbf{H}_0 | \Psi_0 \right\rangle = 0 \end{aligned} \quad (10.27)$$

Variational wave functions thus obey the *Hellmann–Feynman* theorem.

$$\frac{\partial}{\partial \lambda} \langle \Psi | \mathbf{H} | \Psi \rangle = \langle \Psi | \frac{\partial \mathbf{H}}{\partial \lambda} | \Psi \rangle \quad (10.28)$$

In such cases, the expression from first-order perturbation theory (10.21) yields a result identical to the first derivative of the energy with respect to λ . For wave functions that are not completely optimized with respect to all parameters (CI, MP or CC), the Hellmann–Feynman theorem does not hold, and a first-order property calculated as an expectation value will not be identical to that obtained as an energy derivative. Since the Hellmann–Feynman theorem holds for an exact wave function, the difference between the two values becomes smaller as the quality of an approximate wave function increases. However, for practical applications the difference is not negligible.

It has been argued that the derivative technique resembles the physical experiment more, and consequently formula (10.24) should be preferred over (10.21).

The second derivative of the energy can for a real wave function be written as in eq. (10.29).

$$\begin{aligned} \frac{1}{2} \frac{\partial^2 E}{\partial \lambda^2} = & \left\langle \frac{\partial^2 \Psi}{\partial \lambda^2} \left| \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 \right| \Psi \right\rangle + 2 \left\langle \frac{\partial \Psi}{\partial \lambda} \left| \mathbf{P}_1 + 2\lambda \mathbf{P}_2 \right| \Psi \right\rangle + \\ & \left\langle \frac{\partial \Psi}{\partial \lambda} \left| \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 \right| \frac{\partial \Psi}{\partial \lambda} \right\rangle + \langle \Psi | \mathbf{P}_2 | \Psi \rangle \end{aligned} \quad (10.29)$$

In the limit of the perturbation strength going to zero this reduces to eq. (10.30).

$$\begin{aligned} \frac{1}{2} \frac{\partial^2 E}{\partial \lambda^2} \Big|_{\lambda=0} = & \left\langle \frac{\partial^2 \Psi_0}{\partial \lambda^2} \left| \mathbf{H}_0 \right| \Psi_0 \right\rangle + 2 \left\langle \frac{\partial \Psi_0}{\partial \lambda} \left| \mathbf{P}_1 \right| \Psi_0 \right\rangle + \\ & \left\langle \frac{\partial \Psi_0}{\partial \lambda} \left| \mathbf{H}_0 \right| \frac{\partial \Psi_0}{\partial \lambda} \right\rangle + \langle \Psi_0 | \mathbf{P}_2 | \Psi_0 \rangle \end{aligned} \quad (10.30)$$

The implicit wave function dependence on \mathbf{C} allows the derivative to be written as in eq. (10.31).

$$\begin{aligned} \frac{1}{2} \frac{\partial^2 E}{\partial \lambda^2} \Big|_{\lambda=0} = & \frac{\partial^2 \mathbf{C}}{\partial \lambda^2} \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \left| \mathbf{H}_0 \right| \Psi_0 \right\rangle + \left(\frac{\partial \mathbf{C}}{\partial \lambda} \right)^2 \left\langle \frac{\partial^2 \Psi_0}{\partial \mathbf{C}^2} \left| \mathbf{H}_0 \right| \Psi_0 \right\rangle + \\ & 2 \left(\frac{\partial \mathbf{C}}{\partial \lambda} \right) \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \left| \mathbf{P}_1 \right| \Psi_0 \right\rangle + \left(\frac{\partial \mathbf{C}}{\partial \lambda} \right)^2 \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \left| \mathbf{H}_0 \right| \frac{\partial \Psi_0}{\partial \mathbf{C}} \right\rangle + \langle \Psi_0 | \mathbf{P}_2 | \Psi_0 \rangle \end{aligned} \quad (10.31)$$

For a variationally optimized wave function, the first term is again zero (eq. (10.27)). Furthermore, the second term, which involves calculation of the second derivative of the wave function with respect to the parameters, can be avoided. This can be seen by differentiating the stationary condition eq. (10.27) with respect to the perturbation.

$$\begin{aligned} & \frac{\partial}{\partial \lambda} \left\langle \frac{\partial \Psi}{\partial \mathbf{C}} \left| \mathbf{H}_0 + \lambda \mathbf{P}_1 + \lambda^2 \mathbf{P}_2 \right| \Psi_0 \right\rangle \Big|_{\lambda=0} \\ & = \left(\frac{\partial \mathbf{C}}{\partial \lambda} \right) \left\langle \frac{\partial^2 \Psi_0}{\partial \mathbf{C}^2} \left| \mathbf{H}_0 \right| \Psi_0 \right\rangle + \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \left| \mathbf{P}_1 \right| \Psi_0 \right\rangle + \left(\frac{\partial \mathbf{C}}{\partial \lambda} \right) \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \left| \mathbf{H}_0 \right| \frac{\partial \Psi_0}{\partial \mathbf{C}} \right\rangle = 0 \end{aligned} \quad (10.32)$$

The second derivative in eq. (10.31) therefore reduces to eq. (10.33).

$$\frac{1}{2} \frac{\partial^2 E}{\partial \lambda^2} \Big|_{\lambda=0} = \left(\frac{\partial \mathbf{C}}{\partial \lambda} \right) \left\langle \frac{\partial \Psi_0}{\partial \mathbf{C}} \left| \mathbf{P}_1 \right| \Psi_0 \right\rangle + \langle \Psi_0 | \mathbf{P}_2 | \Psi_0 \rangle \quad (10.33)$$

In a more compact notation, this can be written as eq. (10.34).

$$\frac{1}{2} \frac{\partial^2 E}{\partial \lambda^2} \Big|_{\lambda=0} = \left\langle \frac{\partial \Psi_0}{\partial \lambda} \left| \mathbf{P}_1 \right| \Psi_0 \right\rangle + \langle \Psi_0 | \mathbf{P}_2 | \Psi_0 \rangle \quad (10.34)$$

This shows that only the first-order change in the wave function is necessary. For exact wave functions eq. (10.34) becomes identical to the perturbation expression (10.21),

since the first derivative of the wave function may then be expanded in a complete set of eigenfunctions (eq. (4.36)).

$$\frac{\partial \Psi_0}{\partial \lambda} = \sum_{i=1}^{\infty} a_i \Psi_i \quad a_i = \frac{\langle \Psi_i | \mathbf{P}_1 | \Psi_0 \rangle}{E_0 - E_i} \quad (10.35)$$

10.4 Lagrangian Techniques

For variationally optimized wave functions (HF or MCSCF) there is a $2n + 1$ rule, analogous to the perturbational energy expression (eq. (4.35)): knowledge of the n th derivative (also called the response) of the wave function is sufficient for calculating a property to order $2n + 1$. For non-variational wave functions eq. (10.30) suggests that the n th-order wave function response is required for calculating the n th-order property. This may be avoided, however, by a technique first illustrated for CISD geometry derivatives by Handy and Schaefer, often referred to as the *Z-vector* method.² It has later been generalized to cover other types of wave functions and derivatives by formulating it in terms of a Lagrange function.³

The idea is to construct a Lagrange function that has the same energy as the non-variational wave function but which is variational in all parameters. Consider for example a CI wave function, which is variational in the state coefficients (**a**) but not in the orbital coefficients (**c**), since they are determined by the stationary condition for the HF wave function (note that we employ lower case **c** for the orbital coefficients but capital **C** to denote all wave function parameters, i.e. **C** contains both **a** and **c**).

$$\begin{aligned} \frac{\partial E_{\text{CI}}}{\partial \mathbf{a}} &= \frac{\partial}{\partial \mathbf{a}} \langle \Psi_{\text{CI}}(\mathbf{a}, \mathbf{c}) | \mathbf{H} | \Psi_{\text{CI}}(\mathbf{a}, \mathbf{c}) \rangle = 2 \left\langle \frac{\partial \Psi_{\text{CI}}}{\partial \mathbf{a}} \middle| \mathbf{H} | \Psi_{\text{CI}} \right\rangle = 0 \\ \frac{\partial E_{\text{CI}}}{\partial \mathbf{c}} &= \frac{\partial}{\partial \mathbf{c}} \langle \Psi_{\text{CI}}(\mathbf{a}, \mathbf{c}) | \mathbf{H} | \Psi_{\text{CI}}(\mathbf{a}, \mathbf{c}) \rangle = 2 \left\langle \frac{\partial \Psi_{\text{CI}}}{\partial \mathbf{c}} \middle| \mathbf{H} | \Psi_{\text{CI}} \right\rangle \neq 0 \\ \frac{\partial E_{\text{HF}}}{\partial \mathbf{c}} &= \frac{\partial}{\partial \mathbf{c}} \langle \Psi_{\text{HF}}(\mathbf{a}, \mathbf{c}) | \mathbf{H} | \Psi_{\text{HF}}(\mathbf{a}, \mathbf{c}) \rangle = 2 \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \middle| \mathbf{H} | \Psi_{\text{HF}} \right\rangle = 0 \end{aligned} \quad (10.36)$$

Consider now the Lagrange function given in eq. (10.37).

$$L_{\text{CI}} = E_{\text{CI}} + \boldsymbol{\kappa} \frac{\partial E_{\text{HF}}}{\partial \mathbf{c}} \quad (10.37)$$

Here $\boldsymbol{\kappa}$ contains a set of Lagrange multipliers. The derivatives of the Lagrange function with respect to **a**, **c** and $\boldsymbol{\kappa}$ are given in eq. (10.38).

$$\begin{aligned} \frac{\partial L_{\text{CI}}}{\partial \mathbf{a}} &= \frac{\partial E_{\text{CI}}}{\partial \mathbf{a}} = 0 \\ \frac{\partial L_{\text{CI}}}{\partial \mathbf{c}} &= \frac{\partial E_{\text{HF}}}{\partial \mathbf{c}} = 0 \\ \frac{\partial L_{\text{CI}}}{\partial \boldsymbol{\kappa}} &= \frac{\partial E_{\text{CI}}}{\partial \boldsymbol{\kappa}} + \boldsymbol{\kappa} \frac{\partial^2 E_{\text{HF}}}{\partial \mathbf{c}^2} = 0 \end{aligned} \quad (10.38)$$

The first two derivatives are zero due to the properties of the CI and HF wave functions, eq. (10.36). The last equation is zero by virtue of the Lagrange multipliers, i.e. we choose $\boldsymbol{\kappa}$ such that $\partial L_{\text{CI}} / \partial \boldsymbol{\kappa} = 0$. It may be written more explicitly in eq. (10.39).

$$\frac{\partial L_{\text{CI}}}{\partial \mathbf{c}} = 2 \left\langle \frac{\partial \Psi_{\text{CI}}}{\partial \mathbf{c}} \middle| \mathbf{H} \middle| \Psi_{\text{CI}} \right\rangle + 2 \kappa \left[\left\langle \frac{\partial^2 \Psi_{\text{HF}}}{\partial \mathbf{c}^2} \middle| \mathbf{H} \middle| \Psi_{\text{HF}} \right\rangle + \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \middle| \mathbf{H} \middle| \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \right\rangle \right] = 0 \quad (10.39)$$

Note that no new operators are involved, only derivatives of the CI or HF wave function with respect to the MO coefficients. The matrix elements can thus be calculated from the same integrals as the energy itself, as discussed in Sections 3.3 and 4.2.1.

The derivative with respect to a perturbation can now be written as in eq. (10.40).

$$\frac{\partial L_{\text{CI}}}{\partial \lambda} = \frac{\partial E_{\text{CI}}}{\partial \lambda} + \kappa \frac{\partial}{\partial \lambda} \left(\frac{\partial E_{\text{HF}}}{\partial \mathbf{c}} \right) = 0 \quad (10.40)$$

Expanding out the terms gives eq. (10.41).

$$\begin{aligned} \frac{\partial L_{\text{CI}}}{\partial \lambda} &= \frac{\partial}{\partial \lambda} \langle \Psi_{\text{CI}} | \mathbf{H} | \Psi_{\text{CI}} \rangle + 2 \kappa \frac{\partial}{\partial \lambda} \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \middle| \mathbf{H} \middle| \Psi_{\text{HF}} \right\rangle \\ \frac{\partial L_{\text{CI}}}{\partial \lambda} &= \langle \Psi_{\text{CI}} | \frac{\partial \mathbf{H}}{\partial \lambda} | \Psi_{\text{CI}} \rangle + 2 \frac{\partial \mathbf{a}}{\partial \lambda} \left\langle \frac{\partial \Psi_{\text{CI}}}{\partial \mathbf{a}} \middle| \mathbf{H} \middle| \Psi_{\text{CI}} \right\rangle + 2 \frac{\partial \mathbf{c}}{\partial \lambda} \left\langle \frac{\partial \Psi_{\text{CI}}}{\partial \mathbf{c}} \middle| \mathbf{H} \middle| \Psi_{\text{CI}} \right\rangle + \\ &\quad \kappa \left[\left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \middle| \frac{\partial \mathbf{H}}{\partial \lambda} \middle| \Psi_{\text{HF}} \right\rangle + \frac{\partial \mathbf{c}}{\partial \lambda} \left(\left\langle \frac{\partial^2 \Psi_{\text{HF}}}{\partial \mathbf{c}^2} \middle| \mathbf{H} \middle| \Psi_{\text{HF}} \right\rangle + \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \middle| \mathbf{H} \middle| \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \right\rangle \right) \right] \end{aligned} \quad (10.41)$$

The second term disappears since the CI wave function is variational in the state coefficients, eq. (10.36). The three terms involving the derivative of the MO coefficients ($\partial \mathbf{c} / \partial \lambda$) also disappear owing to our choice of the Lagrange multipliers, eq. (10.39). If we furthermore adapt the definition that $\partial \mathbf{H} / \partial \lambda = \mathbf{P}_1$ (eq. (10.20)), the final derivative may be written as eq. (10.42).

$$\frac{\partial L_{\text{CI}}}{\partial \lambda} = \langle \Psi_{\text{CI}} | \mathbf{P}_1 | \Psi_{\text{CI}} \rangle + \kappa \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{c}} \middle| \mathbf{P}_1 \middle| \Psi_{\text{HF}} \right\rangle \quad (10.42)$$

Here the Lagrange multipliers κ are determined from eq. (10.39).

What has been accomplished? The original expression (10.26) contains the derivative of the MO coefficients with respect to the perturbation ($\partial \mathbf{c} / \partial \lambda$), which can be obtained by solving the CPHF equations (Section 10.5 below). For geometry derivatives, for example, there will be $3N_{\text{atom}}$ different perturbations, i.e. we need to solve $3N_{\text{atom}}$ sets of CPHF equations. The Lagrange expression (10.42), on the other hand, contains a set of Lagrange multipliers κ that are *independent* of the perturbation, i.e. we need only to solve one equation for κ , eq. (10.39). Furthermore, the CPHF equations involve derivatives of the basis functions, while the equation for κ only involves integrals of the same type as for calculating the energy itself.

The Lagrange technique may be generalized to other types of non-variational wave functions (MP and CC), and to higher order derivatives. It is found that the $2n + 1$ rule is recovered, i.e. if the wave function response is known to order n , the $(2n + 1)$ th-order property may be calculated for any type of wave function.

10.5 Coupled Perturbed Hartree–Fock

Although a calculation of the wave function response can be avoided for the first derivative, it is necessary for second (and higher) derivatives. Eq. (10.32) gives directly an

equation for determining the (first-order) response, which is structurally the same as eq. (10.39). For a Hartree–Fock wave function, an equation for the change in the MO coefficients may also be formulated from the HF equation, eq. (3.51).

$$\mathbf{F}^{(0)}\mathbf{C}^{(0)} = \mathbf{S}^{(0)}\mathbf{C}^{(0)}\boldsymbol{\epsilon}^{(0)} \quad (10.43)$$

The superscript (0) here denotes the unperturbed system. The orthonormality of the molecular orbitals (eq. (3.20)) can be expressed as in eq. (10.44).

$$\mathbf{C}^{\dagger(0)}\mathbf{S}^{(0)}\mathbf{C}^{(0)} = \mathbf{1} \quad (10.44)$$

Expanding each of the \mathbf{F} , \mathbf{C} , \mathbf{S} and $\boldsymbol{\epsilon}$ matrices in terms of a perturbation parameter (e.g. $\mathbf{F} = \mathbf{F}^{(0)} + \lambda\mathbf{F}^{(1)} + \lambda^2\mathbf{F}^{(2)} + \dots$) and collecting all the first-order terms (analogous to the strategy used in Section 4.8) gives eq. (10.45).

$$\begin{aligned} \mathbf{F}^{(1)}\mathbf{C}^{(0)} + \mathbf{F}^{(0)}\mathbf{C}^{(1)} &= \mathbf{S}^{(1)}\mathbf{C}^{(0)}\boldsymbol{\epsilon}^{(0)} + \mathbf{S}^{(0)}\mathbf{C}^{(1)}\boldsymbol{\epsilon}^{(0)} + \mathbf{S}^{(0)}\mathbf{C}^{(0)}\boldsymbol{\epsilon}^{(1)} \\ [\mathbf{F}^{(0)} - \mathbf{S}^{(0)}\boldsymbol{\epsilon}^{(0)}]\mathbf{C}^{(1)} &= [-\mathbf{F}^{(1)} + \mathbf{S}^{(1)}\boldsymbol{\epsilon}^{(0)} + \mathbf{S}^{(0)}\boldsymbol{\epsilon}^{(1)}]\mathbf{C}^{(0)} \end{aligned} \quad (10.45)$$

The orthonormality condition becomes eq. (10.46).

$$\mathbf{C}^{\dagger(1)}\mathbf{S}^{(0)}\mathbf{C}^{(0)} + \mathbf{C}^{\dagger(0)}\mathbf{S}^{(1)}\mathbf{C}^{(0)} + \mathbf{C}^{\dagger(0)}\mathbf{S}^{(0)}\mathbf{C}^{(1)} = 0 \quad (10.46)$$

Equation (10.45) is the first-order *Coupled Perturbed Hartree–Fock* (CPHF) equation.⁴ The perturbed MO coefficients are given in terms of unperturbed quantities and the first-order Fock, Lagrange ($\boldsymbol{\epsilon}$) and overlap matrices. The $\mathbf{F}^{(1)}$ term is given in eq. (10.47).

$$\mathbf{F}^{(1)} = \mathbf{h}^{(1)} + \mathbf{G}^{(1)}\mathbf{D}^{(0)} + \mathbf{G}^{(0)}\mathbf{D}^{(1)} \quad (10.47)$$

Here \mathbf{h} is the one-electron (core) matrix, \mathbf{D} the density matrix and \mathbf{G} the tensor containing the two-electron integrals. The density matrix is given as a product of MO coefficients (eq. (3.52)).

$$\begin{aligned} \mathbf{D}^{(0)} &= \mathbf{C}^{\dagger(0)}\mathbf{C}^{(0)} \\ \mathbf{D}^{(1)} &= \mathbf{C}^{\dagger(1)}\mathbf{C}^{(0)} + \mathbf{C}^{\dagger(0)}\mathbf{C}^{(1)} \end{aligned} \quad (10.48)$$

The $\mathbf{S}^{(1)}$, $\mathbf{h}^{(1)}$ and $\mathbf{g}^{(1)}$ quantities are (first) derivatives of one- and two-electron integrals over basis functions.

$$\begin{aligned} S_{\alpha\beta}^{(1)} &= \langle \chi_\alpha | \chi_\beta \rangle^{(1)} = \frac{\partial}{\partial \lambda} \langle \chi_\alpha | \chi_\beta \rangle \\ h_{\alpha\beta}^{(1)} &= \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle^{(1)} = \frac{\partial}{\partial \lambda} \langle \chi_\alpha | \mathbf{h} | \chi_\beta \rangle \\ g_{\alpha\beta\gamma\delta}^{(1)} &= \langle \chi_\alpha \chi_\beta | \mathbf{g} | \chi_\gamma \chi_\delta \rangle^{(1)} = \frac{\partial}{\partial \lambda} \langle \chi_\alpha \chi_\beta | \mathbf{g} | \chi_\gamma \chi_\delta \rangle \end{aligned} \quad (10.49)$$

The derivatives of the integrals may involve derivatives of the basis functions or the operator, or both (see Section 10.8). Using eqs (10.48) and (10.47) in eq. (10.45) gives a set of linear equations relating $\mathbf{C}^{(1)}$ to $\mathbf{S}^{(1)}$, $\mathbf{h}^{(1)}$, $\mathbf{g}^{(1)}$ and $\mathbf{C}^{(0)}$.

Just as the variational condition for an HF wave function can be formulated either as a matrix equation or in terms of orbital rotations (Sections 3.5 and 3.6), the CPHF may also be viewed as a rotation of the molecular orbitals. In the absence of a perturbation, the molecular orbitals make the energy stationary, i.e. the derivative of the energy with respect to a change in the MOs is zero. This is equivalent to the statement that the off-diagonal elements of the Fock matrix between the occupied and virtual MOs are zero.

$$\begin{aligned} \langle \phi_i | \mathbf{F} | \phi_a \rangle &= 0 \\ \langle \phi_i | \mathbf{h} | \phi_a \rangle + \sum_{k=1}^{N_{\text{orb}}} [\langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle - \langle \phi_i \phi_k | \mathbf{g} | \phi_k \phi_a \rangle] &= 0 \end{aligned} \quad (10.50)$$

When a perturbation is introduced, the stationary condition means that the orbitals must change, which may be described as a mixing of the unperturbed MOs. In other words, the stationary orbitals in the presence of a perturbation are given by a unitary transformation of the unperturbed orbitals (see also Section 3.6).

$$\phi'_i = \sum_{j=1}^{M_{\text{basis}}} u_{ji} \phi_j \quad (10.51)$$

The \mathbf{U} matrix describes how the MOs change, i.e. it contains the derivatives of the MO coefficients. In the absence of a perturbation, \mathbf{U} is the identity matrix.

Let us now explicitly make $\mathbf{U}^{(1)}$ the matrix containing the first-order changes in the MO coefficients.

$$\phi_i \rightarrow \phi_i + \lambda \sum_{j=1}^{M_{\text{basis}}} u_{ji}^{(1)} \phi_j + \dots \quad (10.52)$$

In terms of the matrix formulation in eqs (10.45) and (10.46), the equivalent of eq. (10.51) is eq. (10.53).

$$\mathbf{C}^{(1)} = \mathbf{U}^{(1)} \mathbf{C}^{(0)} \quad (10.53)$$

An equation for the $\mathbf{U}^{(1)}$ elements can be obtained from the condition that the Fock matrix is diagonal, and by expanding all involved quantities to first order.

$$\begin{aligned} \langle \phi_i | \mathbf{h} | \phi_a \rangle &\rightarrow \langle \phi_i | \mathbf{h} | \phi_a \rangle^{(0)} + \langle \phi_i | \mathbf{h} | \phi_a \rangle^{(1)} \\ \langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle &\rightarrow \langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle^{(0)} + \langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle^{(1)} \end{aligned} \quad (10.54)$$

The $\langle \phi_i | \mathbf{h} | \phi_a \rangle^{(1)}$ and $\langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle^{(1)}$ elements are integral derivatives with respect to the perturbation, analogous to eq. (10.49), but expressed in terms of molecular orbitals. Inserting these expansions into the $\langle \phi_i | \mathbf{F} | \phi_a \rangle = 0$ condition and collecting all terms that are first order in λ gives a matrix equation that can be written as

$$\mathbf{A}^{(0)} \mathbf{U}^{(1)} = \mathbf{B}^{(1)} \quad (10.55)$$

The $\mathbf{A}^{(0)}$ matrix contains only unperturbed quantities ($\langle \phi_i | \mathbf{h} | \phi_a \rangle^{(0)}$ and $\langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle^{(0)}$), while the $\mathbf{B}^{(1)}$ matrix contains first derivatives ($\langle \phi_i | \mathbf{h} | \phi_a \rangle^{(1)}$ and $\langle \phi_i \phi_k | \mathbf{g} | \phi_a \phi_k \rangle^{(1)}$).

Since the energy is independent of a rotation among the occupied or virtual orbitals, only the mixing of occupied and virtual orbitals is determined by requiring that the energy be stationary. The occupied–occupied and virtual–virtual mixing may be fixed from the orthonormality condition (eq. (10.46)) or, equivalently, by requiring the perturbed Fock matrix to be diagonal also in the occupied–occupied and virtual–virtual blocks. Without these additional requirements, the procedure is called *Coupled Hartree–Fock* (CHF), as opposed to CPHF.

The CPHF equations are linear and can be solved by standard matrix operations. The size of the \mathbf{U} matrix is the number of occupied orbitals times the number of virtual orbitals, which in general is quite large, and the CPHF equations are therefore normally solved by iterative methods. Furthermore, as illustrated above, the CPHF equations may be formulated either in an atomic orbital or molecular orbital basis. Although the latter has computational advantages in certain cases, the former is more suitable for use in connection with direct methods (where the atomic integrals are calculated as required), as discussed in Section 3.8.5.

There is one CPHF equation to be solved for each perturbation. If it is an electric or magnetic field, there will in general be three components (F_x, F_y, F_z), if it is a geometry perturbation there will be $3N_{\text{atom}}$ (actually only $3N_{\text{atom}} - 6$ independent) components. Since the $\mathbf{A}^{(0)}$ matrix is independent of the nature of the perturbation, such multiple CPHF equations are often solved simultaneously.

The CPHF procedure may be generalized to higher order. Extending the expansion to second order allows the derivation of an equation for the second-order change in the MO coefficients, by solving a second-order CPHF equation, etc.

For perturbation-dependent basis sets (e.g. geometry derivatives) the (first-order) CPHF equations involve (first) derivatives of the one- and two-electron integrals with respect to the perturbation. For basis functions that are independent of the perturbation (e.g. an electric field), these derivatives are zero. Typically the solution of each CPHF equation (for each perturbation) requires approximately half of the time required for solving the HF equations themselves. For basis set dependent perturbations, the first-order CPHF equations are only needed for calculating second (and higher) derivatives, which have terms involving second (and higher) derivatives of the integrals themselves, and solving the CPHF equations is usually not the computational bottleneck in these cases.

Without the Lagrange technique for non-variational wave functions (CI, MP and CC), the n th-order CPHF is needed for the n th derivative. Consider for example the MP2 energy correction.

$$\text{MP2} = \sum_{i < j}^{N_{\text{occ}}} \sum_{a < b}^{N_{\text{vir}}} \frac{[\langle \phi_i \phi_j | \phi_a \phi_b \rangle - \langle \phi_i \phi_j | \phi_b \phi_a \rangle]^2}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (10.56)$$

The derivative of a molecular integral is given by eq. (10.57).

$$\frac{\partial}{\partial \lambda} \langle \phi_i \phi_j | \phi_a \phi_b \rangle = \frac{\partial}{\partial \lambda} \sum_{\alpha \beta \gamma \delta}^{M_{\text{basis}}} c_{i\alpha} c_{j\beta} c_{a\gamma} c_{b\delta} \langle \chi_\alpha \chi_\beta | \chi_\gamma \chi_\delta \rangle \quad (10.57)$$

This requires both the derivative of the MO coefficients and the two-electron integrals in the AO basis. The denominator leads to derivatives of the MO energies, which can

be obtained by solving the CPHF equations. A straightforward differentiation of eq. (10.56) thus leads to a formula where the first-order response is required.

Let us exemplify some of the above generalizations for the case of an HF wave function.

10.6 Electric Field Perturbation

10.6.1 External electric field

If the perturbation is a homogeneous electric field \mathbf{F} ($\mathbf{F} = F\mathbf{r}$), the perturbation operator \mathbf{P}_1 (eq. (10.21)) is the position vector \mathbf{r} and \mathbf{P}_2 is zero. Assuming that the basis functions are independent of the electric field (as is normally the case), the first-order HF property, the dipole moment, is given by the derivative formula (10.24) as shown in eq. (10.58) (since an HF wave function obeys the Hellmann–Feynman theorem).

$$\boldsymbol{\mu} = -\frac{\partial E_{\text{HF}}}{\partial \mathbf{F}} = -\langle \Psi_{\text{HF}} | \mathbf{r} | \Psi_{\text{HF}} \rangle \quad (10.58)$$

This is equivalent to the expression from first-order perturbation theory, (10.21). For non-variational wave functions the dipole moment calculated by the two approaches will be different, since the derivative of the wave function with respect to the field will not be zero. The second-order property, the dipole polarizability, is given by the derivative formula eq. (10.34) as shown in eq. (10.59).

$$\boldsymbol{\alpha} = -\frac{\partial^2 E_{\text{HF}}}{\partial \mathbf{F}^2} = 2 \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{F}} \left| \mathbf{r} \right| \Psi_{\text{HF}} \right\rangle \quad (10.59)$$

Second-order perturbation theory eq. (10.21) yields eq. (10.60).

$$\boldsymbol{\alpha} = -2 \sum_{i \neq 0} \frac{|\langle \Psi_{\text{HF}} | \mathbf{r} | \Psi_i \rangle|^2}{E_0 - E_i} \quad (10.60)$$

10.6.2 Internal electric field

Although nuclei are often modelled as point charges in quantum chemistry, they do in fact have a finite size. The internal structure of the nucleus leads to a quadrupole moment for nuclei with spin larger than $1/2$ (the dipole and octopole moments vanish by symmetry). This leads to an interaction term that is the product of the quadrupole moment with the field gradient ($\mathbf{F}' = \nabla \mathbf{F}$) created by the electron distribution.

$$\mathbf{H}_Q = - \sum_{A=1}^{N_{\text{nuclei}}} \mathbf{Q}_A \mathbf{F}' \quad (10.61)$$

10.7 Magnetic Field Perturbation

The situation is somewhat more complicated when the perturbation is a magnetic field. An electric field interacts directly with the charged particles (electrons and nuclei), and adds a *potential* energy term to the Hamiltonian operator. A magnetic field,

however, interacts with the magnetic moments generated by the *movement* of the charged particles (electrons), i.e. a magnetic perturbation changes the *kinetic* energy operator. The *generalized* (also called the *canonical*) momentum operator $\boldsymbol{\pi}$ is defined in eq. (10.62).

$$\boldsymbol{\pi} = \mathbf{p} - q\mathbf{A} \quad (10.62)$$

Here q is the charge and \mathbf{A} is the *vector potential* associated with the magnetic field \mathbf{B} (more correctly, the *magnetic induction* or *flux density*, being different from the magnetic field by a factor of $4\pi \cdot 10^{-7} \text{ Hm}^{-1}$), with the latter being given as the curl of the vector potential.

$$\mathbf{B} = \nabla \times \mathbf{A} \quad (10.63)$$

Only the kinetic energy of the electrons is considered within the Born–Oppenheimer approximation, and the generalized momentum becomes ($q = -1$) eq. (10.64).

$$\boldsymbol{\pi} = \mathbf{p} + \mathbf{A} \quad (10.64)$$

The vector potential is not uniquely defined since the gradient of any scalar function may be added (the curl of a gradient is always zero). For an external magnetic field, it is conventional to write it as in eq. (10.65).

$$\mathbf{A}_{\text{ext}}(\mathbf{r}) = \frac{1}{2} \mathbf{B}_{\text{ext}} \times (\mathbf{r} - \mathbf{R}_G) \quad (10.65)$$

Here \mathbf{R}_G is referred to as the *gauge origin*, i.e. the centre of the vector potential. One may verify by explicit calculation that the curl of \mathbf{A}_{ext} in eq. (10.64) indeed gives \mathbf{B}_{ext} .

A nucleus with a non-zero spin acts as a magnetic dipole, giving rise to a vector potential \mathbf{A}_A and producing the associated magnetic field by taking the curl.

$$\begin{aligned} \mathbf{A}_A &= \frac{g_A \mu_N}{c^2} \frac{\mathbf{I}_A \times (\mathbf{r} - \mathbf{R}_A)}{|\mathbf{r} - \mathbf{R}_A|^3} \\ \mathbf{B}_A &= -\frac{g_A \mu_N}{c^2} \left[\frac{\mathbf{I}_A}{|\mathbf{r} - \mathbf{R}_A|^3} - 3 \frac{(\mathbf{r} - \mathbf{R}_A)((\mathbf{r} - \mathbf{R}_A) \cdot \mathbf{I}_A)}{|\mathbf{r} - \mathbf{R}_A|^5} - \frac{8\pi}{3} \mathbf{I}_A \delta(\mathbf{r} - \mathbf{R}_A) \right] \end{aligned} \quad (10.66)$$

Here $g_A \mu_N \mathbf{I}_A$ is the magnetic moment of nucleus A and \mathbf{R}_A is the position (the nucleus is the natural gauge origin). The \mathbf{B}_A expression determines the magnetic field at position \mathbf{r} (not necessarily indicating an electron) due to a magnetic nucleus at position \mathbf{R}_A . The δ function in the last term in the \mathbf{B}_A expression arises from the quantum mechanical possibility of $\mathbf{r} - \mathbf{R}_A = 0$, i.e. the magnetic field directly at the nuclear position. Note that the presence of c^{-2} emerges from the units of magnetic field ($\mu_0/4\pi = c^{-2}$ in atomic units), and does not indicate a relativistic origin.

The spin associated with an electron also acts as a magnetic dipole ($-g_e \mu_B \mathbf{s}_i$), giving rise to a vector potential \mathbf{A}_e and associated magnetic field.

$$\begin{aligned} \mathbf{A}_e &= -\frac{g_e \mu_B}{c^2} \frac{\mathbf{s}_i \times (\mathbf{r} - \mathbf{r}_i)}{|\mathbf{r} - \mathbf{r}_i|^3} \\ \mathbf{B}_e &= -\frac{g_e \mu_B}{c^2} \left[\frac{\mathbf{s}_i}{|\mathbf{r} - \mathbf{r}_i|^3} - 3 \frac{(\mathbf{r} - \mathbf{r}_i)((\mathbf{r} - \mathbf{r}_i) \cdot \mathbf{s}_i)}{|\mathbf{r} - \mathbf{r}_i|^5} - \frac{8\pi}{3} \mathbf{s}_i \delta(\mathbf{r} - \mathbf{r}_i) \right] \end{aligned} \quad (10.67)$$

The \mathbf{B}_e expression similarly determines the magnetic field at position \mathbf{r} due to an electron at position \mathbf{r}_i .

The introduction of the generalized momentum operator in the one-electron kinetic energy part of the Dirac equation leads to three new interaction terms, as shown in eq. (8.29) and (8.30). It should be noted that the last two terms will also show up in a non-relativistic treatment when the magnetic vector potential is included, and only the $\mathbf{s} \cdot \mathbf{B}$ term should be considered a relativistic effect.

$$\begin{aligned} \mathbf{p}^2 &\rightarrow \boldsymbol{\pi}^2 = (\mathbf{p} + \mathbf{A})^2 \\ &\Downarrow \\ g_e \mu_B \mathbf{s} \cdot \mathbf{B}; \quad \mathbf{A} \cdot \mathbf{p}; \quad \frac{1}{2} \mathbf{A}^2 \end{aligned} \quad (10.68)$$

If there is more than one type of magnetic field present there will be an additional mixed $\mathbf{A}\mathbf{A}'$ term. Depending on the type of magnetic interactions present, these give operators as shown below. To simplify the expressions, we will use the notations $\mathbf{r}_{iA} = \mathbf{r}_i - \mathbf{R}_A$, $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ and $\mathbf{R}_{AB} = \mathbf{R}_A - \mathbf{R}_B$.

In addition to the magnetic terms arising from the expansion of the generalized momentum operator, there are also magnetic perturbation terms arising from relativistic corrections, as discussed in Section 8.2. These corrections may be derived by an expansion in the inverse speed of light. For consistency, we will in the following only consider terms up to order c^{-2} , with the exception of the indirect nuclear spin–spin coupling, where the lowest non-vanishing term is of order c^{-4} . Furthermore, in the rest of this section the summation over the number of electrons and nuclei has been omitted for clarity.

10.7.1 External magnetic field

For an external magnetic field, the three terms in eq. (10.68) becomes eqs (10.69)–(10.71).

$$g_e \mu_B \mathbf{s} \cdot \mathbf{B} = g_e \mu_B \mathbf{s} \cdot \mathbf{B}_{\text{ext}} \quad (10.69)$$

$$\begin{aligned} \mathbf{A} \cdot \mathbf{p} &= \left(\frac{1}{2} \mathbf{B}_{\text{ext}} \times \mathbf{r}_{iG} \right) \cdot \mathbf{p} \\ &= \frac{1}{2} \mathbf{B}_{\text{ext}} \cdot (\mathbf{r}_{iG} \times \mathbf{p}) \\ &= \frac{1}{2} \mathbf{B}_{\text{ext}} \cdot \mathbf{L}_G \end{aligned} \quad (10.70)$$

$$\begin{aligned} \frac{1}{2} \mathbf{A}^2 &= \frac{1}{2} \left(\frac{1}{2} \mathbf{B}_{\text{ext}} \times \mathbf{r}_{iG} \right) \cdot \left(\frac{1}{2} \mathbf{B}_{\text{ext}} \times \mathbf{r}_{iG} \right) \\ &= \frac{1}{8} \left(\mathbf{B}_{\text{ext}}^2 \cdot \mathbf{r}_{iG}^2 - (\mathbf{B}_{\text{ext}} \cdot \mathbf{r}_{iG})^2 \right) \\ &= \mathbf{B}_{\text{ext}} \cdot \mathbf{P}_2^{\xi} \cdot \mathbf{B}_{\text{ext}} \end{aligned} \quad (10.71)$$

Here the vector identities $\mathbf{a} \times \mathbf{b} \cdot \mathbf{c} = \mathbf{a} \cdot \mathbf{b} \times \mathbf{c}$ and $(\mathbf{a} \times \mathbf{b}) \cdot (\mathbf{c} \times \mathbf{d}) = (\mathbf{a} \cdot \mathbf{c})(\mathbf{b} \cdot \mathbf{d}) - (\mathbf{a} \cdot \mathbf{d})(\mathbf{b} \cdot \mathbf{c})$ have been used, and the angular momentum operator \mathbf{L}_G is defined implicitly by eq. (10.70). The presence of a magnetic field thus introduces three new terms, two being linear and one being quadratic in the field. The *spin-Zeeman* term $\mathbf{s} \cdot \mathbf{B}_{\text{ext}}$ describes the interaction of the electron spin with the magnetic field, while $\frac{1}{2} \mathbf{B}_{\text{ext}} \cdot \mathbf{L}_G$ is the *orbital-Zeeman* term describing the interaction of the magnetic field with

the magnetic moment associated with the movement of the electron. For a many-electron system, the spin-Zeeman term becomes $\mathbf{S} \cdot \mathbf{B}_{\text{ext}}$, where \mathbf{S} indicates the total molecular spin. The quadratic \mathbf{P}_2^ξ operator arising from $1/2\mathbf{A}_{\text{ext}}^2$ may be written as in eq. (10.72).

$$\mathbf{P}_2^\xi = \frac{1}{8}(\mathbf{r}_{iG}^t \mathbf{r}_{iG} - \mathbf{r}_{iG} \mathbf{r}_{iG}^t) \quad (10.72)$$

Here $\mathbf{r}_{iG}^t \mathbf{r}_{iG}$ is the inner (dot) product times a unit matrix (i.e. $\mathbf{r}_{iG} \cdot \mathbf{r}_{iG} \mathbf{I}$) and $\mathbf{r}_{iG} \mathbf{r}_{iG}^t$ is the outer product, i.e. a 3×3 matrix containing the products of the x, y, z components, analogous to the quadrupole moment, eq. (10.4). Note that both the \mathbf{L}_G and \mathbf{P}_2^ξ operators are gauge dependent.

10.7.2 Nuclear spin

The $g_e \mu_B \mathbf{S} \cdot \mathbf{B}$ term in eq. (10.68) in connection with \mathbf{B}_A in eq. (10.66) gives three terms, which conventionally are collected in two operators.

$$\begin{aligned} \mathbf{H}_{\text{ne}}^{\text{SD}} &= \mathbf{s}_i \cdot \mathbf{P}_{\text{ne}}^{\text{SD}} \cdot \mathbf{I}_A = -\frac{g_e g_A \mu_B \mu_N}{c^2} \mathbf{s}_i \cdot \left(\frac{\mathbf{r}_{iA}^t \mathbf{r}_{iA} - 3\mathbf{r}_{iA} \mathbf{r}_{iA}^t}{r_{iA}^5} \right) \cdot \mathbf{I}_A \\ \mathbf{H}_{\text{ne}}^{\text{FC}} &= \mathbf{s}_i \cdot \mathbf{P}_{\text{ne}}^{\text{FC}} \cdot \mathbf{I}_A = \frac{8\pi g_e g_A \mu_B \mu_N}{3c^2} \delta(\mathbf{r}_{iA}) (\mathbf{s}_i \cdot \mathbf{I}_A) \end{aligned} \quad (10.73)$$

$\mathbf{H}_{\text{ne}}^{\text{SD}}$ is a (one electron) *Spin-Dipolar* and $\mathbf{H}_{\text{ne}}^{\text{FC}}$ is a *Fermi Contact* operator, and their sum is the $\mathbf{H}_{\text{ne}}^{\text{SS}}$ operator in eq. (8.37).

The $\mathbf{A}_A \cdot \mathbf{p}$ term in eq. (10.68) gives the *Paramagnetic Spin-Orbit* operator.

$$\mathbf{H}_{\text{ne}}^{\text{PSO}} \mathbf{I}_A \cdot \mathbf{P}_{\text{ne}}^{\text{PSO}} = \frac{g_A \mu_N}{c^2} \cdot \mathbf{I}_A \cdot \frac{\mathbf{r}_{iA} \times \mathbf{p}_i}{r_{iA}^3} \quad (10.74)$$

$\mathbf{H}_{\text{ne}}^{\text{PSO}}$ is identical to eq. (8.37).

The $1/2\mathbf{A}_A^2$ term in eq. (10.68) gives a *Diamagnetic Spin-Orbit* operator, which is an operator of order c^{-4} . Although we otherwise only consider terms up to order c^{-2} , the nuclear spin-spin coupling constant only contains terms of order c^{-4} , which is why we need to include $\mathbf{H}_{\text{nn}}^{\text{DSO}}$.

$$\mathbf{H}_{\text{nn}}^{\text{DSO}} = \mathbf{I}_A \cdot \mathbf{P}_{\text{nn}}^{\text{DSO}} \cdot \mathbf{I}_B = \frac{g_A g_B \mu_N^2}{2c^4} \mathbf{I}_A \cdot \frac{(\mathbf{r}_{iA}^t \mathbf{r}_{iB} - \mathbf{r}_{iB} \mathbf{r}_{iA}^t)}{r_{iA}^3 r_{iB}^3} \cdot \mathbf{I}_B \quad (10.75)$$

When both nuclear spins and an external magnetic field are present, there is an additional mixed $\mathbf{A}_{\text{ext}} \cdot \mathbf{A}_A$ term arising from the expansion of the generalized momentum operator.

$$\mathbf{H}_{\text{ne}}^{\text{DS}} = \mathbf{B}_{\text{ext}} \cdot \mathbf{P}_{\text{ne}}^{\text{DS}} \cdot \mathbf{I}_A = \frac{g_A \mu_N}{2c^2} \mathbf{B}_{\text{ext}} \cdot \frac{\mathbf{r}_{iG}^t \mathbf{r}_{iA} - \mathbf{r}_{iA} \mathbf{r}_{iG}^t}{r_{iA}^3} \cdot \mathbf{I}_A \quad (10.76)$$

This nuclear *Diamagnetic Shielding* operator contributes to the NMR shielding tensor.

10.7.3 Electron spin

The $g_e\mu_B\mathbf{s} \cdot \mathbf{B}$ term in eq. (10.68) in connection with \mathbf{B}_e in eq. (10.67) gives three terms, which again are collected in two operators.

$$\begin{aligned}\mathbf{H}_{ee}^{\text{SD}} &= \mathbf{s}_i \cdot \mathbf{P}_{ee}^{\text{SD}} \cdot \mathbf{s}_j = \frac{g_e^2\mu_B^2}{2c^2} \mathbf{s}_i \cdot \left(\frac{\mathbf{r}_{ij}^i \mathbf{r}_{ij} - 3\mathbf{r}_{ij} \mathbf{r}_{ij}^t}{r_{ij}^5} \right) \cdot \mathbf{s}_j \\ \mathbf{H}_{ee}^{\text{FC}} &= \mathbf{s}_i \cdot \mathbf{P}_{ee}^{\text{FC}} \cdot \mathbf{s}_j = -\frac{4\pi g_e^2\mu_B^2}{3c^2} \delta(\mathbf{r}_{ij})(\mathbf{s}_i \cdot \mathbf{s}_j)\end{aligned}\quad (10.77)$$

$\mathbf{P}_{ee}^{\text{SD}}$ is a (two-electron) *Spin-Dipolar* and $\mathbf{P}_{ee}^{\text{FC}}$ is a *Fermi Contact* operator, with the sum being equal to the $\mathbf{H}_{ee}^{\text{SS}}$ operator in eq. (8.36).

The $\mathbf{A}_e \cdot \mathbf{p}$ term in eq. (10.67) gives the two-electron part of the spin-orbit operator.

$$\mathbf{H}_{ee}^{\text{SO}} = \mathbf{s}_i \cdot \mathbf{P}_{ee}^{\text{SO}} = -\frac{g_e\mu_B}{2c^2} \mathbf{s}_i \cdot \frac{\mathbf{r}_{ij} \times \mathbf{p}_i + 2\mathbf{r}_{ij} \times \mathbf{p}_j}{r_{ij}^3} \quad (10.78)$$

$\mathbf{H}_{ee}^{\text{SO}}$ is equivalent to the sum of $\mathbf{H}_{ee}^{\text{SO}}$ and $\mathbf{H}_{ee}^{\text{SOO}}$ in eq. (8.36).

The $1/2\mathbf{A}_e^2$ term in eq. (10.68) gives an operator analogous to eq. (10.75), but that depends on two electron spins instead of two nuclear spins. This, however, is an order c^{-4} operator compared with the order c^{-2} operators $\mathbf{P}_{ee}^{\text{SD}}$ and $\mathbf{P}_{ee}^{\text{FC}}$ (eq. (10.77)) describing spin-spin interactions, and is therefore neglected.

The $\mathbf{A}_e \cdot \mathbf{A}_A$ term in eq. (10.68) gives a coupling between the electronic and nuclear spins, and is again an operator of order c^{-4} . Compared with the order c^{-2} operators $\mathbf{P}_{ne}^{\text{SD}}$ and $\mathbf{P}_{ne}^{\text{FC}}$ (eq. (10.73)), it is again neglected.

When both electron spin and an external magnetic field are considered, there is a mixed $\mathbf{A}_{\text{ext}} \cdot \mathbf{A}_e$ term.

$$\mathbf{H}_{ee}^{\text{DS}} = \mathbf{B}_{\text{ext}} \cdot \mathbf{P}_{ee}^{\text{DS}} \cdot \mathbf{s}_i = -\frac{g_e\mu_B}{2c^2} \mathbf{B}_{\text{ext}} \cdot \frac{(\mathbf{r}_{iG}^t \mathbf{r}_{ij} - \mathbf{r}_{ij} \mathbf{r}_{iG}^t)}{r_{ij}^3} \cdot \mathbf{s}_i \quad (10.79)$$

This electronic *Diamagnetic Shielding* operator contributes to the ESR \mathbf{g} -tensor.

10.7.4 Classical terms

The expansion of the generalized momentum operator only involves the magnetic interactions in the electronic part of the wave function. Since the corresponding nuclear part has been separated out by the Born–Oppenheimer approximations, we need to add a few terms corresponding to the (classical) interaction of the nuclear magnetic moments with an external magnetic field and between nuclei.

The nuclear spin-Zeeman term is analogous to the electronic term in eq. (10.69), except that the electron magnetic moment of $-g_e\mu_B\mathbf{s}$ is replaced by the nuclear magnetic moment $g_A\mu_N\mathbf{I}$. The nuclear magneton μ_N is defined analogously to the Bohr magneton μ_B , but using the proton mass m_p instead of the electron mass ($\mu_N = e\hbar/2m_p = 2.723 \times 10^{-4}$ in atomic units), while the g_A factor depends on the specific nucleus (isotope).

$$\mathbf{H}_n^{\text{Zeeman}} = -\mu_N g_A \mathbf{I}_A \cdot \mathbf{B} \quad (10.80)$$

The term involving two nuclei is analogous to the electron–electron spin-dipole term in eq. (10.77). The corresponding Fermi contact term disappears since nuclei cannot occupy the same position at energies relevant for chemistry. Note that the direct spin–spin coupling is independent of the electronic wave function; it only depends on the molecular geometry.

$$\mathbf{H}_{nn}^{\text{SS}} = \frac{\mu_N^2 g_A g_B}{2c^2} \mathbf{I}_A \cdot \left(\frac{R_{AB}^t R_{AB} - 3R_{AB} R_{AB}^t}{R_{AB}^5} \right) \cdot \mathbf{I}_B \quad (10.81)$$

10.7.5 Relativistic terms

The most important relativistic corrections are the one-electron spin–orbit operator, and the relativistic correction to the spin-Zeeman operator.

$$\mathbf{H}_{ne}^{\text{SO}} = \frac{g_e \mu_B Z_A}{2mc^2} \mathbf{s}_i \cdot \frac{\mathbf{r}_{iA} \times \mathbf{p}_i}{r_{iA}^3} \quad (10.82)$$

$$\mathbf{H}_e^{\text{Zeeman-rel}} = \frac{g_e \mu_B}{2mc^2} (\mathbf{s}_i \cdot \mathbf{B}_i) \mathbf{p}_i^2 \quad (10.83)$$

Other relativistic corrections, such as the mass–velocity and Darwin terms, affect the wave function but do not lead to operators associated with molecular properties.

10.7.6 Magnetic properties

Table 10.2 shows the perturbation operators arising from inclusion of a magnetic field and relativistic effects. The last three columns indicate the perturbation order with respect to an external (\mathbf{B}_{ext}) and internal nuclear (\mathbf{I}) magnetic field, and with respect to the inverse speed of light. Only the most important operators up to order c^{-2} have been included, with the exception of the diamagnetic nuclear spin–spin coupling operator, since the leading term for this quantity is of order c^{-4} .

As seen from eq. (10.21), the first-order property is given as an expectation value of operators linear in the perturbation. The second-order property contains two contributions, an expectation value over quadratic (or bilinear) operators and a sum over products of matrix elements involving linear operators connecting the ground and excited states.

The first-order property with respect to an external field is the magnetic dipole moment \mathbf{m} (eq. (10.10)). When field-independent basis functions are used, the HF magnetic dipole moment is given as the expectation value of the $1/2 \mathbf{L}_G$ and \mathbf{S} (total electron spin) operators over the unperturbed wave function, eqs (10.21) and (10.24). Since the \mathbf{L}_G operator is imaginary it can only yield a non-zero result for spatially degenerate wave functions and the expectation value of \mathbf{S} is only non-zero for non-singlet states.

Table 10.2 Magnetic perturbation operators

Origin	Operator	Equation	Name	Perturbation order		
				\mathbf{B}_{ext}	\mathbf{I}	c^{-n}
$\mathbf{S} \cdot \mathbf{B}_{\text{ext}}$	\mathbf{S}	10.69	Electron spin-Zeeman	1	0	0
$\mathbf{A}_{\text{ext}} \cdot \mathbf{p}$	$\frac{1}{2} \mathbf{L}_G$	10.70	Orbital-Zeeman	1	0	0
$\frac{1}{2} \mathbf{A}_{\text{ext}}^2$	\mathbf{P}_2^ζ	10.71	Diamagnetic magnetizability	2	0	0
$\mathbf{s} \cdot \mathbf{B}_A$	$\mathbf{P}_{ne}^{\text{SD}}$	10.73	Nuclear–electron spin-dipole	0	1	2
	$\mathbf{P}_{ne}^{\text{FC}}$	10.73	Nuclear–electron Fermi contact	0	1	2
$\mathbf{A}_A \cdot \mathbf{p}$	$\mathbf{P}_{ne}^{\text{PSO}}$	10.74	Paramagnetic spin–orbit	0	1	2
$\frac{1}{2} \mathbf{A}_A^2$	$\mathbf{P}_{nn}^{\text{DSO}}$	10.75	Diamagnetic nuclear spin–spin	0	2	4
$\mathbf{A}_A \cdot \mathbf{A}_{\text{ext}}$	$\mathbf{P}_{ne}^{\text{DS}}$	10.76	Diamagnetic NMR shielding	1	1	2
$\mathbf{s} \cdot \mathbf{B}_e$	$\mathbf{P}_{ee}^{\text{SD}}$	10.77	Electron–electron spin-dipole	0	0	2
	$\mathbf{P}_{ee}^{\text{FC}}$	10.77	Electron–electron Fermi contact	0	0	2
$\mathbf{A}_e \cdot \mathbf{p}$	$\mathbf{P}_{ee}^{\text{SO}}$	10.78	Two-electron spin–orbit	0	0	2
$\mathbf{A}_e \cdot \mathbf{A}_{\text{ext}}$	$\mathbf{P}_{ee}^{\text{DS}}$	10.79	Diamagnetic ESR shielding	1	0	2
Classic	\mathbf{I}_A	10.80	Nuclear spin-Zeeman	1	1	0
Classic	$\mathbf{P}_{nn}^{\text{SS}}$	10.81	Nuclear spin–spin coupling	0	2	0
Relativistic	$\mathbf{P}_{ne}^{\text{SO}}$	10.82	One-electron spin–orbit	0	0	2
Relativistic	$\mathbf{p}^2 \mathbf{s}$	10.83	Electron spin-Zeeman, rel. corr.	1	0	2

Note that the spin-Zeeman term involves the total electron spin \mathbf{S} .

$$\mathbf{m} = -\frac{\partial E_{\text{HF}}}{\partial \mathbf{B}_{\text{ext}}} = -\langle \Psi_{\text{HF}} | \frac{1}{2} \mathbf{L}_G + g_e \beta_B \mathbf{S} | \Psi_{\text{HF}} \rangle \quad (10.84)$$

The second-order term, the magnetizability ξ , has two components. The derivative expression (10.34) is given by eq. (10.85).

$$\xi = -\frac{\partial^2 E_{\text{HF}}}{\partial \mathbf{B}_{\text{ext}}^2} = -2 \langle \Psi_{\text{HF}} | \mathbf{P}_2^\xi | \Psi_{\text{HF}} \rangle + 2 \left\langle \frac{\partial \Psi_{\text{HF}}}{\partial \mathbf{B}_{\text{ext}}} \middle| \frac{1}{2} \mathbf{L}_G | \Psi_{\text{HF}} \right\rangle \quad (10.85)$$

Second-order perturbation theory (eq. (10.21)) yields eq. (10.86).

$$\xi = -2 \langle \Psi_{\text{HF}} | \mathbf{P}_2^\xi | \Psi_{\text{HF}} \rangle - 2 \sum_{i \neq 0} \frac{|\langle \Psi_{\text{HF}} | \frac{1}{2} \mathbf{L}_G | \Psi_i \rangle|^2}{E_{\text{HF}} - E_i} \quad (10.86)$$

The first term is referred to as the *diamagnetic* contribution, while the latter is the *paramagnetic* part of the magnetizability. The total spin operator \mathbf{S} gives no contribution to the paramagnetic term, since the ground and excited states are orthogonal in the spatial part. Each of the two components depends on the selected gauge origin. However, for exact wave functions the gauge dependencies cancel exactly. For approximate wave functions this is not guaranteed, and as a result the total property may depend on where the origin for the vector potential (eq. (10.65)) has been chosen.

The first-order term with respect to a nuclear magnetic moment \mathbf{I} is the *hyperfine coupling* tensor \mathbf{A} , giving the coupling between the nuclear and electronic magnetic moments.⁵ The leading order term (c^{-2}) can be evaluated as a simple expectation

value of the $\mathbf{P}_{\text{ne}}^{\text{FC}}$ and $\mathbf{P}_{\text{ne}}^{\text{DS}}$ operators. The former provides the isotropic part of the tensor, while the latter gives the anisotropic part.

$$\mathbf{A} = \langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ne}}^{\text{FC}} + \mathbf{P}_{\text{ne}}^{\text{SD}} | \Psi_{\text{HF}} \rangle \quad (10.87)$$

The second-order property related to two nuclear spins \mathbf{I}_A and \mathbf{I}_B is the nuclear spin–spin coupling tensor. The *direct* interaction is determined entirely by the molecular geometry and given by eq. (10.81). For rapidly tumbling molecules (solution or gas phase) this contribution averages out to zero, but it is significant for solid-state NMR.

The *indirect* spin–spin coupling between nuclei A and B, which is the one observed in solution phase NMR, contains several contributions, all being of order c^{-4} .

$$\mathbf{J}_{\text{AB}} = \langle \Psi_{\text{HF}} | \mathbf{P}_{\text{nn}}^{\text{DSO}} | \Psi_{\text{HF}} \rangle + \sum_{i \neq 0} \frac{\langle \Psi_{\text{HF}} | \mathbf{P}_{1,A} | \Psi_i \rangle \langle \Psi_i | \mathbf{P}_{1,B} | \Psi_{\text{HF}} \rangle}{E_{\text{HF}} - E_i} \quad (10.88)$$

The first part can be evaluated as the expectation value of $\mathbf{P}_{\text{nn}}^{\text{DSO}}$ (eq. (10.75)). The second part corresponds to all combinations of operators that are linear in the nuclear spin, i.e. $\mathbf{P}_{\text{ne}}^{\text{FC}}$, $\mathbf{P}_{\text{ne}}^{\text{SD}}$ and $\mathbf{P}_{\text{ne}}^{\text{PSO}}$ (eqs (10.73) and (10.74)). $\mathbf{P}_{\text{ne}}^{\text{FC}}$ and $\mathbf{P}_{\text{ne}}^{\text{SD}}$ contain the electron spin operators and for a singlet ground state (as is usually the case), this means that the excited state Ψ_i in the summation must be a triplet state. Since $\mathbf{P}_{\text{ne}}^{\text{PSO}}$ does not depend on electron spin, the combination of $\mathbf{P}_{\text{ne}}^{\text{PSO}}$ with either $\mathbf{P}_{\text{ne}}^{\text{FC}}$ or $\mathbf{P}_{\text{ne}}^{\text{SD}}$ gives zero contribution. For rapidly tumbling molecules, it can be shown the cross term between $\mathbf{P}_{\text{ne}}^{\text{FC}}$ and $\mathbf{P}_{\text{ne}}^{\text{SD}}$ averages out. For the trace (sum of the diagonal terms) of the 3×3 coupling matrix \mathbf{J} , which is the observed coupling constant, only the three “diagonal” terms ($\mathbf{P}_1 = \mathbf{P}_1$) in eq. (10.88) thus survive. The Fermi contact term is the most important for one-bond couplings ($^1\mathbf{J}$) in singly bonded systems, but the other three contributions become important for multiple-bonded systems and for longer-range couplings ($^2\mathbf{J}$, $^3\mathbf{J}$).⁶

The first-order term with respect to total electron spin \mathbf{S} is the spin–orbit splitting of open-shell molecules with a net angular momentum, as for example NO with a $^2\Pi$ ground state.

$$\text{SO-splitting} = \langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ne}}^{\text{SO}} + \mathbf{P}_{\text{ee}}^{\text{SO}} | \Psi_{\text{HF}} \rangle \quad (10.89)$$

The second-order property related to two electron spins is the *Zero Field Splitting* (ZFS) tensor \mathbf{D} , which is responsible for making the three individual components of a triplet state non-degenerate, with a typical magnitude being of the order of a few cm^{-1} .⁷ It again contains second- and first-order contributions, the latter arising from the spin–orbit operator containing both the one- and two-electron contributions.

$$\mathbf{D} = \langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ee}}^{\text{SD}} + \mathbf{P}_{\text{ee}}^{\text{FC}} | \Psi_{\text{HF}} \rangle + \sum_{i \neq 0} \frac{|\langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ne}}^{\text{SO}} + \mathbf{P}_{\text{ee}}^{\text{SO}} | \Psi_i \rangle|^2}{E_{\text{HF}} - E_i} \quad (10.90)$$

The $\mathbf{P}_{\text{ee}}^{\text{FC}}$ operator just introduces a uniform shift of all energy levels, and thus produces no observable effect on the splitting of the energy levels.

The interaction of an external magnetic field with a nuclear spin gives a Zeeman splitting of the energy levels by the $\mathbf{I}_A \cdot \mathbf{B}_{\text{ext}}$ term (Table 10.2). The details of the splitting, however, depend on the molecular environment, since the local magnetic field at a nuclear position is shielded by the electrons relative to the external field, $\mathbf{B}_{\text{local}} =$

$(\mathbf{1} - \boldsymbol{\sigma})\mathbf{B}_{\text{ext}}$. The *NMR shielding tensor* $\boldsymbol{\sigma}$, which is the mixed second derivative with respect to a nuclear spin and an external magnetic field, has, by analogy with the magnetizability, a diamagnetic and paramagnetic part.^{6,8} The diamagnetic part arises from $\mathbf{P}_{\text{ne}}^{\text{DS}}$, while the paramagnetic contribution contains products of matrix elements involving operators linear in \mathbf{B} or \mathbf{I} . These are given by the angular momentum operator $1/2\mathbf{L}_G$, eq. (10.70) and by the paramagnetic spin-orbit operator $\mathbf{P}_{\text{ne}}^{\text{PSO}}$, eq. (10.74). Written in terms of the perturbation formula (10.21), the expression for the nuclear shielding for atom A becomes eq. (10.91).

$$\boldsymbol{\sigma}_A = \langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ne}}^{\text{DS}} | \Psi_{\text{HF}} \rangle - \sum_{i \neq 0} \frac{\langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ne}}^{\text{PSO}} | \Psi_i \rangle \langle \Psi_i | 1/2 \mathbf{L}_G | \Psi_{\text{HF}} \rangle + \langle \Psi_{\text{HF}} | 1/2 \mathbf{L}_G | \Psi_i \rangle \langle \Psi_i | \mathbf{P}_{\text{ne}}^{\text{PSO}} | \Psi_{\text{HF}} \rangle}{E_{\text{HF}} - E_i} \quad (10.91)$$

All the operators $\mathbf{P}_{\text{ne}}^{\text{DS}}$, $\mathbf{P}_{\text{ne}}^{\text{PSO}}$ and \mathbf{L}_G are gauge dependent, and each of the dia- and paramagnetic terms consequently depends on the chosen gauge. The shielding tensor is a 3×3 matrix, which can be diagonalized to give three eigenvalues. These principal components can be observed by solid-state NMR, but for rapidly tumbling molecules, as in solution phase, only the average can be observed, corresponding to one-third of the trace of the shielding tensor.

The ESR equivalent of the NMR shielding is called the *g-tensor*, and can be considered as the mixed second derivative with respect to the electron spin and an external magnetic field.⁹ It is a 3×3 tensor, and can be written as the diagonal component for the free electron plus a small correction due to the molecular environment.

$$\mathbf{g} = g_e \mathbf{1} + \Delta \mathbf{g} \quad (10.92)$$

The diagonal term $g_e \mathbf{1}$ arises from the spin-Zeeman term $\mathbf{S} \cdot \mathbf{B}_{\text{ext}}$, while the anisotropic part has two contributions. The direct term arises from the $\mathbf{P}_{\text{ee}}^{\text{DS}}$ operator (eq. (10.79)) and the relativistic correction from the spin-Zeeman term (eq. (10.83)), with additional contributions coming from the combination of $1/2\mathbf{L}_G$ and the spin-orbit operators $\mathbf{P}_{\text{ne}}^{\text{SO}}$ and $\mathbf{P}_{\text{ee}}^{\text{SO}}$.

$$\Delta \mathbf{g} = \langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ee}}^{\text{DS}} | \Psi_{\text{HF}} \rangle - \frac{g_e \mu_B}{2c^2} \langle \Psi_{\text{HF}} | \mathbf{P}^2 | \Psi_{\text{HF}} \rangle - \sum_{i \neq 0} \frac{\langle \Psi_{\text{HF}} | \mathbf{P}_{\text{ne}}^{\text{SO}} + \mathbf{P}_{\text{ee}}^{\text{SO}} | \Psi_i \rangle \langle \Psi_i | 1/2 \mathbf{L}_G | \Psi_{\text{HF}} \rangle + \langle \Psi_{\text{HF}} | 1/2 \mathbf{L}_G | \Psi_i \rangle \langle \Psi_i | \mathbf{P}_{\text{ne}}^{\text{SO}} + \mathbf{P}_{\text{ee}}^{\text{SO}} | \Psi_{\text{HF}} \rangle}{E_{\text{HF}} - E_i} \quad (10.93)$$

It should be noted that the formulas in eqs (10.84)–(10.93) have been derived by considering only the operators from the leading order in the inverse speed of light. One may obtain relativistic corrections by carrying out the expansion to higher orders, but this rapidly becomes quite involved,^{8,9} as many different operators and their combinations can make contributions to a given property. For systems where relativistic effects are important, a full four-component type calculation (Section 8.4) becomes attractive, at least conceptually, since it automatically includes all effects without the necessity of multiple perturbation operators.

10.7.7 Gauge dependence of magnetic properties

There are two factors that make the calculation of magnetic properties somewhat more complicated than the corresponding electric properties. First, the angular momentum operator \mathbf{L}_G is imaginary (eq. (10.10)), implying that the wave function must be allowed to be complex. Second, the presence of the gauge origin in the operators means that the results may be origin dependent. An exact wave function will of course give origin-independent results, as will a Hartree–Fock wave function if a complete basis set is employed. In practice, however, a finite basis must be employed, and standard basis sets will yield results that depend on where the user has chosen the origin of the gauge. The centre of mass is often used in actual calculations, but this is by no means a unique choice. The gauge error depends on the distance between the wave function and the gauge origin, and some methods try to minimize the error by selecting separate gauges for each (localized) molecular orbital. Two such methods are known as *Individual Gauge for Localized Orbitals* (IGLO)¹⁰ and *Localized Orbital/local oRiGin* (LORG).¹¹

A more recent implementation, which eliminates the gauge dependence for properties, is to make the basis functions explicitly dependent on the magnetic field by inclusion of a complex phase factor referring to the position of the basis function (usually the nucleus).

$$\begin{aligned} X_A(\mathbf{r} - \mathbf{R}_A) &= e^{-\frac{i}{c} \mathbf{A}_A \cdot \mathbf{r}} \chi_A(\mathbf{r} - \mathbf{R}_A) \\ \chi_A(\mathbf{r} - \mathbf{R}_A) &= N e^{-\alpha(\mathbf{r} - \mathbf{R}_A)^2} \\ \mathbf{A}_A &= \frac{1}{2} \mathbf{B} \times (\mathbf{R}_A - \mathbf{R}_G) \end{aligned} \quad (10.94)$$

Such orbitals are known as *London Atomic Orbitals* (LAO) or *Gauge Including/Invariant Atomic Orbitals* (GIAO).¹² The effect is that matrix elements involving GIAOs only contain a *difference* in vector potentials, thereby removing the reference to an absolute gauge origin. For the overlap and potential energy, it is straightforward to see that matrix elements become independent of the gauge origin.

$$\begin{aligned} \langle X_A | X_B \rangle &= \langle \chi_A | e^{\frac{i}{c}(\mathbf{A}_A - \mathbf{A}_B) \cdot \mathbf{r}} | \chi_B \rangle \\ \langle X_A | \mathbf{V} | X_B \rangle &= \langle \chi_A | e^{\frac{i}{c}(\mathbf{A}_A - \mathbf{A}_B) \cdot \mathbf{r}} \mathbf{V} | \chi_B \rangle \\ \mathbf{A}_A - \mathbf{A}_B &= \frac{1}{2} \mathbf{B} \times (\mathbf{R}_A - \mathbf{R}_B) \end{aligned} \quad (10.95)$$

The kinetic energy is slightly more complicated, but it can be shown that the relation shown in eq. (10.96) holds.

$$\begin{aligned} \langle X_A | \pi^2 | X_B \rangle &= \langle X_A | (\mathbf{p} + \frac{1}{2} \mathbf{B} \times (\mathbf{r} - \mathbf{R}_G))^2 | X_B \rangle \\ &= \langle \chi_A | e^{\frac{i}{c}(\mathbf{A}_A - \mathbf{A}_B) \cdot \mathbf{r}} (\mathbf{p} + \frac{1}{2} \mathbf{B} \times (\mathbf{r} - \mathbf{R}_B))^2 | \chi_B \rangle \end{aligned} \quad (10.96)$$

Note that \mathbf{R}_G has been replaced by \mathbf{R}_B in the last bracket. The use of GIAOs as basis functions makes all matrix elements, and therefore all properties, independent of the gauge origin. The wave function itself, however, is expressed in term of the basis functions, and therefore becomes gauge dependent, by means of a complex phase factor. The use of perturbation-dependent basis functions has the further advantage of greatly

reducing the need for high angular momentum basis functions, i.e. the property is typically calculated with an accuracy comparable to that of the unperturbed system.¹³

While LAOs/GIAOs were proposed well before the advent of modern computational chemistry, it was only owing to developments in calculating (geometrical) derivatives of the energy (and wave function) that it became practical to use field-dependent orbitals.¹⁴

10.8 Geometry Perturbations

The general formula for the first derivative of the energy with respect to a change in geometry, the molecular (nuclear) gradient, is given by eq. (10.24).

$$\mathbf{g} = \frac{\partial E}{\partial \mathbf{R}} = \langle \Psi | \frac{\partial \mathbf{H}}{\partial \mathbf{R}} | \Psi \rangle + 2 \left\langle \frac{\partial \Psi}{\partial \mathbf{R}} | \mathbf{H} | \Psi \right\rangle \quad (10.97)$$

The first term is the Hellmann–Feynman force and the second is the wave function response. The latter contains contributions from a change in the basis functions, the state and the MO coefficients.

$$\frac{\partial \Psi}{\partial \mathbf{R}} = \frac{\partial \Psi}{\partial \boldsymbol{\chi}} \frac{\partial \boldsymbol{\chi}}{\partial \mathbf{R}} + \frac{\partial \Psi}{\partial \mathbf{a}} \frac{\partial \mathbf{a}}{\partial \mathbf{R}} + \frac{\partial \Psi}{\partial \mathbf{c}} \frac{\partial \mathbf{c}}{\partial \mathbf{R}} \quad (10.98)$$

The state and MO dependence disappears for HF, DFT and MCSCF type wave functions owing to the variational nature ($\partial \Psi / \partial \mathbf{a} = 0$ and $\partial \Psi / \partial \mathbf{c} = 0$). For traditional basis sets consisting of nuclear-centred Gaussian functions, the basis functions are clearly perturbation dependent since the functions move along with the nuclei, and standard perturbation theory is therefore not suitable for calculating molecular gradients. For a plane wave basis, however, the basis functions are independent of a geometry perturbation, and the molecular gradient is just the Hellmann–Feynman term.

Since geometry derivatives are important for optimizing geometries, it may be useful to look in more detail at the quantities involved in calculating first and second derivatives of a Hartree–Fock wave function with a Gaussian type basis set, with the expressions for density functional methods being very similar. These formulas are most easily derived directly from the HF energy expressed in terms of the atomic quantities (eq. (3.54)).¹⁵

$$E_{\text{HF}} = \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} h_{\alpha\beta} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} (\langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle - \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\delta} \chi_{\beta} \rangle) + V_{\text{nn}} \quad (10.99)$$

Differentiation (using λ as a general geometrical displacement of a nucleus) yields eq. (10.100).

$$\begin{aligned} \frac{\partial E_{\text{HF}}}{\partial \lambda} = & \sum_{\alpha\beta}^{M_{\text{basis}}} \left(\frac{\partial D_{\alpha\beta}}{\partial \lambda} h_{\alpha\beta} + D_{\alpha\beta} \frac{\partial h_{\alpha\beta}}{\partial \lambda} \right) + \\ & \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \left(\frac{\partial D_{\alpha\beta}}{\partial \lambda} D_{\gamma\delta} + D_{\alpha\beta} \frac{\partial D_{\gamma\delta}}{\partial \lambda} \right) (\langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle - \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\delta} \chi_{\beta} \rangle) + \\ & \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} \frac{\partial}{\partial \lambda} (\langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle - \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\delta} \chi_{\beta} \rangle) + \frac{\partial V_{\text{nn}}}{\partial \lambda} \end{aligned} \quad (10.100)$$

The third and fourth terms are identical and may be collected to cancel the factor of $1/2$. Rearranging the terms gives eq. (10.101).

$$\begin{aligned} \frac{\partial E_{\text{HF}}}{\partial \lambda} = & \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} \frac{\partial h_{\alpha\beta}}{\partial \lambda} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} \frac{\partial}{\partial \lambda} (\langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle - \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\delta} \chi_{\beta} \rangle) \\ & + \sum_{\alpha\beta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} h_{\alpha\beta} + \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} D_{\gamma\delta} (\langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle - \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\delta} \chi_{\beta} \rangle) \\ & + \frac{\partial V_{\text{nn}}}{\partial \lambda} \end{aligned} \quad (10.101)$$

The first two terms involve products of the density matrix with derivatives of the atomic integrals, while the two next terms can be recognized as derivatives of the density matrix times the Fock matrix (eq. (3.52)).

$$\begin{aligned} \frac{\partial E_{\text{HF}}}{\partial \lambda} = & \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} \frac{\partial h_{\alpha\beta}}{\partial \lambda} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} \frac{\partial}{\partial \lambda} (\langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle - \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\delta} \chi_{\beta} \rangle) \\ & + \frac{\partial V_{\text{nn}}}{\partial \lambda} + \sum_{\alpha\beta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} F_{\alpha\beta} \end{aligned} \quad (10.102)$$

The derivative in eq. (10.102) of the nuclear repulsion (third term) is trivial since it does not involve electron coordinates. The one-electron derivatives are given in eq. (10.103).

$$\begin{aligned} h_{\alpha\beta} &= \langle \chi_{\alpha} | \mathbf{h} | \chi_{\beta} \rangle \\ \frac{\partial h_{\alpha\beta}}{\partial \lambda} &= \left\langle \frac{\partial \chi_{\alpha}}{\partial \lambda} \middle| \mathbf{h} \middle| \chi_{\beta} \right\rangle + \langle \chi_{\alpha} | \frac{\partial \mathbf{h}}{\partial \lambda} | \chi_{\beta} \rangle + \left\langle \chi_{\alpha} | \mathbf{h} \middle| \frac{\partial \chi_{\beta}}{\partial \lambda} \right\rangle \end{aligned} \quad (10.103)$$

The central term is recognized as the Hellmann–Feynman force. The two-electron derivatives in eq. (10.102) become eq. (10.104).

$$\begin{aligned} \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle &= \langle \chi_{\alpha} \chi_{\gamma} | \mathbf{g} | \chi_{\beta} \chi_{\delta} \rangle \\ \frac{\partial}{\partial \lambda} \langle \chi_{\alpha} \chi_{\gamma} | \chi_{\beta} \chi_{\delta} \rangle &= \left\langle \frac{\partial \chi_{\alpha}}{\partial \lambda} \chi_{\gamma} \middle| \mathbf{g} \middle| \chi_{\beta} \chi_{\delta} \right\rangle + \left\langle \chi_{\alpha} \frac{\partial \chi_{\gamma}}{\partial \lambda} \middle| \mathbf{g} \middle| \chi_{\beta} \chi_{\delta} \right\rangle \\ &\quad + \langle \chi_{\alpha} \chi_{\gamma} | \frac{\partial \mathbf{g}}{\partial \lambda} | \chi_{\beta} \chi_{\delta} \rangle + \langle \chi_{\alpha} \chi_{\gamma} | \mathbf{g} | \frac{\partial \chi_{\beta}}{\partial \lambda} \chi_{\delta} \rangle \\ &\quad + \langle \chi_{\alpha} \chi_{\gamma} | \mathbf{g} | \chi_{\beta} \frac{\partial \chi_{\delta}}{\partial \lambda} \rangle \end{aligned} \quad (10.104)$$

The central term is again the Hellmann–Feynman force, which vanishes since the two-electron operator \mathbf{g} is independent of the nuclear positions.

The last term in eq. (10.102) involves a change in the density matrix, i.e. the MO coefficients.

$$D_{\alpha\beta} = \sum_{i=1}^{N_{\text{elec}}} n_i c_{\alpha i} c_{\beta i} \quad (10.105)$$

$$\frac{\partial D_{\alpha\beta}}{\partial \lambda} = \sum_{i=1}^{N_{\text{elec}}} n_i \left(\frac{\partial c_{\alpha i}}{\partial \lambda} c_{\beta i} + c_{\alpha i} \frac{\partial c_{\beta i}}{\partial \lambda} \right)$$

Since the HF wave function is variationally optimized, the explicit calculation of the density derivatives can be avoided, as first derived by Pulay.¹⁶ The last term in eq. (10.102) may with eq. (10.105) be written as in eq. (10.106).

$$\sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} F_{\alpha\beta} = \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \sum_{i=1}^{N_{\text{elec}}} n_i \left(\frac{\partial c_{\alpha i}}{\partial \lambda} F_{\alpha\beta} c_{\beta i} + \frac{\partial c_{\beta i}}{\partial \lambda} F_{\alpha\beta} c_{\alpha i} \right) \quad (10.106)$$

By virtue of the HF condition ($\mathbf{FC} = \mathbf{SC}\epsilon$), eq. (10.106) may be written in terms of overlap integrals and MO energies.

$$\sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} F_{\alpha\beta} = \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \sum_{i=1}^{N_{\text{elec}}} n_i \left(\frac{\partial c_{\alpha i}}{\partial \lambda} S_{\alpha\beta} \epsilon_i c_{\beta i} + \frac{\partial c_{\beta i}}{\partial \lambda} S_{\alpha\beta} \epsilon_i c_{\alpha i} \right) \quad (10.107)$$

Finally, since the MOs are orthonormal, the derivatives of the coefficients may be replaced by derivatives of the overlap matrix.

$$\langle \phi_i | \phi_j \rangle = \sum_{\alpha\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta j} \langle \chi_\alpha | \chi_\beta \rangle = \sum_{\alpha\beta}^{M_{\text{basis}}} c_{\alpha i} c_{\beta j} S_{\alpha\beta} = \delta_{ij}$$

$$\frac{\partial}{\partial \lambda} \langle \phi_i | \phi_j \rangle = \left(\frac{\partial c_{\alpha i}}{\partial \lambda} c_{\beta j} S_{\alpha\beta} + c_{\alpha i} \frac{\partial c_{\beta j}}{\partial \lambda} S_{\alpha\beta} + c_{\alpha i} c_{\beta j} \frac{\partial S_{\alpha\beta}}{\partial \lambda} \right) = 0 \quad (10.108)$$

$$2 \frac{\partial c_{\alpha i}}{\partial \lambda} c_{\beta j} S_{\alpha\beta} = -c_{\alpha i} c_{\beta j} \frac{\partial S_{\alpha\beta}}{\partial \lambda}$$

The final derivative of the energy may thus be written as in eq. (10.109).

$$\begin{aligned} \frac{\partial E_{\text{HF}}}{\partial \lambda} &= \sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} \frac{\partial h_{\alpha\beta}}{\partial \lambda} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} \frac{\partial}{\partial \lambda} (\langle \chi_\alpha \chi_\gamma | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \chi_\delta \chi_\beta \rangle) \\ &\quad + \frac{\partial V_{\text{nn}}}{\partial \lambda} - \sum_{\alpha\beta}^{M_{\text{basis}}} W_{\alpha\beta} \frac{\partial S_{\alpha\beta}}{\partial \lambda} \end{aligned} \quad (10.109)$$

Here the energy-weighted density matrix \mathbf{W} has been introduced.

$$W_{\alpha\beta} = \sum_{i=1}^{N_{\text{elec}}} \epsilon_i c_{\alpha i} c_{\beta i} \quad (10.110)$$

Consider now the case where the perturbation λ is a specific nuclear displacement, $X_k \rightarrow X_k + \Delta X_k$. The derivatives of the one- and two-electron integrals are of two types, those involving derivatives of the basis functions, and those involving derivatives of the operators. The latter are given in eq. (10.111).

$$\begin{aligned}
\frac{\partial \mathbf{h}}{\partial X_k} &= \frac{\partial}{\partial X_k} \left(-\frac{1}{2} \nabla_i^2 - \sum_A^{N_{\text{nuclei}}} \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} \right) = -Z_k \frac{(x_i - X_k)}{|\mathbf{r}_i - \mathbf{R}_k|^3} \\
\frac{\partial \mathbf{g}}{\partial X_k} &= 0 \\
\frac{\partial \mathbf{V}_{\text{nn}}}{\partial X_k} &= \frac{\partial}{\partial X_k} \left(\sum_{A>B}^{N_{\text{nuclei}}} \frac{Z_A Z_B}{|\mathbf{R}_A - \mathbf{R}_B|} \right) = \sum_{A \neq k}^{N_{\text{nuclei}}} Z_A Z_k \frac{(X_A - X_k)}{|\mathbf{R}_A - \mathbf{R}_k|^3}
\end{aligned} \tag{10.111}$$

The derivative of the core operator \mathbf{h} is a one-electron operator similar to the nuclear-electron attraction required for the energy itself (eq. (3.56)). The two-electron part yields zero, and the \mathbf{V}_{nn} term is independent of the electronic wave function. The remaining terms in eqs (10.103), (10.104) and (10.109) all involve derivatives of the basis functions. When these are Gaussian functions (as is usually the case) the derivative can be written in terms of two other Gaussian functions, having one lower and one higher angular momentum.

$$\begin{aligned}
\chi_\alpha(\mathbf{R}_k) &= N(x - X_k)^l (y - Y_k)^m (z - Z_k)^n e^{-\alpha(\mathbf{r} - \mathbf{R}_k)^2} \\
\frac{\partial \chi_\alpha}{\partial X_k} &= -N(x - X_k)^{l-1} (y - Y_k)^m (z - Z_k)^n e^{-\alpha(\mathbf{r} - \mathbf{R}_k)^2} \\
&\quad + 2N\alpha(x - X_k)^{l+1} (y - Y_k)^m (z - Z_k)^n e^{-\alpha(\mathbf{r} - \mathbf{R}_k)^2}
\end{aligned} \tag{10.112}$$

The derivative of a p -function can thus be written in terms of an s - and a d -type Gaussian function. The one- and two-electron integrals involving derivatives of basis functions are therefore of the same type as those used in the energy expression itself, the only difference is the angular momentum, and the fact that there are roughly three times as many of these derivative integrals than for the energy itself. Of all the terms in eq. (10.109), the only significant computational effort is the derivatives of the two-electron integrals. Note, however, that the density matrix elements are known at the time when these integrals are calculated, and screening procedures analogous to those used in direct SCF techniques (Section 3.8.5) can be used to avoid calculating integrals that make insignificant contributions to the final result.

The second derivative of the energy with respect to a geometry change can be written as in eq. (10.113).

$$\begin{aligned}
\mathbf{H} &= \frac{\partial^2 E_{\text{HF}}}{\partial \lambda^2} = \\
&\sum_{\alpha\beta}^{M_{\text{basis}}} D_{\alpha\beta} \frac{\partial^2 h_{\alpha\beta}}{\partial \lambda^2} + \frac{1}{2} \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} D_{\alpha\beta} D_{\gamma\delta} \frac{\partial^2}{\partial \lambda^2} (\langle \chi_\alpha \chi_\gamma | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \chi_\delta \chi_\beta \rangle) \\
&+ \frac{\partial^2 V_{\text{nn}}}{\partial \lambda^2} - \sum_{\alpha\beta}^{M_{\text{basis}}} W_{\alpha\beta} \frac{\partial^2 S_{\alpha\beta}}{\partial \lambda^2} + \sum_{\alpha\beta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} \frac{\partial h_{\alpha\beta}}{\partial \lambda} \\
&+ \sum_{\alpha\beta\gamma\delta}^{M_{\text{basis}}} \frac{\partial D_{\alpha\beta}}{\partial \lambda} D_{\gamma\delta} \frac{\partial}{\partial \lambda} (\langle \chi_\alpha \chi_\gamma | \chi_\beta \chi_\delta \rangle - \langle \chi_\alpha \chi_\gamma | \chi_\delta \chi_\beta \rangle) \\
&- \sum_{\alpha\beta}^{M_{\text{basis}}} \frac{\partial W_{\alpha\beta}}{\partial \lambda} \frac{\partial S_{\alpha\beta}}{\partial \lambda}
\end{aligned} \tag{10.113}$$

The first four terms only involve derivatives of operators and AO integrals. However, for the last three terms we need the derivative of the density matrix and MO energies. These can be obtained by solving the first-order CPHF equations (Section 10.5).

The calculation of the second derivatives is substantially more involved than calculating the first derivative, typically by an order of magnitude, and for large systems, a full calculation of the Hessian (force constant) matrix may thus be prohibitively expensive. If the second derivatives are required only for characterizing the nature of a stationary point (minimum or saddle point), the full Hessian is not required: only a few of the lowest eigenvalues are of interest. As shown by Deglmann and Furche, the lowest eigenvalues may be extracted by iterative techniques without explicit construction of the full second derivatives, leading to a substantial saving for large systems.¹⁷

10.9 Response and Propagator Methods

The perturbation and derivative approaches in sections 10.2 and 10.3 are not suitable for time-dependent properties since there is no well-defined energy function in such cases. The equivalent of eq. (10.20) for a time-dependent perturbation is eq. (10.114).

$$\mathbf{H}(t) = \mathbf{H}_0 + \mathbf{V}(t) \quad (10.114)$$

The perturbation is usually an oscillating electric field, which we can write as in eq. (10.115).

$$\mathbf{V}(t) = \sum_k e^{-i\omega_k t} \mathbf{Q} \mathbf{F}_k \quad (10.115)$$

Here ω_k is the frequency of the field, \mathbf{F}_k is the corresponding field strength and \mathbf{Q} is the perturbation operator. The $\mathbf{Q} \mathbf{F}_k$ term should again be interpreted as a sum over all products of components. In most cases the field can be represented by its linear approximation, i.e. \mathbf{Q} is the dipole operator \mathbf{r} and \mathbf{F}_k is a vector containing the x , y and z components of the field. Concentrating on a uniform field of strength F with a single frequency, eq. (10.115) reduces to eq. (10.116).

$$\mathbf{V}(t) = \mathbf{r} F \cos(\omega t) \quad (10.116)$$

The expectation value of a given operator \mathbf{P} can be expanded according to perturbations $\mathbf{Q}, \mathbf{R}, \dots$

$$\langle \mathbf{P} \rangle(t) = \langle \mathbf{P} \rangle(0) + \sum_k e^{-i\omega_k t} \langle \langle \mathbf{P}; \mathbf{Q} \rangle \rangle_{\omega_k} \mathbf{F}_k + \frac{1}{2} \sum_{k,l} e^{-i(\omega_k + \omega_l)t} \langle \langle \mathbf{P}; \mathbf{Q}, \mathbf{R} \rangle \rangle_{\omega_k, \omega_l} \mathbf{F}_k \mathbf{F}_l + \dots \quad (10.117)$$

The *first*-, *second*-, *third*-, etc., order terms are called the *linear*, *quadratic*, *cubic*, etc., responses, and may be interpreted as the change in the property \mathbf{P} due to the perturbations $\mathbf{Q}, \mathbf{R}, \dots$. Note that the linear response is a second-order quantity in the terminology of eq. (10.1), while the quadratic response is a third-order quantity, etc.

For the case where $\mathbf{P} = \mathbf{Q} = \mathbf{R} = \mathbf{r}$ (the position operator), the expansion describes the response of the expectation value of the dipole operator $\langle \Psi_0 | \mathbf{r} | \Psi_0 \rangle$ to a uniform electric field, with the first-order term being the polarizability, eq. (10.5). In the limit where $\omega \rightarrow 0$ (i.e. where the perturbation is time independent), the linear response is

identical to the second-order perturbation formula for a constant electric field (eq. (10.60)), i.e. the $\langle\langle\mathbf{r};\mathbf{r}\rangle\rangle_0$ term determines the static polarizability α . Choosing a non-zero value for ω corresponds to a time-dependent field, i.e. $\langle\langle\mathbf{r};\mathbf{r}\rangle\rangle_\omega$ determines the frequency-dependent polarizability. Similarly, the second-order term $\langle\langle\mathbf{r};\mathbf{r},\mathbf{r}\rangle\rangle_0$ determines the first hyperpolarizability β for a constant field. In the dynamic case, the higher order properties may involve several different frequencies. The corresponding property may be written as $\beta(-\omega;\omega_1,\omega_2)$, with $\omega = \omega_1 + \omega_2$, where ω_1 and ω_2 are associated with the two perturbations. The $\beta(-2\omega;\omega,\omega)$ quantity, for example, determines the second harmonic generation (frequency doubling), while $\beta(-\omega;\omega,0)$ is associated with the electro-optical Pockels effect.

By suitable choices for the $\mathbf{P}, \mathbf{Q}, \mathbf{R}, \dots$ operators a whole variety of properties may be calculated.¹⁸ The polarizability corresponding to imaginary frequencies, for example, provides the van der Waals dispersion coefficients, with the leading term depending on the inverse sixth power of the interatomic distance between atoms A and B.¹⁹

$$E_{\text{dispersion}}(R_{AB}) = -\frac{C_6}{R_{AB}^6} - \frac{C_8}{R_{AB}^8} - \dots$$

$$C_6 = \frac{3}{\pi} \int_0^\infty \alpha_A(i\omega) \alpha_B(i\omega) d\omega$$
(10.118)

An alternative formulation of response theory is in terms of *propagators*, also known as *Greens functions*.²⁰ For two time-dependent operators $\mathbf{P}(t)$ and $\mathbf{Q}(t)$, a propagator may be defined as in eq. (10.119).

$$\langle\langle\mathbf{P}(t); \mathbf{Q}(t')\rangle\rangle = -i\theta(t-t')\langle\Psi_0|\mathbf{P}(t)\mathbf{Q}(t')|\Psi_0\rangle \pm i\theta(t'-t)\langle\Psi_0|\mathbf{Q}(t)\mathbf{P}(t')|\Psi_0\rangle$$
(10.119)

Here the \pm sign depends on whether \mathbf{P} and \mathbf{Q} are number-conserving operators or not, and $\theta(x)$ is the Heaviside step function ($\theta(x) = 0$ for $x < 0$ and $\theta(x) = 1$ for $x > 0$). The propagator may be Fourier transformed to an *energy representation*, also called a *spectral* or *frequency representation*.

$$\langle\langle\mathbf{P}; \mathbf{Q}\rangle\rangle_\omega = \sum_{i \neq 0} \frac{\langle\Psi_0|\mathbf{P}|\Psi_i\rangle\langle\Psi_i|\mathbf{Q}|\Psi_0\rangle}{\omega - E_i + E_0 + i\eta} \pm \frac{\langle\Psi_0|\mathbf{Q}|\Psi_i\rangle\langle\Psi_i|\mathbf{P}|\Psi_0\rangle}{\omega + E_i - E_0 - i\eta}$$
(10.120)

Here η is an infinitesimally small number that ensures that the transformation is also valid when $\omega = \pm(E_i - E_0)$.

If the \mathbf{P}/\mathbf{Q} operators correspond to removal or addition of an electron, the propagator is called an *electron propagator*. The *poles* of the propagator (where the denominator is zero) correspond to ionization potentials and electron affinities.

If the \mathbf{P}/\mathbf{Q} operators are number-conserving operators, the propagator is called a *Polarization Propagator* (PP), and is completely equivalent to the response formulation in eq. (10.117). The $\langle\langle\mathbf{r};\mathbf{r}\rangle\rangle_\omega$ propagator, for example, is given in eq. (10.121) and can be compared with the expression for the static case in eq. (10.60).

$$\langle\langle\mathbf{r}; \mathbf{r}\rangle\rangle_\omega = \sum_{i \neq 0} \frac{|\langle\Psi_0|\mathbf{r}|\Psi_i\rangle|^2}{\omega - E_i + E_0} - \frac{|\langle\Psi_0|\mathbf{r}|\Psi_i\rangle|^2}{\omega + E_i - E_0}$$
(10.121)

The poles correspond to excitation energies, and the residues (numerator at the poles) to transition moments between the reference and excited states (excitation or deexcitation).

Although eq. (10.120) for the propagator appears to involve the same effort as the perturbation approach (sum over all excited states, eq. (10.21)), the actual calculation of the propagator is somewhat different. Returning to the time representation of the polarization propagator, it may be written in terms of a *commutator*.

$$\begin{aligned}\langle\langle\mathbf{P}(t); \mathbf{Q}(t')\rangle\rangle &= -i\theta(t-t')\langle\Psi_0[\mathbf{P}(t), \mathbf{Q}(t')]\Psi_0\rangle \\ [\mathbf{P}(t), \mathbf{Q}(t')] &= \mathbf{P}(t)\mathbf{Q}(t') - \mathbf{Q}(t)\mathbf{P}(t')\end{aligned}\quad (10.122)$$

The Heisenberg equation of motion is shown in eq. (10.123).

$$i\frac{d\mathbf{P}(t)}{dt} = [\mathbf{P}(t), \mathbf{H}] \quad (10.123)$$

When used for the propagator, it yields eq. (10.124).

$$i\frac{d\langle\langle\mathbf{P}(t); \mathbf{Q}(t')\rangle\rangle}{dt} = \delta(t-t')\langle\Psi_0[\mathbf{P}(t), \mathbf{Q}(t')]\Psi_0\rangle + \langle\langle[\mathbf{P}(t), \mathbf{H}]; \mathbf{Q}(t')\rangle\rangle \quad (10.124)$$

Moving back to the frequency representation, and using the fact that $\langle\langle[\mathbf{P}, \mathbf{H}]; \mathbf{Q}\rangle\rangle = \langle\langle\mathbf{P}; [\mathbf{H}, \mathbf{Q}]\rangle\rangle$ allows eq. (10.124) to be written as in eq. (10.125).

$$\omega\langle\langle\mathbf{P}; \mathbf{Q}\rangle\rangle_\omega = \langle\Psi_0[\mathbf{P}, \mathbf{Q}]\Psi_0\rangle + \langle\langle\mathbf{P}; [\mathbf{H}, \mathbf{Q}]\rangle\rangle_\omega \quad (10.125)$$

This shows that a propagator may be written as an expectation value of a commutator plus another propagator involving a commutator with the Hamiltonian. Applying this formula iteratively gives eq. (10.126).

$$\begin{aligned}\langle\langle\mathbf{P}; \mathbf{Q}\rangle\rangle_\omega &= \omega^{-1}\langle\Psi_0[\mathbf{P}, \mathbf{Q}]\Psi_0\rangle + \omega^{-2}\langle\Psi_0[\mathbf{P}, [\mathbf{H}, \mathbf{Q}]]\Psi_0\rangle + \\ &\omega^{-3}\langle\Psi_0[\mathbf{P}, [\mathbf{H}, [\mathbf{H}, \mathbf{Q}]]]\Psi_0\rangle + \cdots\end{aligned}\quad (10.126)$$

The propagator may thus be written as an infinite series of expectation values of increasingly complex operators over the reference wave function.

We now define identity and Hamiltonian *superoperators* as in eq. (10.127).

$$\begin{aligned}\hat{\mathbf{I}}\mathbf{Q} &= \mathbf{Q} \\ \hat{\mathbf{H}}\mathbf{Q} &= [\mathbf{H}, \mathbf{Q}] \\ \hat{\mathbf{H}}^2\mathbf{Q} &= [\mathbf{H}, [\mathbf{H}, \mathbf{Q}]] \\ \hat{\mathbf{H}}^3\mathbf{Q} &= [\mathbf{H}, [\mathbf{H}, [\mathbf{H}, \mathbf{Q}]]]\end{aligned}\quad (10.127)$$

Here the “super” reflects that the $\hat{}$ -operators work on operators rather than functions. The binary product corresponding to a bracket is in superoperator space defined as in eq. (10.128).

$$(\mathbf{P}|\mathbf{Q}) = \langle\Psi_0[\mathbf{P}^\dagger, \mathbf{Q}]\Psi_0\rangle \quad (10.128)$$

The infinite sum eq. (10.126) can then be written as an inverse.

$$\langle\langle\mathbf{P}; \mathbf{Q}\rangle\rangle_\omega = \langle\Psi_0[\mathbf{P}, (\omega\hat{\mathbf{I}} - \hat{\mathbf{H}})^{-1}\mathbf{Q}]\Psi_0\rangle \quad (10.129)$$

This may be further transformed by an “inner projection” onto a complete set of excitation and deexcitation operators, \mathbf{h} . This is equivalent to inserting a “resolution

of the identity” in the operator space (remember that superoperators work on operators).

$$\langle\langle \mathbf{P}; \mathbf{Q} \rangle\rangle_{\omega} = (\mathbf{P}|\mathbf{h})(\mathbf{h}|\omega\hat{\mathbf{I}} - \hat{\mathbf{H}}|\mathbf{h})^{-1}(\mathbf{h}|\mathbf{Q}) \quad (10.130)$$

For the electron propagator we may write \mathbf{h} as in eq. (10.131).

$$\mathbf{h} = \{\mathbf{h}_1, \mathbf{h}_3, \mathbf{h}_5, \dots\} \quad (10.131)$$

Here \mathbf{h}_1 corresponds to addition or removal of an electron, \mathbf{h}_3 to addition or removal of an electron while simultaneously generating a single excitation or deexcitation, \mathbf{h}_5 to addition or removal of an electron while simultaneously generating a double excitation or deexcitation, etc.

For the polarization propagator we may write \mathbf{h} as in eq. (10.132).

$$\mathbf{h} = \{\mathbf{h}_2, \mathbf{h}_4, \mathbf{h}_6, \dots\} \quad (10.132)$$

Here \mathbf{h}_2 generates all single excitations and deexcitations, \mathbf{h}_4 all double excitations and deexcitations, etc.

So far everything is exact. A complete manifold of excitation operators, however, means that all excited states are considered, i.e. a “full CI” approach. Approximate versions of propagator methods may be generated by restricting the excitation level, i.e. truncating \mathbf{h} . A complete specification furthermore requires a selection of the reference, normally taken as an HF, MCSCF or MP wave function.

The simplest polarization propagator corresponds to choosing an HF reference and including only the \mathbf{h}_2 operator, known as the *Random Phase Approximation* (RPA), which is identical to *Time-Dependent Hartree-Fock* (TDHF), with the corresponding density functional version called *Time-Dependent Density Functional Theory* (TDDFT).²¹ For the static case ($\omega = 0$) the resulting equations are identical to those obtained from a coupled Hartree-Fock approach (Section 10.5). When used in conjunction with coupled cluster wave functions, the approach is usually called *Equation Of Motion* (EOM) methods.²²

Splitting the \mathbf{h}_2 operator into an excitation and deexcitation part, $\mathbf{h}_2 = \{\mathbf{e}, \mathbf{d}\}$, allows the propagator to be written as two property vectors times an inverse matrix, often called the *principal propagator*.

$$\langle\langle \mathbf{P}; \mathbf{Q} \rangle\rangle_{\omega} = ((\mathbf{P}|\mathbf{e}) \quad (\mathbf{P}|\mathbf{d})) \begin{pmatrix} \omega\mathbf{1} - \mathbf{A} & -\mathbf{B} \\ -\mathbf{B} & -\omega\mathbf{1} - \mathbf{A} \end{pmatrix}^{-1} \begin{pmatrix} (\mathbf{e}|\mathbf{Q}) \\ (\mathbf{d}|\mathbf{Q}) \end{pmatrix} \quad (10.133)$$

The \mathbf{A} and \mathbf{B} matrices and \mathbf{P}/\mathbf{Q} vectors are defined in eq. (10.134).

$$\begin{aligned} \mathbf{A} &= (\mathbf{e}|\hat{\mathbf{H}}|\mathbf{e}) = \langle\Psi_0|[\mathbf{e}^\dagger, [\mathbf{H}, \mathbf{e}]]\Psi_0\rangle = [\mathbf{d}, [\mathbf{H}, \mathbf{e}]]\Psi_0\rangle \\ \mathbf{B} &= (\mathbf{d}|\hat{\mathbf{H}}|\mathbf{e}) = \langle\Psi_0|[\mathbf{d}^\dagger, [\mathbf{H}, \mathbf{e}]]\Psi_0\rangle = [\mathbf{e}, [\mathbf{H}, \mathbf{e}]]\Psi_0\rangle \\ (\mathbf{P}|\mathbf{e}) &= \langle\Psi_0|[\mathbf{P}^\dagger, \mathbf{e}]\Psi_0\rangle \end{aligned} \quad (10.134)$$

The \mathbf{A} matrix involves elements between singly excited states while \mathbf{B} is given by matrix elements between doubly excited states and the reference. The \mathbf{P}/\mathbf{Q} elements

are matrix elements of the operator between the reference and a singly excited state. If $\mathbf{P} = \mathbf{r}$ this is a transition moment, and in the general case it is often denoted a “property gradient”, by analogy with the case where the operator is the Hamiltonian (eq. (3.68)).

$$\begin{aligned} A_{ij}^{ab} &= \langle \Psi_i^a | \mathbf{H} | \Psi_j^b \rangle - E_0 \delta_{ij} \delta_{ab} \\ B_{ij}^{ab} &= -\langle \Psi_0 | \mathbf{H} | \Psi_{ij}^{ab} \rangle \\ P_i^a &= \langle \Psi_0 | \mathbf{P}^t | \Psi_i^a \rangle \end{aligned} \quad (10.135)$$

The matrix elements may be reduced to orbital energies and two-electron integrals, as described in Section 4.2.1. Although it is not clear from this derivation, the principal propagator in eq. (10.123) is related to the \mathbf{A} matrix in the CHF eq. (10.55), i.e. $(\mathbf{A} - \mathbf{B})$ in eq. (10.133) is the same as \mathbf{A} in eq. (10.55).

Since the dimension of the principal propagator matrix may be large, it is impractical to calculate the inverse matrix in eq. (10.133) directly. In practice, the propagator is therefore calculated in two step, by first solving for an intermediate vector \mathbf{X} (corresponding to \mathbf{U} in eq. (10.53)).

$$\begin{pmatrix} \omega \mathbf{1} - \mathbf{A} & -\mathbf{B} \\ -\mathbf{B} & -\omega \mathbf{1} - \mathbf{A} \end{pmatrix}^{-1} \begin{pmatrix} (\mathbf{e}|\mathbf{Q}) \\ (\mathbf{d}|\mathbf{Q}) \end{pmatrix} = \mathbf{X} \quad (10.136)$$

$$\begin{pmatrix} \omega \mathbf{1} - \mathbf{A} & -\mathbf{B} \\ -\mathbf{B} & -\omega \mathbf{1} - \mathbf{A} \end{pmatrix} \mathbf{X} = \begin{pmatrix} (\mathbf{e}|\mathbf{Q}) \\ (\mathbf{d}|\mathbf{Q}) \end{pmatrix}$$

Multiplying it onto the property gradient gives eq. (10.137).

$$\langle \langle \mathbf{P}; \mathbf{Q} \rangle \rangle_\omega = ((\mathbf{P}|\mathbf{e}) \quad (\mathbf{P}|\mathbf{d})) \mathbf{X} \quad (10.137)$$

The \mathbf{X} vector in eq. (10.136) may be determined by iterative techniques, analogous to those used in direct CI (Section 4.2.4), i.e. the principal propagator matrix is never constructed explicitly. If the \mathbf{Q} vector is set equal to zero in eq. (10.136), the equation corresponds to determining the poles of the principal propagator, i.e. the excitation energies. This is an eigenvalue problem, and finding the principal propagator for a CI wave function is equivalent to diagonalizing the CI Hamiltonian matrix (Section 4.2). For other types of reference wave functions (e.g. HF, MCSCF or MP), the propagator formulation allows a generalization of calculating excitation energies.

The RPA method may be improved either by choosing an MCSCF reference wave function, leading to the MCRPA method, or by extending the operator manifold beyond \mathbf{h}_2 . By expanding the two parts of the propagator (property vector and principal propagator) as a function of the fluctuation potential (difference between the HF and exact electron–electron repulsion), it may be shown that RPA corresponds to terminating the expansion at first order. In the *Second-Order Polarization Propagator Approximation* (SOPPA), the expansion is carried out through second order, which may be shown to require inclusion of the \mathbf{h}_4 operator, and corresponds to choosing an MP wave function as the reference. The *Higher RPA* (HRPA) method may be considered as an approximation to SOPPA, where the part involving the \mathbf{h}_4 operator is neglected. A full third-order propagator model has not been implemented, but a hybrid

method called SOPPA(CCSD) has been proposed, where the first- and second-order perturbation coefficients are replaced with the corresponding coupled cluster amplitudes (eqs (4.36), (4.39) and (4.52)) and this incorporates some of the higher order effects.²³ It tends to perform somewhat better for cases where the HF reference contains significant multi-reference character.

Although the *formal* expression for the propagator and the second-order perturbation formula are identical, involving a sum over all excited states, the final practical expressions for the propagator refer only to the reference wave function, and the basic computational problem involves matrix elements between Slater determinants, and matrix manipulations. Modern implementations of propagator methods are computationally related to the derivative techniques discussed in Section 10.3. The significance is that propagator methods allow a calculation of a property *directly*, without having to construct all the excited states explicitly, i.e. avoiding the sum over states method. This also means that there are no excited wave functions directly associated with a given propagator method. The RPA method includes all singly and some doubly excited states, and typically generates results that are better than those from a CI calculation with single excitations only (if the **B** matrix is neglected in eq. (10.136), the results are identical to CIS), but not as good as CISD. Similarly, the SOPPA method involves an expansion through second order, and typically gives results of MP2 quality, or slightly better.

10.10 Property Basis Sets

The basis set requirements for obtaining a certain accuracy of a given molecular property is usually different from that required for a corresponding accuracy in energy. There is no analogy to the variational principle for properties, since the value in general is not bound. Basis sets for properties must therefore be tailored by adding functions until the desired accuracy is obtained. Given the nature of the perturbation, the specific needs may be very different. An electric field, for example, measures how easily the wave function distorts, i.e. it is primarily dependent on the most loosely bound electrons since they are the ones that are most easily polarized. The important part of the wave function is thus the “tail”, necessitating diffuse function in the basis set. Furthermore, an electric field polarizes the electron cloud, and polarization functions are therefore also important. For perturbation-independent basis functions, there is a “ $2n + 1$ ” rule, i.e. if the unperturbed system is reasonably described by basis functions up to angular momentum L , then a basis set that includes functions up to angular momentum $L + n$ can predict properties up to order $2n + 1$. A minimum description of molecules containing first and second row atoms require *s*- and *p*-functions, implying that *d*-functions are necessary for the polarizability and the first hyperpolarizability, and *f*-functions should be included for the second and third hyperpolarizability. A more realistic description, however, would include *d*-functions for the unperturbed system, necessitating *f*-functions for the polarizability.

A completely different type of property is for example spin–spin coupling constants, which contain information about the interactions of electronic and nuclear spins. One of the operators is a δ function (Fermi contact, eq. (10.73)), which measures the quality of the wave function at a single point, the nuclear position. Since Gaussian functions have an incorrect behaviour at the nucleus (zero derivative compared with the “cusp”

displayed by an exponential function), this requires the addition of a number of very “tight” functions (large exponents) in order to predict coupling constants accurately.

References

1. Y. Yamaguchi, Y. Osamura, J. D. Goddard, H. F. Schaefer III, *A New Dimension to Quantum Chemistry*, Oxford University Press, 1994; C. E. Dykstra, J. D. Augspurger, B. Kirtman, D. J. Malik, *Rev. Comp. Chem.*, **1** (1990), 83; D. B. Chesnut, *Rev. Comp. Chem.*, **8** (1996), 245; R. McWeeny, *Methods of Molecular Quantum Mechanics*, Academic Press, 1992; J. Olsen, P. Jørgensen, *Modern Electronic Structure Theory*, Part II, D. Yarkony, Ed., World Scientific, 1995, pp. 857–990.
2. N. C. Handy, H. F. Schaefer III, *J. Chem. Phys.*, **81** (1984), 5031.
3. T. Helgaker, P. Jørgensen, *Theor. Chim. Acta.*, **75** (1989), 111.
4. J. Gerratt, I. M. Mills, *J. Chem. Phys.*, **49** (1968), 1719.
5. S. P. Karma, *Phys. Rev. Lett.*, **79** (1997), 379.
6. T. Helgaker, M. Jaszunski, K. Ruud, *Chem. Rev.*, **99** (1999), 293.
7. O. Vahtras, O. Loboda, B. Minaev, H. Ågren, K. Ruud, *Chem. Phys.*, **279** (2002), 133.
8. J. Gauss, *J. Chem. Phys.*, **99** (1993), 3629; H. Fukui, T. Baba, H. Inomata, *J. Chem. Phys.*, **105** (1996), 3175; J. Gauss, J. F. Stanton, *Adv. Chem. Phys.*, **123** (2002), 355.
9. P. Manninen, J. Vaara, K. Ruud, *J. Chem. Phys.*, **121** (2004), 1258; D. Jayatilaka, *J. Chem. Phys.*, **108** (1998), 7587.
10. M. Schindler, W. Kutzelnigg, *J. Chem. Phys.*, **76** (1982), 1919.
11. A. E. Hansen, T. D. Bouman, *J. Chem. Phys.*, **82** (1985), 5035.
12. F. London, *J. Phys. Radium*, **8** (1937), 397; R. Ditchfield, *Mol. Phys.*, **27** (1974), 789.
13. K. Ruud, T. Helgaker, K. L. Bak, P. Jørgensen, H. J. Aa. Jensen, *J. Chem. Phys.*, **99** (1993), 3847; K. L. Bak, P. Jørgensen, T. Helgaker, K. Ruud, H. J. Aa. Jensen, *J. Chem. Phys.*, **100** (1994), 6620.
14. K. Wolinski, J. F. Hinton, P. Pulay, *J. Am. Chem. Soc.*, **112** (1990), 8251.
15. J. A. Pople, R. Krishnan, H. B. Schlegel, J. S. Binkley, *Int. J. Quant. Chem. Symp.*, **13** (1979), 225.
16. P. Pulay, *Mol. Phys.*, **17** (1969), 197.
17. P. Deglmann, F. Furche, *J. Chem. Phys.*, **117** (2002), 9535.
18. J. Oddershede, J. R. Sabin, *Int. J. Quant. Chem.*, **39** (1991), 371.
19. J. F. Stanton, *Phys. Rev. A.*, **49** (1994), 1698.
20. J. Oddershede, *Adv. Chem. Phys.*, **69** (1987), 201.
21. F. Furche, K. Burke, *Ann. Rev. Comp. Chem.*, **1** (2005), 19.
22. J. F. Stanton, R. J. Bartlett, *J. Chem. Phys.*, **99** (1993), 5178; H. Koch, P. Jørgensen, *J. Chem. Phys.*, **93** (1990), 3333.
23. S. P. A. Sauer, *J. Phys. B At. Mol. Opt. Phys.*, **30** (1997), 3773.

11 Illustrating the Concepts

In this chapter we will illustrate some of the methods described in the previous sections. It is of course impossible to cover all types of bonding and geometries, but for highlighting the features we will look at the H₂O molecule. This is small enough that we can employ the full spectrum of methods and basis sets, and illustrate some general trends.

11.1 Geometry Convergence

The experimental geometry for H₂O has a bond length of 0.9578 Å and an angle of 104.49°. ^{1,2} Let us investigate how the calculated geometry changes as a function of the theoretical sophistication.

11.1.1 *Ab Initio methods*

We will look at the convergence as a function of basis set and amount of electron correlation (Figure 4.3). For independent-particle methods (HF and DFT) we will use the polarization consistent basis sets (pc-*n*), while for correlated methods (MP2 and CCSD(T)) we will use the correlation consistent basis sets (cc-pVXZ, X = D, T, Q, 5, 6). Table 11.1 shows how the geometry changes as a function of basis set at the HF level of theory, where the quality of the basis sets is indicated by the maximum angular momentum function (L_{\max}) included in the basis set.

The HF results are clearly converged with the cc-pV5Z and pc-3 basis sets, and the HF limit predicts a bond length that is too short, reminiscent of the incorrect dissociation of the single-determinant wave function (Section 4.3). As a consequence, the bond angle becomes too large, owing to an overestimation of the repulsion between the two hydrogens. The underestimation of bond lengths at the HF level is quite general for covalent bonds, while the overestimation of bond angles is not. Although the increased repulsion/attraction between atom pairs in general is overestimated owing to too short bond lengths and too large charge polarization, these factors may

Table 11.1 H₂O geometry as a function of basis set at the HF level of theory

L_{\max}	Basis	R_{OH} (Å)	θ_{HOH} (°)	Basis	R_{OH} (Å)	θ_{HOH} (°)
1				pc-0	0.9619	113.08
2	cc-pVDZ	0.9463	104.61	pc-1	0.9464	105.59
3	cc-pVTZ	0.9406	106.00	pc-2	0.9392	106.41
4	cc-pVQZ	0.9396	106.22	pc-3	0.9396	106.34
5	cc-pV5Z	0.9396	106.33	pc-4	0.9396	106.34
6	cc-pV6Z	0.9396	106.33			

pull in different directions for a larger molecule, and bond angles may either be too large or too small. Note that the bond length decreases as the basis set is enlarged, thus a minimum or DZP type basis may give bond lengths that are longer than the experimental value for some systems. At the HF *limit*, however, covalent bond lengths will normally be too short.

The geometry variation at the MP2 level is shown in Table 11.2, with the change relative to the HF level given as Δ values.

Table 11.2 H₂O geometry as a function of basis set at the MP2 level of theory

Basis	R_{OH} (Å)	θ_{HOH} (°)	ΔR_{OH} (Å)	$\Delta \theta_{\text{HOH}}$ (°)
cc-pVDZ	0.9649	101.90	0.0186	-2.71
cc-pVTZ	0.9591	103.59	0.0185	-2.48
cc-pVQZ	0.9577	104.02	0.0181	-2.20
cc-pV5Z	0.9579	104.29	0.0184	-2.04
cc-pV6Z	0.9581	104.34	0.0185	-1.99

Including electron correlation at the MP2 level increases the bond length by about 0.018 Å, fairly independently of the basis set. As a consequence, the bond angle decreases, by about 2°. Note that the convergence in terms of basis set is much slower than at the HF level. From the observed behaviour the MP2 basis set limit may be estimated as 0.9582 ± 0.0001 Å and $104.40^\circ \pm 0.04^\circ$, which is already in good agreement with the experimental values. H₂O at the equilibrium geometry is a system where the HF is a good zeroth-order wave function, and perturbation methods should consequently converge fast. Indeed, the MP2 method recovers ~94% of the electron correlation energy, as shown in Table 11.7.

The variation at the CCSD(T) level is shown in Table 11.3, with the change relative to the MP2 level given as Δ values.

Table 11.3 H₂O geometry as a function of basis set at the CCSD(T) level of theory

Basis	R_{OH} (Å)	θ_{HOH} (°)	ΔR_{OH} (Å)	$\Delta \theta_{\text{HOH}}$ (°)
cc-pVDZ	0.9663	101.91	0.0014	0.01
cc-pVTZ	0.9594	103.58	0.0003	0.06
cc-pVQZ	0.9579	104.12	0.0002	0.10
cc-pV5Z	0.9580	104.38	0.0001	0.09

Additional correlation with the CCSD(T) method gives only small changes relative to the MP2 level, and the effect of higher order correlation diminishes as the basis set is enlarged. For H₂O the CCSD(T) method is virtually indistinguishable from CCSDT, and presumably very close to the full CI limit.³

The HF wave function contains equal amounts of ionic and covalent contributions (Section 4.3). For covalently bonded systems, such as H₂O, the HF wave function is too ionic, and the effect of electron correlation is to increase the covalent contribution. Since the ionic dissociation limit is higher in energy than the covalent, the effect is that the equilibrium bond length increases when correlation methods are used. For dative bonds, such as metal–ligand compounds, the situation is reversed. In this case the HF wave function dissociates correctly, and bond lengths are normally too long. Inclusion of electron correlation adds attraction between ligands (dispersion interaction), which causes the metal–ligand bond lengths to contract.

The MP2 and CCSD(T) values in Tables 11.2 and 11.3 are for correlation of the valence electrons only, i.e. the frozen-core approximation. In order to assess the effect of core electron correlation, the basis sets need to be augmented with tight polarization functions. The corresponding MP2 results are shown in Table 11.4, where the Δ values refer to the change relative to the valence-only MP2 with the same basis set. Essentially identical changes are found at the CCSD(T) level.

Table 11.4 H₂O geometry as a function of basis set at the MP2 level of theory including all electrons in the correlation

Basis	R_{OH} (Å)	θ_{HOH} (°)	ΔR_{OH} (Å)	$\Delta \theta_{\text{HOH}}$ (°)
cc-pCVDZ	0.9643	101.91	−0.0005	0.04
cc-pCVTZ	0.9580	103.63	−0.0008	0.11
cc-pCVQZ	0.9569	104.14	−0.0009	0.12
cc-pCV5Z	0.9570	104.41	−0.0009	0.12
cc-pCV6Z	0.9572	104.47	−0.0009	0.13

The effect of core electron correlation is small: a small decrease of the bond length and a corresponding small increase in bond angle. Addition of the CCSD(T)-MP2 changes (Table 11.3) to the MP2 basis set limiting results in Table 11.4 gives a bond length of 0.9573 Å and an angle of 104.56°. Further basis set increases will presumably lead to increases of ~ 0.0002 Å and $\sim 0.08^\circ$. Relativistic effects at the Dirac–Fock–Breit level of theory have been reported to give changes of 0.00016 Å and -0.07° .² Including these corrections allows a final predicted structure of 0.9577 Å and 104.57°, which can be compared with the experimental values of 0.9578 Å and 104.49°.

These results show that *ab initio* methods can give results of very high accuracy, provided that sufficiently large basis sets are used. Unfortunately, the combination of highly correlated methods, such as CCSD(T), and large basis sets means that such calculations are computationally expensive. For the H₂O system a CCSD(T) calculation

with the cc-pV5Z basis is already quite demanding. The results also show, however, that a quite respectable level of accuracy is reached at the MP2/cc-pVTZ level, which is applicable to a much larger variety of molecules. Furthermore, the errors at a given level are quite systematic, and relative values (comparing for example changes in geometries upon introduction of substituents) will be predicted with a substantially higher accuracy.

It should also be noted that the effect of electron correlation at the MP2 level (relative to HF) is largely independent of the basis set, but there is a significant coupling between the basis set and the higher order correlation (beyond MP2) effect. The importance of higher order electron correlation decreases as the basis set is enlarged. This suggests that it is better to invest a given amount of computer time in performing a large basis set MP2 calculation than a highly correlated calculation with a modest basis, at least when the HF is a good zeroth-order wave function.

11.1.2 Density functional methods

The two variables in DFT methods are the basis set and the choice of the exchange–correlation potential. The performance for six popular functionals on the geometry for the pc-*n* basis sets is given in Tables 11.5 and 11.6. The LSDA functional employs the uniform electron gas approximation, the BLYP, PBE and HCTH functionals are of the gradient-corrected type, while the B3LYP and PBE0 are hybrid types that contain a fraction of Hartree–Fock exchange (Section 6.5.4). The grid size for the numerical integration of the exchange–correlation energy is sufficiently large that the error from incomplete grids can be neglected.

Table 11.5 H₂O bond distances (Å) as a function of basis set with various DFT functionals

Basis	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0
pc-0	0.9878	0.9962	0.9936	0.9854	0.9841	0.9806
pc-1	0.9764	0.9791	0.9763	0.9656	0.9683	0.9645
pc-2	0.9696	0.9706	0.9689	0.9589	0.9604	0.9574
pc-3	0.9700	0.9704	0.9689	0.9589	0.9604	0.9576
pc-4	0.9700	0.9704	0.9689	0.9590	0.9604	0.9576

Table 11.6 H₂O bond angles (°) as a function of basis set with various DFT functionals

Basis	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0
pc-0	111.82	109.27	109.40	109.43	110.72	110.93
pc-1	104.15	103.24	103.09	103.22	104.06	103.99
pc-2	105.10	104.56	104.27	104.52	105.19	104.98
pc-3	104.98	104.52	104.21	104.44	105.13	104.90
pc-4	104.98	104.52	104.21	104.42	105.13	104.90

The geometry displays a convergence characteristic similar to the wave mechanics HF method (Table 11.1). A TZP type basis (pc-2) gives good results, and a QZP type (pc-3) is essentially converged to the basis set limiting value. The basis set limiting values can be compared with the experimental values of 0.9578 Å and 104.48°, and the deviations are thus the inherent errors associated with the functionals. For this particular molecule and property, the HCTH and PBE0 functionals perform best.

11.2 Total Energy Convergence

The total energy in *ab initio* theory is given relative to the separated particles, i.e. bare nuclei and electrons. The experimental value for an atom is the sum of all the ionization potentials; for a molecule there are in addition contributions from the molecular bonds and associated zero point energies. The experimental value for the total energy of H₂O is −76.480 au and the estimated contribution from relativistic effects is −0.045 au. Including also a mass correction of 0.0028 au (a non-Born–Oppenheimer effect that accounts for the difference between finite and infinite nuclear masses) allows the “experimental” non-relativistic energy to be estimated as -76.438 ± 0.003 au.⁴

For the cc-pVDZ basis set, the full CI result is available,⁵ which allows an assessment of the performance of various approximate methods. The percent of the electron correlation recovered by different methods is shown in Table 11.7.

Table 11.7 Percent electron correlation recovered by various methods in the cc-pVDZ basis

Method	%EC
MP2	94.0
MP3	97.0
MP4	99.5
MP5	99.8
CCSD	98.3
CCSD(T)	99.7
CISD	94.5
CISDT	95.8
CISDTQ	99.9

As already mentioned, the H₂O molecule is an easy system, where the HF wave function provides a good reference. Furthermore, since there are only ten electrons in H₂O the effect of higher order electron correlation is small. The intraorbital correlation between electron pairs dominates the correlation energy for such a small system, and the doubly excited configurations, which mainly describe the pair correlation, accounts for a large fraction of the total correlation energy. Consequently even the simple MP2 method performs exceedingly well, and the CCSD(T) result is for practical purposes identical to the full CI result. For such simple systems, the MP2 and MP3 percent correlations are probably significantly higher than would be expected for a larger system.

The calculated total energy as a function of basis set and electron correlation (valence electrons only) at the experimental geometry is given in Table 11.8. As the

cc-pVXZ basis sets are fairly systematic in how they are extended from one level to the next, there is some justification for extrapolating the results to the “infinite” basis set limit (Section 5.4.6). The HF energy is expected to have an exponential behaviour, and a functional form⁶ of the type $A + B(L + 1)\exp(-C\sqrt{L})$ with $L = 4, 5$ and 6 yields an infinite basis set limit of -76.0675 au, in good agreement with the estimated HF limit of -76.0674 au.⁷

Table 11.8 Total energy (+76 au) as a function of basis set and electron correlation (valence only)

Method	cc-pVDZ	cc-pVTZ	cc-pVQZ	cc-pV5Z	cc-pV6Z	cc-pV ∞ Z
HF	-0.0268	-0.0571	-0.0648	-0.0670	-0.0674	-0.0675
MP2	-0.228	-0.319	-0.348	-0.359	-0.363	-0.369
MP3	-0.235	-0.323	-0.349	-0.358	-0.361	-0.365
MP4	-0.241	-0.333	-0.361	-0.371	-0.374	-0.378
MP5	-0.241	-0.332	-0.359			
CCSD	-0.238	-0.325	-0.351	-0.360	-0.362	-0.365
CCSD(T)	-0.241	-0.332	-0.360	-0.369	-0.372	-0.376
CISD	-0.230	-0.314	-0.339	-0.348	-0.350	-0.353

The correlation energy is expected to have an inverse power dependence on the highest angular momentum once the basis set reaches a sufficient (large) size. Extrapolating the correlation contribution for $L = 5$ and 6 with a function of the type $A + BL^{-3}$ yield the cc-pV ∞ Z values in Table 11.8. The extrapolated CCSD(T) energy is -76.376 au, yielding a valence correlation energy of -0.308 au.

The magnitude of the core correlation can be evaluated by including the oxygen 1s-electrons and using the cc-pCVXZ basis sets, with the results shown in Table 11.9.

Table 11.9 Total energy (+76 au) as a function of basis set and electron correlation (all electrons)

Method	cc-pCVDZ	cc-pCVTZ	cc-pCVQZ	cc-pCV5Z	cc-pCV ∞ Z	%EC
HF	-0.0272	-0.0573	-0.0649	-0.0671	-0.0683	0.0
MP2	-0.269	-0.375	-0.408	-0.419	-0.429	97.3
MP3	-0.276	-0.380	-0.410	-0.420	-0.429	97.3
MP4	-0.282	-0.391	-0.422	-0.433	-0.443	101.1
MP5	-0.282	-0.389				
CCSD	-0.279	-0.382	-0.411	-0.421	-0.430	97.6
CCSD(T)	-0.282	-0.390	-0.421	-0.431	-0.440	100.3
CISD	-0.269	-0.368	-0.397	-0.406	-0.414	93.3

The HF energies change very little upon inclusion of the tight basis functions, but the HF limit is estimated with less accuracy as the extrapolation is done using basis sets with lower angular momentum functions than in Table 11.8. Using the HF result from Table 11.8, the extrapolated CCSD(T) correlation energy is -0.372 au. Assuming that the CCSD(T) method provides 99.7% of the full CI value, as indicated in Table 11.7, the extrapolated correlation energy becomes -0.373 au, well within the error limits on the estimated experimental value of -0.371 ± 0.003 au. The core (and core–valence) electron correlation is thus 0.065 au, which is comparable to the

value for the valence electrons (i.e. 0.308 divided between four electron pairs is 0.077 au).

The percent of the total correlation energy is given in parenthesis in Table 11.9; in the infinite basis set limit the MP2 method recovers 97.3%. Notice, however, that while the perturbation series is smoothly convergent with the cc-pVDZ basis, it becomes oscillating with the larger basis set. With the cc-pCVTZ basis, the MP5 result is higher in energy than MP4, and with the cc-pCV5Z the MP3 result is higher than the MP2 value. This may be an indication that the perturbation series is actually divergent in a sufficiently large basis set. The extrapolated MP4 value is in perfect agreement with the experimental estimate, but this is probably fortuitous. The CISD method performs rather poorly, yielding results that are worse than MP2 but at a cost similar to an MP4 calculation.

Since the CCSD(T) result is essentially equivalent to a full CI (Table 11.7), the data show that the cc-pCVDZ basis is able to provide 69% of the *total* correlation energy. The corresponding values for the cc-pCVTZ, cc-pCVQZ and cc-pCV5Z basis sets are 90%, 96% and 98%, respectively. Slightly lower percentages have been found in other systems.⁸ This illustrates the slow convergence of the correlation energy as a function of basis set. Each step up in basis set quality roughly doubles the number of functions. The cc-pCVDZ basis is capable of recovering 69% of the correlation energy, and improving the basis from cc-pCVDZ to cc-pCVTZ allows an additional 21% to be calculated. The next step up gives only 6% and the expansion from cc-pCVQZ to cc-pCV5Z only 2%. The last 5–10% of the correlation energy is therefore hard to get, requiring very large basis sets. This slow convergence is the principal limitations of traditional *ab initio* methods. The CCSD(T)/cc-pCV5Z total energy is still 18 kJ/mol off the experimentally derived non-relativistic value, with the remaining error being distributed roughly equally between incomplete basis set and incomplete electron correlation effects. These errors are comparable to the Born–Oppenheimer correction of 7 kJ/mol, and substantially smaller than the relativistic correction of 118 kJ/mol. Calculating the total energy with an accuracy of a few kJ/mol is thus only borderline possible for this simple system.

Although the total energy calculated by DFT methods should in principle converge to the “experimental” value (−76.438 au), there are no upper or lower bounds for the currently employed methods with approximate exchange–correlation functionals. Indeed, all the gradient-corrected methods used here (BLYP, PBE and HCTH) give total energies well below the “experimental” value with the pc-4 basis set.

11.3 Dipole Moment Convergence

As examples of molecular properties we will look at how the dipole moment and harmonic vibrational frequencies converge as a function of level of theory.

11.3.1 *Ab Initio* methods

The experimental value for the dipole moment is 1.847 debye,⁹ and the calculated value at various levels of theory is shown in Table 11.10.

The dipole moment may be considered as the response of the wave function (energy) to the presence of an external electric field, in the limit where the field strength is van-

Table 11.10 H₂O dipole moment (debye) as a function of theory (valence correlation only); the experimental value is 1.847 debye

Basis	HF	MP2	CCSD(T)	Basis	HF	MP2	CCSD(T)
cc-pVDZ	2.057	1.964	1.936	aug-cc-pVDZ	2.000	1.867	1.848
cc-pVTZ	2.026	1.922	1.903	aug-cc-pVTZ	1.984	1.852	1.839
cc-pVQZ	2.008	1.904	1.890	aug-cc-pVQZ	1.982	1.858	1.848
cc-pV5Z	2.003	1.895	1.884	aug-cc-pV5Z	1.982	1.861	

ishingly small (Section 10.1.1). It is consequently sensitive to the representation of the wave function “tail”, i.e. far from the nuclei, and diffuse functions are therefore expected to be important. Although the results with the regular cc-pVXZ basis sets may be converging, the rate of convergence is slow, as compared with the results for the basis sets augmented with diffuse functions. This illustrates that care must be taken when calculating properties other than the total energy, as standard basis sets may not be able to describe important aspects of the wave function.

The HF dipole moment is too large, which is quite general, as the HF wave function overestimates the ionic contribution. The MP2 procedure recovers the large majority of the correlation effect, but the convergence with the aug-cc-pVXZ basis sets is not smooth, and does not readily allow an extrapolation. The CCSD(T) result with the aug-cc-pVQZ basis is very close to the experimental value, although remaining basis set effects and further correlation may change the value slightly. As expected for this property, the effect of core correlation is small, as shown by MP2 calculations in Table 11.11.

Table 11.11 H₂O dipole moment (debye) as a function of theory (all electrons)

Basis	HF	MP2	CCSD(T)
aug-cc-pCVDZ	2.001	1.868	1.849
aug-cc-pCVTZ	1.983	1.857	1.843

11.3.2 Density functional methods

Table 11.10 establishes that diffuse functions are mandatory for calculating dipole moments, and only the aug-pc-*n* basis set have been used with DFT methods. The calculated results are given in Table 11.12.

The calculated dipole moment is remarkably insensitive to the size of the basis set, once polarization functions have been included (i.e. at least aug-pc-1). Note that the LSDA value in this case is substantially better than the GGA functionals (BLYP, PBE and HCTH), i.e. this is a case where the theoretically “poorer” method provides better results than the more advanced gradient methods. Inclusion of “exact” exchange (B3LYP and PBE0) again improves the performance, and provides results very close to the experimental value, even with relatively small basis sets.

Table 11.12 H₂O dipole moment (debye) as a function of DFT functional and basis set; the experimental value is 1.847 debye

Basis	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0
aug-pc-0	2.721	2.602	2.618	2.605	2.655	2.675
aug-pc-1	1.823	1.762	1.760	1.758	1.821	1.826
aug-pc-2	1.837	1.781	1.779	1.782	1.837	1.842
aug-pc-3	1.834	1.778	1.774	1.779	1.833	1.837
aug-pc-4	1.833	1.778	1.774	1.777	1.833	1.836

11.4 Vibrational Frequency Convergence

The experimental values for the fundamental vibrational frequencies are 1595, 3657 and 3756 cm⁻¹, while the corresponding *harmonic* values are 1649, 3832 and 3943 cm⁻¹.¹⁰ The differences due to anharmonicity are thus 54, 175 and 187 cm⁻¹, i.e. 3–5% of the harmonic values.

11.4.1 *Ab Initio* methods

The calculated harmonic frequencies at the HF level are given in Table 11.13.

Table 11.13 H₂O HF harmonic frequencies (cm⁻¹) as a function of basis set

Basis	ν_1	ν_2	ν_3
pc-0	1690	3966	4145
pc-1	1751	4120	4233
pc-2	1744	4138	4239
pc-3	1748	4131	4232
pc-4	1748	4130	4231
Experimental	1649	3832	3943

Vibrational frequencies are examples of a slightly more complicated property. The frequencies are obtained from the force constant matrix (second derivative of the energy), *evaluated at the equilibrium geometry* (Section 16.2.2). Both the equilibrium geometry and the shape of the energy surface depend on the theoretical level. Part of the change in frequencies is due to changes in the geometry since the force constant in general decreases with increasing bond length.

The HF vibrational frequencies are too high by about 7% relative to the experimental harmonic values, and by 10–13% relative to the anharmonic values. This overestimation is due to the incorrect dissociation and the corresponding bond lengths being too short (Table 11.1), and is consequently quite general. Vibrational frequencies at the HF level are therefore often scaled by ~0.9 to partly compensate for these systematic errors.¹¹

The inclusion of electron correlation normally lowers the force constants, since the correlation energy increases as a function of bond length. This usually means that

Table 11.14 H₂O MP2 harmonic frequencies (cm⁻¹) as a function of basis set (only valence electrons are correlated)

Basis	ν_1	ν_2	ν_3
cc-pVDZ	1678	3852	3971
cc-pVTZ	1651	3855	3976
cc-pVQZ	1643	3855	3978
cc-pV5Z	1636	3849	3974
Experimental	1649	3832	3943

Table 11.15 H₂O CCSD(T) harmonic frequencies (cm⁻¹) as a function of basis set (only valence electrons are correlated)

Basis	ν_1	ν_2	ν_3
cc-pVDZ	1690	3822	3928
cc-pVTZ	1669	3841	3946
cc-pVQZ	1659	3845	3952
Experimental	1649	3832	3943

vibrational frequencies decrease, although there are exceptions (vibrational frequencies also depend on off-diagonal force constants). The values calculated the MP2 and CCSD(T) levels are shown in Tables 11.14 and 11.15.

The MP2 treatment recovers the majority of the correlation effect, and the CCSD(T) results with the cc-pVQZ basis sets are in good agreement with the experimental values. The remaining discrepancies of 9, 13 and 10 cm⁻¹ are mainly due to basis set inadequacies, as indicated by the MP2/cc-pV5Z results. The MP2 values are in respectable agreement with the experimental harmonic frequencies, but of course still overestimate the experimental fundamental ones by the anharmonicity. For this reason, calculated MP2 harmonic frequencies are often scaled by ~ 0.97 for comparing with experimental results.¹¹

The effect of core electron correlation is small, as shown in Table 11.16. It should be noted that the valence and core correlation energy per electron pair is of the same magnitude. However, the core correlation is almost constant over the whole energy surface, and consequently contributes very little to properties depending on relative energies, such as vibrational frequencies. It should be noted that relativistic corrections for the frequencies are expected to be of the order of 1 cm⁻¹ or less.¹²

Table 11.16 H₂O MP2 harmonic frequencies (cm⁻¹) as a function of basis set (all electrons are correlated)

Basis	ν_1	ν_2	ν_3
cc-pCVDZ	1679	3853	3973
cc-pCVTZ	1651	3857	3976

For comparing with experimental frequencies (which necessarily are anharmonic), there is normally little point in improving the theoretical level beyond MP2 with a TZP type basis set unless anharmonicity constants are calculated explicitly. Although anharmonicity can be approximately accounted for by scaling the harmonic frequencies by ~ 0.97 , the remaining errors in the harmonic force constants at this level are normally smaller than the corresponding errors due to variations in anharmonicity.

11.4.2 Density functional methods

The harmonic frequencies calculated with various DFT methods as a function of basis set are shown in Tables 11.17–11.19.

Table 11.17 H₂O lowest harmonic frequency (cm⁻¹) as a function of basis set with various DFT functionals; the experimental value is 1649 cm⁻¹

Basis	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0
pc-0	1474	1535	1539	1567	1565	1578
pc-1	1548	1597	1596	1627	1628	1635
pc-2	1544	1595	1590	1617	1625	1630
pc-3	1549	1597	1594	1621	1629	1635
pc-4	1550	1598	1594	1621	1629	1635

Table 11.18 H₂O second lowest harmonic frequency (cm⁻¹) as a function of basis set with various DFT functionals; the experimental value is 3832 cm⁻¹

Basis	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0
pc-0	3588	3453	3510	3574	3620	3691
pc-1	3690	3611	3664	3760	3767	3835
pc-2	3730	3669	3710	3796	3811	3870
pc-3	3718	3667	3707	3794	3807	3865
pc-4	3718	3666	3706	3794	3807	3865

Table 11.19 H₂O highest harmonic frequency (cm⁻¹) as a function of basis set with various DFT functionals; the experimental value is 3943 cm⁻¹

Basis	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0
pc-0	3787	3641	3699	3767	3806	3879
pc-1	3817	3727	3784	3884	3882	3955
pc-2	3842	3772	3816	3909	3915	3977
pc-3	3827	3768	3811	3905	3909	3970
pc-4	3827	3768	3811	3905	3909	3970

The convergence as a function of basis set is similar to that observed for the HF method. The hybrid B3LYP and PBE0 functionals again show the best performance. At the basis set limit, the deviations from the experimental harmonic frequencies are $\sim 30\text{ cm}^{-1}$, which is comparable to the results obtained with the MP2 method (Table 11.14). It is also clear from Tables 11.17–11.19 that inclusion of “exact” exchange (B3LYP and PBE0) substantially improves the performance. The “pure” DFT gradient methods, BLYP and PBE, have errors of $\sim 150\text{ cm}^{-1}$ for the stretching frequencies and $\sim 50\text{ cm}^{-1}$ for the angle bending.

11.5 Bond Dissociation Curves

As seen in Table 11.9, it is very difficult to converge the total energy to an accuracy of a few kJ/mol. The total energy, however, is in almost all cases irrelevant; the important quantity is the relative energy. Let us now examine how the *shape* of a potential energy surface depends on the theoretical level. We will look at two cases: stretching one of the O—H bonds in H_2O , and the HOH bending potential. The O—H dissociation curve is a case where the main change is associated with the difference in electron correlation between the two electrons in the bond being stretched. It should be noted that transition structures typically have bonds that are elongated by $0.5\text{--}0.8\text{ \AA}$, and the performance for the dissociation curve in this range will model the behaviour for describing bond breaking/forming reactions. The HOH bending energy, on the other hand, does not involve any bond breaking, and should therefore be less sensitive to the level of theory.

11.5.1 Basis set effect at the Hartree–Fock level

Figure 11.1 shows the bond dissociation curves at the HF level with the STO-3G, 3-21G, 6-31G(d,p), cc-pVDZ and cc-pVQZ basis sets.

The total energy drops considerably upon going from the STO-3G to the 3-21G and again to the 6-31G(d,p) basis. This is primarily due to the improved description of the oxygen 1s-orbital. The two different types of DZP basis sets, 6-31G(d,p) and cc-pVDZ, give very similar results, and the improvement upon going to the cc-pVQZ basis is relatively minor.

More important than the total energy is the shape of the curve, i.e. the energy relative to the equilibrium value which is shown in Figure 11.2.

The minimal STO-3G basis increases much more steeply than the other basis sets, while the 3-21G is slightly too low in the ΔR range of $0.3\text{--}1.3\text{ \AA}$. Considering that the STO-3G and 3-21G basis set have the same number of primitive GTOs (Section 5.4.1), it is clear that uncontraction of the valence orbitals greatly improves the flexibility. The 6-31G(d,p), cc-pVDZ and cc-pVQZ basis sets give essentially identical curves, i.e. improvement of the basis set beyond DZP has a very minor effect at the HF level. Note also that the total energy for the 6-31G(d,p) basis is $\sim 0.05\text{ au}$ ($\sim 130\text{ kJ/mol}$, Figure 11.1) above the HF limit, but this error is constant to within a few kJ/mol over the whole range.




Figure 11.1 Bond dissociation curves for H₂O at the HF level (absolute energies)




Figure 11.2 Bond dissociation curves for H₂O at the HF level (relative energies)

11.5.2 Performance of different types of wave function

We will now look at how different types of wave function behave when the O—H bond is stretched. The basis set used in all cases is the aug-cc-pVTZ, and the reference curve is taken as the [8,8]-CASSCF result, which is slightly larger than a full valence CI. As mentioned in Section 4.6, this allows a correct dissociation, and since all the valence electrons are correlated, it will generate a curve close to the full CI limit. The bond dissociation energy calculated at this level is 511 kJ/mol, which is comparable to the experimental value of 527 kJ/mol.

H₂O is a closed shell singlet and the HF wave function near the equilibrium geometry is of the RHF type. As one of the bonds is stretched, however, a UHF type will become lower in energy at some point (Section 4.4). Beyond this instability point, electron correlation methods may be based either on the RHF or UHF reference. The UHF wave function will be spin contaminated, which has some consequences as shown below. It should be noted that for open-shell species one similarly has the option of using either a ROHF or UHF reference wave function, but in such cases they will be different at all geometries, also near the equilibrium. In many cases, however, the UHF wave function is only slightly spin contaminated, and both approaches will then give similar results.

Figure 11.3 illustrates the behaviour of the single-determinant wave functions, RHF, UHF and PUHF (projected UHF, Section 4.4). The RHF energy continues to increase as the bond is stretched since it has the wrong dissociation limit, while the UHF converges to a value of 366 kJ/mol. At the equilibrium geometry the two electrons in the O—H bonding orbital are correlated, but this correlation energy disappears once the bond is broken. The UHF wave function correctly describes the dissociation limit in




Figure 11.3 RHF, UHF and PUHF dissociation curves for H₂O

terms of energy, but does not recover any of the electron correlation at equilibrium (by definition, since UHF = RHF here). The difference between the UHF dissociation energy and the CASSCF value is therefore a measure of the amount of electron correlation in the O—H bond. With the present basis set this is 140 kJ/mol, a typical value for the correlation energy between two electron in the same spatial MO.

At the dissociation limit the UHF wave function is essentially an equal mixture of a singlet and triplet state, as discussed in Section 4.4. Removal of the triplet state by projection (PUHF) lowers the energy in the intermediate range, but has no effect when the bond is completely broken since the singlet and triplet states are degenerate here.

The RHF/UHF instability point with this basis occurs when the bond is stretched 0.42 Å. Figure 11.4 shows the behaviour of the energy curves in more detail in this region. It is seen that the PUHF has a discontinuous derivative at the instability point, and there is furthermore a shallow minimum right after the instability point, at an elongation of ~0.50 Å.




Figure 11.4 RHF, UHF and PUHF dissociation curves for H₂O near the instability point

Since the RHF curve is too high in the transition structure region ($\Delta R \sim 0.5\text{--}0.8$ Å), it is clear that RHF activation energies in general will be too large. UHF activation energies may either be too high or too low, but the PUHF value will essentially always be too low. Furthermore, the shape of a spin contaminated UHF energy surface will be too flat, and PUHF surfaces will be qualitatively wrong in the TS region. Spin contaminated UHF wave functions should consequently not be used for geometry optimizations.

The corresponding difference between restricted, unrestricted and projected unrestricted wave functions at the MP2 level is shown in Figure 11.5.

The RMP2 rises too high, owing to the wrong dissociation limit of the underlying RHF. Both the UMP2 and PUMP2 dissociation energies are in reasonable agreement




Figure 11.5 RMP2, UMP2 and PUMP2 dissociation curves for H_2O

with the CASSCF value, but it is clear that the UMP2 energy is too high in the “intermediate” range owing to spin contamination. The PUMP2 curve on the other hand, traces the reference CASSCF values closely. Figure 11.6 shows the curves in more detail near the RHF/UHF instability point.




Figure 11.6 RMP2, UMP2 and PUMP2 dissociation curves for H_2O near the instability point

The UMP2 energy is *higher* than the RMP2, although the UHF energy is *lower* than the RHF. At the HF level, the UHF energy is lowest owing to a combination of spin contamination and inclusion of electron correlation (Section 4.8.2). Since the MP2 procedure recovers most of the electron correlation, only the energy rising effect due to spin contamination remains, and the UMP2 energy becomes higher than RMP2. Removing the unwanted spin components makes the PUMP2 energy very similar to RMP2 for elongations less than ~ 1 Å, but is significantly better at longer bond lengths owing to the correct dissociation of the UHF wave function. The RMP2 energy follows the “exact” curve closely out to a ΔR of ~ 0.5 Å, and is in respectable agreement out to ~ 1.0 Å. RMP2 activation energies are therefore often in quite reasonable agreement with experimental or higher level theoretical values. It should also be noted that the discontinuity at the PUHF level essentially disappears when the projection is carried out on the MP2 wave function.

Figure 11.7 and 11.8 show the effect of extending the perturbation series at the RMP and UMP levels.

Addition of more terms in the perturbation series improves the results, although the effect of MP3 compared with MP2 is minute. As the bond is stretched more than ~ 1.5 Å, the perturbation series breaks down owing to the RHF wave function becoming a too poor reference, and the energies start to decrease. The RMP4 method performs well out to an elongation of ~ 1.0 Å, and in the TS region where the bond is stretched 0.5 – 0.8 Å, the MP4 error is less than a few kJ/mol. Although real transition structures usually have more than one breaking/forming bond, and therefore are more sensitive to correlation effects, it is often found that the MP4 method with a suitable large basis can reproduce activation energies to within a few kJ/mol.

The improvement by extending the perturbation series beyond second order is small when a UHF wave function is used as the reference, i.e. the higher order terms do very




Figure 11.7 RMP2, RMP3 and RMP4 dissociation curves for H₂O




Figure 11.8 UMP2, UMP3 and UMP4 dissociation curves for H₂O

little to reduce the spin contamination. In the dissociation limit the spin contamination is inconsequential, and the MP2, MP3 and MP4 results are all in reasonable agreement with the “exact” CASSCF result (but too high compared with the experimental result due to basis set limitations).

Figures 11.9 and 11.10 compare the performance of the CCSD and CCSD(T) methods, based on either an RHF or UHF reference wave function.




Figure 11.9 RCCSD and RCCSD(T) dissociation curves for H₂O




Figure 11.10 UCCSD and UCCSD(T) dissociation curves for H_2O

Compared with the RMP_n curves (Figure 11.7), it can be seen that the infinite nature of coupled cluster performs somewhat better as the reference wave function becomes increasingly poor. While the RMP_4 energy curve follows the “exact” out to an elongation of $\sim 1.0 \text{ \AA}$, the CCSD(T) has the same accuracy out to $\sim 1.5 \text{ \AA}$. Eventually, however, the wrong dissociation limit of the RHF wave also makes the coupled cluster methods break down, and the energy starts to decrease.

The spin contamination makes the UCC energy curves somewhat too high in the intermediate region, but the infinite nature of coupled cluster methods is significantly better at removing unwanted spin states as compared with UMP_n methods (Figure 11.8).

The only generally applicable CI method is CISD, where the singly and doubly excited configurations are treated variationally. These are also part of the MP_4 method, which additionally has a term arising from disconnected quadruples, i.e. products of D-configurations, as well as a term due to (connected) triples. The CCSD method includes effects due to higher order products of singles and doubles, i.e. sextuples, octuples, etc. It is the inclusion of the product excitations that makes the MP and CC methods size extensive. Considering only the single and double excitations, and products thereof, allows a comparison between methods, and the performance of the CISD, $\text{MP}_4(\text{SDQ})$ and CCSD models are shown in Figure 11.11.




Figure 11.11 RMP4(SDQ), RCCSD and CISD dissociation curves for H_2O

It can be clearly seen that the CISD curve is worse than either of the other two, which are essentially identical out to a ΔR of 1.3 \AA . The size inconsistency of the CISD method also has consequences for the energy curve where the bond is only half broken. Figure 11.11 illustrates why the use of CI methods has declined over the years: they normally give less accurate results compared with MP or CC methods, but at a similar or higher computational cost. Furthermore, it is difficult to include the important triply excited configurations in CI methods (CISDT scales as M_{basis}^8), but relatively easy in MP or CC methods (MP4 and CCSD(T) scales as M_{basis}^7).

11.5.3 Density functional methods

The performance of various DFT methods resembles the HF results. A restricted type determinant leads to an incorrect dissociation, while an unrestricted determinant has the energetically correct dissociation limit. Figure 11.12 shows the performance of restricted and unrestricted type determinants with the BLYP and B3LYP functionals.

It is immediately clear that DFT methods do not have the “spin contamination” problem in the intermediate region; indeed, spin contamination is not well defined in DFT.¹³ Furthermore, as electron correlation is implicitly included, DFT methods are closer in shape to the CASSCF curve. Removing the “spin contamination” by projection methods results in discontinuous derivatives and artificial minima, analogously to the PUHF case in Figures 11.3 and 11.4, and should consequently not be employed.¹⁴




Figure 11.12 Bond dissociation curve for DFT methods

11.6 Angle Bending Curves

The angle bending in H_2O occurs without breaking any bonds, and the electron correlation energy is therefore relatively constant over the whole curve. The HF, MP2 and MP4 bending potentials are shown in Figure 11.13, where the reference curve is taken from a parametric fit to a large number of spectroscopic data.¹⁵

The HF and MP2 methods underestimate the barrier for linearity by 1 and 2 kJ/mol, respectively, while the CCSD(T) result is too high by 1 kJ/mol. The HF curve is slightly too high for small bond angles, while both the MP2 and CCSD(T) results are within a few tenths of a kJ/mol of the exact result over the whole curve. Compared with the bond dissociation discussed above, it is clear that relative energies of conformations that have similar bonding are fairly easy to calculate. While the HF and MP2 total energies with the aug-cc-pVTZ basis are ~1000 and ~300 kJ/mol higher than the exact values at the equilibrium geometry (Table 11.8), these errors are essentially constant over the whole surface.

11.7 Problematic Systems

The H_2O case is an example of a system where it is relatively easy to obtain good results. Nature is not always so kind; let us look at a couple of “theoretically difficult” cases.




Figure 11.13 Angle bending curves for H_2O

11.7.1 The geometry of FOOF

The FOOF molecule has an experimental geometry with an O—O bond length of 1.217 Å and an F—O bond of 1.575 Å.¹⁶ The calculated bond distances at different levels of theory with the aug-cc-pVTZ basis set are given in Table 11.20.

The results in Table 11.20 clearly show that the results are very sensitive to the inclusion of electron correlation. The MP4(SDQ) geometry is very similar to the CCSD one

Table 11.20 Bond distance (Å) in FOOF with the aug-cc-pVTZ basis set

Method	R_{OO}	R_{FO}
HF	1.300	1.356
MP2	1.166	1.619
MP3	1.300	1.427
MP4(SDQ)	1.291	1.453
CCSD	1.288	1.449
CCSD(T)	1.234	1.545
CISD	1.295	1.389
SVWN	1.188	1.559
BLYP	1.206	1.632
PBE	1.198	1.606
HCTH	1.176	1.609
B3LYP	1.227	1.523
PBE0	1.233	1.479
Experimental	1.217	1.575

but inclusion of the triply excited configurations in the full MP4(SDTQ) method has a huge effect. The F—O bonds are elongated to the point ($>2.5 \text{ \AA}$) where perturbation theory breaks down since the underlying RHF wave function becomes extremely poor. The MP4(SDTQ) model basically does not predict a stable FOOF molecule. The triples also have a large effect at the CCSD(T) level, but it is clear that the effect is wildly overestimated with the MP4 method. Although the results are not converged with respect to basis set (aug-cc-pVTZ), the remaining changes are of the order of a few thousandths of an angstrom.¹⁷ Even with the sophisticated CCSD(T) model, the geometry errors are thus $\sim 0.03 \text{ \AA}$.

The DFT methods are all well behaved and perform surprisingly well for such a difficult system, with the B3LYP results being comparable to CCSD. The main problem is of course that there is no way of systematically improving the structure, or knowing beforehand whether DFT will be able to give a good description for the specific problem.

11.7.2 The dipole moment of CO

The experimental value for the dipole moment of CO is 0.122 debye, with the polarity C^-O^+ , for a bond length of 1.1281 \AA .¹⁸ Calculated values with the aug-cc-pVXZ basis sets¹⁹ are given in Table 11.21.

Table 11.21 Dipole moment (debye) for CO; the experimental value is 0.122 debye

Method	aug-cc-pVDZ	aug-cc-pVTZ	aug-cc-pVQZ	aug-cc-pV5Z
HF	-0.259	-0.266	-0.265	-0.264
MP2	0.296	0.280	0.275	0.273
MP3	0.076	0.047	0.036	0.032
MP4	0.220	0.222	0.216	0.214
CCSD	0.097	0.070	0.059	0.055
CCSD(T)	0.141	0.127	0.118	0.115
CISD	0.050	0.023	0.011	0.008
LSDA	0.232	0.226	0.229	0.229
BLYP	0.187	0.184	0.185	0.185
PBE	0.229	0.224	0.224	0.224
HCTH	0.194	0.181	0.175	0.179
B3LYP	0.091	0.086	0.087	0.088
PBE0	0.107	0.101	0.102	0.102

The HF level (as usual) overestimates the polarity, in this case leading to an incorrect direction of the dipole moment. The MP perturbation series oscillates, and it is clear that the MP4 result is far from converged. The CCSD(T) method apparently recovers the most important part of the electron correlation, and is very close to the full CCSDT result in an augmented DZP basis.²⁰ However, even with the aug-cc-pV5Z basis sets, there is still a discrepancy of ~ 0.01 debye relative to the experimental value. The DFT methods are not particularly accurate, although for this specific problem the PBE0 method gives a reasonably good result.

11.7.3 The vibrational frequencies of O_3

Ozone is an example of a molecule where the single-reference RHF is quite poor, since there is considerable biradical character in the wave function (as illustrated in Figure 4.9). The harmonic vibrational frequencies derived from experiments are 716, 1089 and 1135 cm^{-1} , where the band at 1089 cm^{-1} corresponds to an asymmetric stretch.²¹ As this nuclear motion changes the relative weights of the ionic and biradical structures, the frequency is very sensitive to the quality of the wave function. Although the wave function is equally poor for all the frequencies, the two other vibrations (symmetric stretch and angle bending) conserve the C_{2v} symmetry, and thus benefit from a significant cancellation of errors. The calculated frequencies at different levels of theory with the cc-pVTZ basis are given in Table 11.22 together with the mean absolute deviation (MAD).

Table 11.22 Harmonic frequencies (cm^{-1}) for O_3 with the cc-pVTZ basis

Method	ν_1	ν_2	ν_3	MAD
HF	867	1418	1537	294
MP2	743	2241	1166	403
MP3	798	1713	1364	312
MP4	695	1592	1107	184
CCSD	762	1266	1278	122
CCSD(T)	716	1054	1153	18
CCSDT ^a	717	1117	1163	19
CCSDT(Q) ^a	709	1112	1133	6
CISD	815	1535	1407	272
[2,2]-CASSCF	799	1497	1189	182
[2,2]-CASPT2 ^b	737	1268	1318	128
[12,9]-CASPT2 ^b	692	1003	1092	51
LSDA	744	1147	1248	66
BLYP	683	980	1129	49
PBE	710	1057	1184	29
HCTH	742	1111	1227	47
B3LYP	746	1193	1251	83
PBE0	777	1295	1322	151
Experimental	716	1089	1135	

^a Data from Kucharski and Bartlett²²

^b Data from Ljubic and Sabljic²³

The simple picture with ozone as a resonance structure between ionic and biradical forms suggests that a two-configuration wave function should be able to give a qualitatively correct description. The [2,2]-CASSCF and [2,2]-CASPT2 results, however, show that dynamical correlation is also very important. The poor RHF reference wave function is clearly seen by the MP_n results, with the MP2 value being in error by a factor of 2 for the asymmetric stretch, and the MP4 result is in error by $\sim 500\text{ cm}^{-1}$ for ν_2 , despite reproducing ν_1 and ν_3 to within 30 cm^{-1} . The coupled cluster methods are less sensitive to the quality of the HF wave function, and are in somewhat better

agreement with the experimental values. The CCSD(T) results are within $\sim 20\text{cm}^{-1}$ of the experimental values, but part of this agreement is accidental as seen by the CCSDT and CCSDT(Q) results, and even the CCSDT(Q) model has errors of $\sim 25\text{cm}^{-1}$. Part of this discrepancy may be due to basis set errors, although the results for the CASPT2 method indicate that larger basis sets will further increase the value of the vibrational frequencies.²³ The DFT methods perform well, yielding results comparable to those at the CCSD or CCSD(T) levels, at a fraction of the computational cost. Even the local density functional gives acceptable results, but this is a case where the hybrid DFT methods (B3LYP and PBE0) perform worse than the pure DFT ones. It can be noted that the cc-pVTZ basis set is sufficiently large that the DFT results are essentially converged, and the results in Table 11.22 thus reflect the intrinsic accuracy of the different DFT methods.

11.8 Relative Energies of C_4H_6 Isomers

The elaborate treatment for the H_2O system is only possible because of its small size. For larger systems, less rigorous methods must be employed. Let us as a more realistic example consider a determination of the relative stability of the C_4H_6 isomers shown in Figure 11.14.

There are experimental values for the first eight structures,²⁴ which allows an evaluation of the performance of different methods. This in turn enables an estimate of how much trust should be put in the predicted values for structures **9**, **10** and **11**.




Figure 11.14 C_4H_6 isomers

An investigation may start by optimizing the geometries by semi-empirical methods, as this will give initial estimates of the energetics and provide reasonable starting geometries for higher level *ab initio* calculations. Relative energies and associated errors relative to the experimental values for different semi-empirical methods are shown in Table 11.23.

Table 11.23 Energies (kJ/mol) relative to **1** calculated by semi-empirical methods

Isomer	MINDO/3	MNDO	AM1	PM3	SAM1	Exp
2	-83	-17	9	-5	-23	36
3	5	9	66	28	58	47
4	-14	19	30	29	8	52
5	-16	30	32	20	15	55
6	8	37	74	56	63	91
7	74	147	202	160	185	107
8	42	103	145	110	118	133
9	84	132	158	122	143	
10	232	349	354	318	255	
11	227	354	378	336	336	
MAD	72	39	31	33	39	

MINDO/3 clearly has severe problems with some of the conjugated systems. The MNDO/AM1/PM3 family performs somewhat better, although none of them can predict the correct ordering. The SAM1 method is not an improvement for this case. The mean absolute deviation (MAD) for the predicted stabilities is ~30 kJ/mol, which is a typical accuracy for semi-empirical methods.

The next step up in terms of theory is *ab initio* HF with increasingly larger basis sets. Table 11.24 shows the results for various basis sets, where the geometries have been optimized with the STO-3G, 3-21G and 6-31G(d,p) basis sets, but the latter used for the pc-*n* basis sets.

Table 11.24 Energies (kJ/mol) relative to **1** calculated at the HF level with various basis sets

Isomer	STO-3G	3-21 G	6-31 G(d,p)	pc-1 ^a	pc-2 ^a	ZPE	Exp
2	-54	15	30	35	31	-3	36
3	-52	75	54	74	66	+4	47
4	36	47	54	59	56	-3	52
5	-22	38	56	62	55	-1	55
6	24	107	85	98	94	0	91
7	48	191	126	148	140	+4	107
8	73	182	136	151	145	-2	133
9	93	198	155	168	163	-3	
10	216	350	304	320	314	-7	
11	188	324	295	317	312	-3	
MAD	67	31	6	15	11		

^a HF/6-31 G(d,p) geometry

The minimum STO-3G basis performs worse than the semi-empirical methods, at a substantially higher computational cost. From experience, it is known that the geometry usually changes little beyond a DZP type basis, and relative energies change little beyond a TZP type basis set. Indeed, the change by increasing the basis set beyond pc-2 is less than 1 kJ/mol, i.e. the pc-2 results reflect the inherent error of the HF model. Note that the 6-31G(d,p) basis set yields smaller errors than the larger pc-2, i.e. with a medium-sized basis set there are some (fortuitous) cancellations of errors from incomplete basis and neglect of electron correlation. The HF method underestimates the stability of some isomers (**3**, **7** and **8**), and the singlet–triplet energy difference between **10** and **11** is qualitatively incorrect. Since **11** has one fewer electron pair than **10**, this stability is reversed once correlation is taken into account. With errors up to ~20 kJ/mol, there is little point in including for example differences in zero point energies (HF/6-31G(d,p) values scaled¹¹ with a factor of 0.92), as these are only a few kJ/mol. Including them in the above data in general only changes the MAD values by ~1 kJ/mol.

The next level up for improving the results would be to include electron correlation, and the MP2 method clearly is an obvious first choice. Correlated calculations require polarization functions, and at least a DZP type basis is mandatory. The results at the MP2/6-31G(d,p) level using the HF/6-31G(d,p) optimized geometries is shown in the first column in Table 11.25. The largest error is now reduced to less than ~25 kJ/mol. Furthermore, as a good fraction of the correlation energy is recovered, the singlet carbene **10** is stabilized relative to **11** by ~35 kJ/mol. The change by further increasing the basis set (e.g. cc-pVQZ) is of the order of 1 kJ/mol, i.e. the last column reflects the inherent error of the MP2 model. The mean error for the first eight systems is approximately the same for the MP2 and HF methods, but the singlet–triplet energy difference for the carbenes **10** and **11** is significantly improved by the MP2 treatment. Including the differences in zero point energies (MP2/6-31G(d,p) values scaled¹¹ with a factor of 0.97) again only lead to changes in the MAD values of 1–2 kJ/mol.

Table 11.25 MP2 energies (kJ/mol) relative to **1**

Isomer	6-31 G(d,p) ^a	6-31 G(d,p) ^b	cc-pVDZ ^b	cc-pVTZ ^b	ZPE	Exp
2	23	20	24	22	–2	36
3	32	33	36	39	+5	47
4	51	51	55	53	–1	52
5	45	42	45	42	–1	55
6	67	69	76	74	+2	91
7	84	85	93	89	+6	107
8	109	109	118	117	0	133
9	125	125	134	132	–1	
10	340	343	342	347	–5	
11	375	376	379	393	0	
MAD	16	16	11	13		

^a HF/6-31 G(d,p) geometry

^b MP2/6-31 G(d,p) geometry

Table 11.26 Energies (kJ/mol) relative to **1** at different levels calculated with the cc-pVDZ basis sets at the MP2/6-31G(d,p) optimized geometry

Isomer	HF	MP2	MP3	MP4	CISD	CCSD	CCSD(T)	G3	Exp
2	38	24	41	33	35	38	40	36	36
3	9	36	42	44	42	45	47	54	47
4	56	55	54	54	54	54	55	50	52
5	63	45	61	54	58	5	60	55	55
6	91	76	81	84	79	83	86	83	91
7	134	93	109	112	108	115	118	117	107
8	144	118	129	129	128	132	133	131	133
9	161	134	144	144	145	146	148	144	
10	304	342	326	331	316	321	328	331	
11	300	379	372	381	348	370	379	395	
MAD	9	11	5	4	4	4	4	4	

Addition of electron correlation beyond MP2 improves the agreement with experiments to ~4kJ/mol, as shown in Table 11.26, and essentially all of the advanced methods provide similar accuracy for these (uncomplicated) systems. The composite G3 method that tries to estimate the QCISD(T)/6-311++G(2df,2p) results by additivity of lower level calculations provides similar results. From the observed accuracy of ~4kJ/mol for structures **2–8**, the energetics of species **9–11** may be assumed to be reliable to the same level of accuracy.

If further refinements are required, several factors must be considered:

- The MP2/6-31G(d,p) geometry should be re-optimized at a better level, for example coupled cluster and/or with a better basis set.
- Larger basis sets should be used with for example the CCSD(T) method.
- Zero point energy corrections should be included, perhaps evaluated with a better method than MP2/6-31G(d,p).

Results from various DFT methods with the pc-2 basis set are shown in Table 11.27. In general they give results of MP2 quality or better.

Table 11.27 Energies relative to **1** calculated at DFT levels with the pc-2 basis set

Isomer	LSDA	BLYP	PBE	HCTH	B3LYP	PBE0	Exp
2	35	41	37	35	38	34	36
3	25	74	45	51	65	36	47
4	40	45	42	42	46	44	52
5	66	68	67	66	65	62	55
6	52	92	67	63	85	62	91
7	66	145	96	85	131	84	107
8	104	141	117	108	136	112	133
9	123	161	136	129	155	131	
10	337	353	336	331	346	326	
11	385	401	373	358	394	361	
MAD	22	14	11	14	10	14	

Calculating the relative energies of a series of hydrocarbons is of course well suited for force field methods, although a comparison of stabilities for isomers containing different number of “functional” groups (CH_3 , CH_2 , etc.) means that only force fields that are able to convert steric energies to heat of formation can be used (Section 2.2.10). Even for these relatively simple compounds, however, there are several “unusual” features for which adequate parameters are lacking. The straight MM2 and MM3 force fields lack parameters for the cyclopropenes **8** and **9**, the methylene-cyclopropane **6** and the allene **4** for MM3. The carbenes **10** and **11** are of course outside the capabilities of force field methods. Table 11.28 compares the performance of the MM2 and MM3 methods, along with MMX, which is a modified MM2 model,²⁵ where parameters have been added to allow calculations on **4**, **6**, **8**, and **9**.

Table 11.28 Energies relative to **1** calculated by force field methods

Isomer	MM2	MM3	MMX	Exp
2	46	44	47	36
3	54	53	48	47
4	57		53	52
5	61	61	62	55
6			91	91
7	112	113	100	107
8			142	133
9			133	
MAD	(7)	(7)	5	

The performance is (as expected) very good. MMX provides relative (and absolute) stabilities with a MAD of only 5 kJ/mol, which are on a par with the results from the sophisticated G3 composite model in Table 11.26. Considering that force field calculations require a factor of $\sim 10^5$ less computer time for these systems than the *ab initio* methods combined in Table 11.26, this clearly stresses that knowledge of the strengths and weakness of different theoretical tools is important in selecting a proper model for answering a given question.

References

1. A. R. Hoy, P. R. Bunker, *J. Mol. Spect.*, **74** (1979), 1.
2. A. G. Csaszar, G. Czako, T. Furtenbacher, J. Tennyson, V. Szalay, S. V. Shirin, N. F. Zobov, O. L. Polyansky, *J. Chem. Phys.*, **122** (2005), 214305.
3. A. Halkier, P. Jørgensen, J. Gauss, T. Helgaker, *Chem. Phys. Lett.*, **274** (1997), 235.
4. A. Lüchow, J. B. Anderson, D. Feller, *J. Chem. Phys.*, **106** (1997), 7706.
5. J. Olsen, P. Jørgensen, H. Koch, A. Balkova, R. J. Bartlett, *J. Chem. Phys.*, **104** (1996), 8007.
6. A. Karton, J. M. L. Martin, *Theor. Chem. Acc.*, **115** (2006), 330.
7. T. Helgaker, W. Klopper, H. Koch, J. Noga, *J. Chem. Phys.*, **106** (1997), 9639.
8. A. K. Wilson, T. H. Dunning Jr, *J. Chem. Phys.*, **106** (1997), 8718.
9. S. A. Clough, Y. Beers, G. P. Klein, L. S. Rothman, *J. Chem. Phys.*, **59** (1973), 2254.
10. W. S. Benedict, N. Gailar, E. K. Plyler, *J. Chem. Phys.*, **24** (1956), 1139.
11. A. P. Scott, L. Radom, *J. Phys. Chem.*, **100** (1996), 16502.

12. J. M. L. Martin, T. J. Lee, *Chem. Phys. Lett.*, **225** (1994), 473.
13. J. A. Pople, P. M. W. Gill, N. C. Handy, *Int. J. Quant. Chem.*, **56** (1995), 303.
14. J. M. Wittbrodt, H. B. Schlegel, *J. Chem. Phys.*, **105** (1996), 6574.
15. P. Jensen, *J. Mol. Spect.*, **133** (1989), 438.
16. R. H. Jackson, *J. Chem. Soc.* (1962) 4585.
17. E. Kraka, Y. He, D. Cremer, *J. Phys. Chem. A*, **105** (2001), 3269.
18. J. S. Muentert, *J. Mol. Spectrosc.*, **55** (1970), 490.
19. K. A. Peterson, T. H. Dunning Jr, *J. Mol. Struct.*, **400** (1997), 93.
20. L. A. Barnes, B. Liu, R. Lindh, *J. Chem. Phys.*, **98** (1993), 3972; G. E. Scuseria, M. D. Miller, F. Jensen, J. Geertsens, *J. Chem. Phys.*, **94** (1991), 6660.
21. A. Barbe, C. Secroun, P. Jouve, *J. Mol. Spectrosc.*, **49** (1974), 171.
22. S. A. Kucharski, R. J. Bartlett, *J. Chem. Phys.*, **110** (1999), 8233.
23. I. Ljubic, A. Sabljic, *Chem. Phys. Lett.*, **385** (2004), 214.
24. S. W. Benzon, F. R. Cruickshank, D. M. Golden, G. R. Haugen, H. E. O'Neal, A. S. Rodgers, R. Shaw, R. Walsh, *Chem. Rev.*, **69** (1969), 279.
25. K. E. Gilbert, PCMODEL, Serena Software.

12 Optimization Techniques

Many problems in computational chemistry can be formulated as an optimization of a multi-dimensional function.¹ Optimization is a general term for finding *stationary* points of a function, i.e. points where the first derivative is zero. In the majority of cases, the desired stationary point is a minimum, i.e. all the second derivatives are positive. In some cases, the desired point is a first-order saddle point, i.e. the second derivative is negative in one, and positive in all other, directions. In a few special cases, a higher order saddle point is desired.

Most optimization methods determine the *nearest* stationary point, but a multi-dimensional function may contain many (in some cases very many!) different station-




Figure 12.1 Illustrating a multi-dimensional energy surface

ary points of the same kind. The minimum with the lowest value is called the *global* minimum, while all the others are *local* minima. Some examples:

- The energy as a function of nuclear coordinates. Both minima and first-order saddle points (transition structures) are of interest. The energy function may be of the force field type, or from solving the electronic Schrödinger equation.
- An error function depending on parameters. Only minima are of interest, and the global minimum is usually (but not always) desired. This may for example be determination of parameters in a force field, a set of atomic charges, or a set of localized molecular orbitals.
- The energy of a wave function containing variational parameters, such as a Hartree–Fock (one Slater determinant) or multi-configurational (many Slater determinants) wave function. Parameters are typically the molecular orbital and configurational state coefficients, but may also be for example basis function exponents. Usually only minima are desired, although in some cases saddle points may also be of interest (excited states).

When the parameters enter the function to be optimized in a *quadratic* fashion, the stationary points can be obtained by solving a set of *linear* equations by standard matrix techniques. In most cases, however, the parameters enter the function in a *non-linear* fashion, which require *iterative* methods for locating the stationary points. The latter can be further divided into methods for locating minima, and methods for locating saddle points. The problem of optimizing quadratic functions is treated in Section 12.1, minima and saddle points of general functions in Sections 12.2 and 12.4, respectively, while methods for global optimization are covered in Section 12.6. Optimization of functions subject to external constraints is dealt with in Section 12.5. Finally, Section 12.8 describes methods for following reactions paths, which may be considered either as solving partial differential equations or as a constrained optimization problem.

12.1 Optimizing Quadratic Functions

Data fitting is a typical example of an optimization problem where the parameters enter the function in a quadratic fashion. Consider for example the problem of determining a set of force field partial charges Q_i by minimizing an error function measuring the fit to the electrostatic potential sampled at a number of points surrounding the molecule (Section 2.2.6).

$$\text{ErrF}(\mathbf{Q}) = \sum_j^{N_{\text{points}}} \left(\phi_{\text{ESP}}(\mathbf{r}_j) - \sum_i^{N_{\text{atoms}}} \frac{Q_i(\mathbf{R}_i)}{|\mathbf{R}_i - \mathbf{r}_j|} \right)^2 \quad (12.1)$$

We will at present ignore the constraint that the sum of the charges must be equal to the total molecular charge; this can be dealt with by the techniques in Section 12.5. The ErrF in eq. (12.1) can be generalized as in eq. (12.2), with y_j being the reference values, a_i the fitting parameters (Q_i) and x_{ij} the coefficients corresponding to the inverse distances $|\mathbf{R}_i - \mathbf{r}_j|$.

$$\text{ErrF}(\mathbf{a}) = \sum_j^M \left(y_j - \sum_i^N a_i x_{ij} \right)^2 \quad (12.2)$$

The best set of parameters is determined by setting all the first derivatives of ErrF to zero.

$$\begin{aligned}
 \frac{\partial \text{ErrF}}{\partial a_1} &= -2 \sum_j^M x_{1j} \left(y_j - \sum_i^N a_i x_{ij} \right) = 0 \\
 \frac{\partial \text{ErrF}}{\partial a_2} &= -2 \sum_j^M x_{2j} \left(y_j - \sum_i^N a_i x_{ij} \right) = 0 \\
 &\vdots \\
 \frac{\partial \text{ErrF}}{\partial a_N} &= -2 \sum_j^M x_{Nj} \left(y_j - \sum_i^N a_i x_{ij} \right) = 0
 \end{aligned} \tag{12.3}$$

By rearrangement, this gives a set of coupled linear equations.

$$\begin{aligned}
 \sum_i^N a_i \sum_j^M x_{1j} x_{ij} &= \sum_j^M x_{1j} y_j \\
 \sum_i^N a_i \sum_j^M x_{2j} x_{ij} &= \sum_j^M x_{2j} y_j \\
 &\vdots \\
 \sum_i^N a_i \sum_j^M x_{Nj} x_{ij} &= \sum_j^M x_{Nj} y_j
 \end{aligned} \tag{12.4}$$

These N equations with N unknowns can also be written in a matrix-vector notation.

$$\mathbf{Xa} = \mathbf{b} \tag{12.5}$$

The formal solution can be obtained by multiplying with the inverse \mathbf{X} matrix.

$$\mathbf{a} = \mathbf{X}^{-1} \mathbf{b} \tag{12.6}$$

In actual applications, the \mathbf{X} matrix may be singular, or nearly so, and the inverse matrix either does not exist or is prone to numerical errors. A singular matrix indicates that at least one of the linear equations can be written as a combination of the other equations, and such cases can be handled by singular value decomposition methods, as discussed in Section 16.2. Indeed, for the example of determining partial charges by fitting to the electrostatic potential, the equations determining the charges on the atoms far from the molecular surface are often poorly conditioned, i.e. the external electrostatic potential is only weakly dependent on the charges on the buried atoms.

Other examples of optimizing functions that depend quadratically of the parameters include the energy of Hartree–Fock (HF) and configuration interaction (CI) wave functions. Minimization of the energy with respect to the MO or CI coefficients leads to a set of linear equations. In the HF case, the x_{ij} coefficients depend on the parameters a_i , and must therefore be solved iteratively. In the CI case, the number of parameters is typically 10^6 – 10^9 , and a direct solution of the linear equations is therefore prohibitive, and special iterative methods are used instead. The use of iterative techniques for solving the CI equations is not due to the mathematical nature of the problem, but due to computational efficiency considerations.

12.2 Optimizing General Functions: Finding Minima

The simple-minded approach for minimizing a function is to step one variable at a time until the function has reached a minimum, and then switch to another variable. This requires only the ability to calculate the function value for a given set of variables. As the variables are not independent, however, several cycles through the whole set are necessary for finding a minimum. This is impractical for more than five–ten variables, and may not work anyway.

The *Simplex* method represents a more efficient approach using only function values for constructing an irregular polyhedron in parameters space, and moving this polyhedron towards the minimum, while allowing the size to contract or expands to improve the convergence.² It is better than the simple-minded “one-variable-at-a-time” approach, but becomes too slow for many-dimensional functions.

Since optimization problems in computational chemistry tend to have many variables, essentially all commonly used methods assume that at least the first derivative of the function with respect to all variables, the gradient \mathbf{g} , can be calculated analytically (i.e. directly, and not as a numerical differentiation by stepping the variables). Some methods also assume that the second derivative matrix, the Hessian \mathbf{H} , can be calculated.

It should be noted that the target function and its derivative(s) are calculated with a finite precision, which depends on the computational implementation. A stationary point can therefore not be located exactly, the gradient can only be reduced to a certain value. Below this value, the numerical inaccuracies due to the finite precision will swamp the “true” functional behaviour. In practice, the optimization is considered converged if the gradient is reduced below a suitable “cutoff” value, or if the function change between two iterations becomes sufficiently small. Both these criteria may in some cases lead to problems, as a function with a very flat surface in a certain region may meet the criteria without containing a stationary point.

There are three classes of commonly used optimization methods for finding minima, each having their advantages and disadvantages.

12.2.1 Steepest descent

The gradient vector \mathbf{g} points in the direction where the function increases most, i.e. the function value can always be lowered by stepping in the opposite direction. In the *Steepest Descent* (SD) method, a series of function evaluations are performed in the negative gradient direction, i.e. along a search direction defined as $\mathbf{d} = -\mathbf{g}$. Once the function starts to increase, an approximate minimum may be determined by interpolation between the calculated points. At this interpolated point, a new gradient is calculated and used for the next line search.

The steepest descent algorithm is sure-fire. If the line minimization is carried out sufficiently accurately, it will *always* lower the function value, and is therefore *guaranteed* to approach a minimum. It has, however, two main problems. Two subsequent line searches are necessarily perpendicular to each other; if there was a gradient component along the previous search direction, the energy could be further lowered in this direction. The steepest descent algorithm therefore has a tendency for each line search to partly spoil the function lowering obtained by the previous search. The steepest




Figure 12.2 Steepest descent minimization

descent path oscillates around the minimum path, as illustrated in Figure 12.2, and this is particularly problematic for surfaces having long narrow valleys.

Furthermore, as the minimum is approached, the rate of convergence slows down. The steepest descent will actually never reach the minimum, it will crawl towards it at an ever decreasing speed.

An accurate line search requires several function evaluations along each search direction. Often the minimization along the line is only carried out fairly crudely, or a single step is simply taken along the negative gradient direction. In the latter case, the step size is varied dynamically during the optimization; if the previous step reduced the function value, the next step is taken with a slightly longer step size, but if the function values increased, the step size is reduced. Without an accurate line search, the guarantee for lowering of the function value is lost, and the optimization may potentially end up in an oscillatory state.

By its nature, the steepest descent method can only locate function minima. The advantage is that the algorithm is very simple, and requires only storage of a gradient vector. It is furthermore one of the few methods that is guaranteed to lower the function value. Its main use is to quickly relax a poor starting point, before some of the more advanced algorithms take over, or as a “backup” algorithm if the more sophisticated methods are unable to lower the function value.

12.2.2 Conjugate gradient methods

The main problem with the steepest descent method is the partial “undoing” of the previous step. The *Conjugate Gradient* (CG) method tries to improve on this by performing each line search not along the current gradient but along a line that is constructed such that it is “conjugate” to the previous search direction(s). If the surface is purely quadratic, the conjugate direction criterion guarantees that each successive minimization will not generate gradient components along any of the previous directions, and the minimum is reached after at most N_{var} steps. The first step is equivalent to a steepest descent step, but subsequent searches are performed along a line formed as a mixture of the current negative gradient and the previous search direction.

$$\mathbf{d}_i = -\mathbf{g}_i + \beta_i \mathbf{d}_{i-1} \quad (12.7)$$

There are several ways of choosing the β value. Some of the names associated with these methods are *Fletcher–Reeves* (FR), *Polak–Ribiere* (PR) and *Hestenes–Stiefel* (HS). Their definitions of β are given in eq. (12.8).

$$\begin{aligned}\beta_i^{\text{FR}} &= \frac{\mathbf{g}_i^t \mathbf{g}_i}{\mathbf{g}_{i-1}^t \mathbf{g}_{i-1}} \\ \beta_i^{\text{PR}} &= \frac{\mathbf{g}_i^t (\mathbf{g}_i - \mathbf{g}_{i-1})}{\mathbf{g}_{i-1}^t \mathbf{g}_{i-1}} \\ \beta_i^{\text{HS}} &= \frac{\mathbf{g}_i^t (\mathbf{g}_i - \mathbf{g}_{i-1})}{\mathbf{d}_{i-1}^t (\mathbf{g}_i - \mathbf{g}_{i-1})}\end{aligned}\tag{12.8}$$

For non-quadratic surfaces, the conjugate property does not hold rigorously and, for real problems, the CG algorithm must often be restarted (i.e. setting $\beta = 0$) during the optimization. The PR formula for β has a tendency of restarting the procedure more gracefully than the other two, and is usually preferred in practice. The conjugate property holds best for near-quadratic surfaces, and the convergence properties of CG methods can be improved by scaling the variables by a suitable pre-conditioner matrix, for example containing (approximate) inverse second derivatives. Conjugate gradient methods have much better convergence characteristics than the steepest descent, but they are again only able to locate minima. They do require slightly more storage than the steepest descent, since two (current gradient and previous search direction) vectors must be stored, but this is rarely a problem.

12.2.3 Newton–Raphson methods

The *Newton–Raphson* (NR) method expands the true function to second order around the current point \mathbf{x}_0 .

$$f(\mathbf{x}) \approx f(\mathbf{x}_0) + \mathbf{g}^t(\mathbf{x} - \mathbf{x}_0) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_0)^t \mathbf{H}(\mathbf{x} - \mathbf{x}_0)\tag{12.9}$$

Requiring the gradient of the second-order approximation in eq. (12.9) to be zero produces the step in eq. (12.10).

$$(\mathbf{x} - \mathbf{x}_0) = -\mathbf{H}^{-1} \mathbf{g}\tag{12.10}$$

In the coordinate system (\mathbf{x}') where the Hessian is diagonal (i.e. performing a unitary transformation, see Section 16.2), the NR step may be written as in eq. (12.11).

$$\begin{aligned}\Delta \mathbf{x}' &= (\Delta x'_1, \Delta x'_2, \Delta x'_3, \dots, \Delta x'_N)^t \\ \Delta x'_i &= -\frac{f_i}{\varepsilon_i}\end{aligned}\tag{12.11}$$

Here f_i is the projection of the gradient along the Hessian eigenvector with eigenvalue ε_i (the gradient component pointing in the direction of the i th eigenvector).

As the real function contains terms beyond second order, the NR formula can be used iteratively for stepping towards a stationary point. Near a minimum, all the Hessian eigenvalues are positive (by definition), and the step direction is opposite to the gradient direction, as it should be. If, however, one of the Hessian eigenvalues is negative, the step in this direction will be *along* the gradient component, and thus

increase the function value. In this case, the optimization may end up at a stationary point with one negative Hessian eigenvalue, a first-order saddle point. In general, the NR method will attempt to converge on the “nearest” stationary point, regardless of whether this is a minimum, saddle point or maximum.

Another problem is the use of the inverse Hessian for determining the step size. If one of the Hessian eigenvalues becomes close to zero, the step size goes toward infinity (except if the corresponding gradient component f_i is exactly zero). The NR step is thus without bound, and it may take the variables far outside the region where the second-order Taylor expansion is valid. The latter region is often described by a “*Trust Radius*”. In some cases, the NR step is taken as a search direction along which the function is minimized, analogously with the steepest descent and conjugate gradient methods. The augmented Hessian methods described below are normally more efficient.

The advantage of the NR method is that the convergence is second order near a stationary point. If the function only contains terms up to second order, the NR step will go to the stationary point in a single step. In general, the function contains higher order terms, but the second-order approximation becomes better and better as the stationary point is approached. Sufficiently close to the stationary point, the gradient is reduced quadratically, i.e. if the gradient norm is reduced by a factor of 2 between two iterations, it will go down by a factor of 4 in the next iteration, and a factor of 16 in the next. The quadratic convergence, however, is often only observed very close to the stationary point, and the NR method typically only displays linear convergence.

Besides the abovementioned problems with step control, there are also other computational aspects that tend to make the straightforward NR problematic for many problem types. The true NR method requires calculation of the full second derivative matrix, which must be stored and inverted (diagonalized). For some types of function, a calculation of the Hessian is computationally demanding. For others, the number of variables is so large that manipulating a matrix the size of the number of variables squared is impossible. The following two sections address some solutions to these problems.

12.2.4 Step control

There are two aspects in step control, one is controlling the total *length* of the step, such that it does not exceed the region in which the second-order Taylor expansion is valid, and the second is making sure that the step *direction* is correct. If the optimization is towards a minimum, the Hessian should have all positive eigenvalues in order for the step to be in the correct direction. If, however, the starting point is in a region where the Hessian has negative eigenvalues, the NR step will take it towards a saddle point or maximum. Both these problems can be solved by introducing a shift parameter λ (compare to eq. (12.11)).

$$\Delta x'_i = -\frac{f_i}{\epsilon_i - \lambda} \quad (12.12)$$

If λ is chosen to be below the lowest Hessian eigenvalue, the denominator is always positive, and the step direction will thus be correct. Furthermore, if λ goes towards

$-\infty$, the step size goes towards zero, i.e. the step size can be made arbitrarily small. Methods that modify the nature of the Hessian matrix by a shift parameter are known by names such as “*augmented Hessian*”, “*level-shifted Newton–Raphson*”, “*norm-extended Hessian*” or “*Eigenvector Following*” (EF), depending on how λ is chosen. We will here mention two popular methods for choosing λ .

The *Rational Function Optimization* (RFO) expands the function in terms of a rational approximation instead of a straight second-order Taylor series (eq. (12.9)).³

$$f(\mathbf{x}) \approx \frac{f(\mathbf{x}_0) + \mathbf{g}^t(\mathbf{x} - \mathbf{x}_0) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_0)^t \mathbf{H}(\mathbf{x} - \mathbf{x}_0)}{1 + \frac{1}{2}(\mathbf{x} - \mathbf{x}_0)^t \mathbf{S}(\mathbf{x} - \mathbf{x}_0)} \quad (12.13)$$

The \mathbf{S} matrix is eventually set equal to a unit matrix which leads to the following equation for λ .

$$\sum_i \frac{f_i^2}{\varepsilon_i - \lambda} = \lambda \quad (12.14)$$

This is a one-dimensional equation in λ which can be solved by standard (iterative) methods. There will in general be one more solution than the number of degrees of freedom, but by choosing the lowest λ solution, it is ensured that the resulting step will be towards a minimum. The RFO step calculated from eq. (12.12) will always be shorter than the pure NR step (eq. (12.11)), but there is no guarantee that it will be within the trust radius. If the RFO step is too long, it may be scaled down by a simple multiplicative factor, however, if the factor is much smaller than 1, it follows that the resulting step may not be the optimum for the given trust radius.

Another way of choosing λ is to require that the step length be equal to the trust radius R , which is in essence the best step on a hypersphere with radius R . This is known as the *Quadratic Approximation* (QA) method.⁴

$$|\Delta \mathbf{x}'|^2 = \sum_i \left(\frac{f_i}{\varepsilon_i - \lambda} \right)^2 = R^2 \quad (12.15)$$

This may again have multiple solutions, but by choosing the lowest λ value, the minimization step is selected. The maximum step size R may be taken as a fixed value, or allowed to change dynamically during the optimization. If the actual energy change between two steps agrees well with that predicted from the second-order Taylor expansion, the trust radius for the next step may be increased, and vice versa.

12.2.5 Obtaining the Hessian

The second problem, the computational aspect of calculating the Hessian, is often encountered in electronic structure calculations. Here the calculation of the second derivative matrix can be an order of magnitude more demanding than calculating the gradient. In such cases, an *updating* scheme may be used instead. The idea is to start off with an approximation to the Hessian, maybe just a unit matrix. The initial step will thus resemble a steepest descent step. As the optimization proceeds, the gradients at the previous and current points are used for making the Hessian a better approximation for the actual system. After two steps, the updated Hessian is a rather

good approximation to the exact Hessian in the direction defined by these two points (but not in the other directions). There are many such updating schemes, some of the commonly used are associated with the names *Davidon–Fletcher–Powell* (DFP), *Broyden–Fletcher–Goldfarb–Shanno* (BFGS) and *Powell*. For minimizations, the BFGS update eq. (12.16) is usually preferred, as it tends to keep the Hessian positive definite.

$$\mathbf{H}_n = \mathbf{H}_{n-1} + \Delta\mathbf{H}$$

$$\Delta\mathbf{H}_{\text{BFGS}} = \frac{\Delta\mathbf{g}\Delta\mathbf{g}^t}{\Delta\mathbf{g}^t\Delta\mathbf{x}} - \frac{\mathbf{H}\Delta\mathbf{x}\Delta\mathbf{x}^t\mathbf{H}}{\Delta\mathbf{x}^t\mathbf{H}\Delta\mathbf{x}} \quad (12.16)$$

For saddle point searches, the updating must allow the Hessian to develop negative eigenvalues, and the Powell or updates based on combining several methods are usually employed.⁵

The use of approximate Hessians within the NR method is known as *pseudo-Newton–Raphson* or *variable metric* methods. It is clear that they do not converge as fast as true NR methods, where the exact Hessian is calculated in each step, but if for example five steps can be taken for the same computational cost as one true NR step, the overall computational effort may be less. True NR methods converge quadratically near a stationary point, while pseudo-NR methods display a linear convergence. Far from a stationary point, however, the true NR method will typically also only display linear convergence.

Pseudo-NR methods are usually the best choice in geometry optimizations using an energy function calculated by electronic structure methods. The quality of the initial Hessian of course affects the convergence when an updating scheme is used. The use of an exact Hessian at the first point often gives a good convergence, however, this may not be the most cost-efficient strategy. In many cases, a quite reasonable Hessian for a minimum search may be generated by simple rules connecting for example bond lengths and force constants.⁶ Alternatively, the initial Hessian may be taken from a calculation at a lower level of theory. As an initial exploration of an energy surface is often carried out at a low level of theory, followed by frequency calculations to establishing the nature of the stationary points, the resulting force constants can be used for starting an optimization at higher levels. This is especially useful for transition structure searches that require a quite accurate Hessian. The success of this strategy relies on the fact that the qualitative structure of an energy surface is often fairly insensitive to the level of theory, although there certainly are many examples where this is not the case.

12.2.6 Storing and diagonalizing the Hessian

The last potential problem of all NR-based methods is the storage and handling of the Hessian matrix. For methods where the calculation of the Hessian is easy but the number of variables is large, this may be a problem. A prime example here is geometry optimization using a force field energy function. The computational effort for calculating the Hessian goes up roughly as the *square* of the number of atoms. Diagonalization of the Hessian matrix required for the NR optimization, however, depends on the *cube* of the matrix size, i.e. it goes up as the *cube* of the number of atoms. Since matrix diagonalization becomes a significant factor around a size of

1000 × 1000, it is clear that NR methods should not be used for force field optimizations beyond a few hundred atoms. For large systems the computational effort for predicting the geometry step will completely overwhelm the calculation of the energy, gradient and Hessian. The conjugate gradient method avoids handling of the Hessian and only requires storage of two vectors, and it is therefore usually the method of choice for force field optimizations.

For large systems many of the off-diagonal elements in the Hessian are very small, essentially zero (the coupling between distant atoms is very small), and the Hessian for large systems is therefore a *sparse matrix*. NR methods that take advantage of this fact by neglecting off-diagonal blocks, are denoted *truncated NR*. Some force field programs use an extreme example of this where only the 3 × 3 submatrices along the diagonal are retained. These 3 × 3 matrices contain the coupling elements between the *x*, *y* and *z* coordinates for a single atom. The task of inverting say a 3000 × 3000 matrix is thus replaced by inverting 1000 3 × 3 matrices, reducing the computational cost for the diagonalization by a factor of 10⁶. If the NR step is not taken directly, but rather used as a direction along which the function is minimized, truncated NR methods start to resemble the conjugate gradient method, although it is somewhat more complicated to implement.

12.2.7 Extrapolations: the GDIIS method

Newton–Raphson methods can be combined with extrapolation procedures, and the best known of these is perhaps the *Geometry Direct Inversion in the Iterative Subspace* (GDIIS),⁷ which is directly analogous to the DIIS for electronic wave functions described in Section 3.8.1. In the GDIIS method, the NR step is not taken from the last geometry but from an interpolated point with a corresponding interpolated gradient based on the previously calculated points on the surface.

$$\begin{aligned}\mathbf{x}_n^* &= \sum_i^n c_i \mathbf{x}_i \\ \mathbf{g}_n^* &= \sum_i^n c_i \mathbf{g}_i\end{aligned}\tag{12.17}$$

The interpolated geometry and gradient are generated by requiring that the norm of an error vector is minimum, subject to a normalization condition.

$$\begin{aligned}\text{ErrF}(\mathbf{e}) &= \left| \sum_i^n c_i \mathbf{e}_i \right| \\ \sum_i^n c_i &= 1\end{aligned}\tag{12.18}$$

There are two common choices for the error vector, either a “geometry” or “gradient” vector, with the latter being preferred in more recent work.⁸

$$\mathbf{e}_i = \mathbf{H}_n^{-1} \mathbf{g}_i \quad \text{or} \quad \mathbf{e}_i = \mathbf{g}_i\tag{12.19}$$

The DIIS approach attempts to find a low-gradient point within the subspace already searched. For optimizing electronic wave functions, there is usually only one minimum

and the DIIS extrapolation significantly improve both the convergence rate and stability. For geometry optimizations, however, the target function is usually complicated and contains many minima and saddle points, making DIIS extrapolations much less useful or even disadvantageous. It is not uncommon for an optimization to move across a flat part of the surface before entering the local minimum region. This will result in the gradient being small for several steps, and then increasing as the minimum is approached. In such cases, the DIIS procedure will attempt to pull the structure back to the flat energy region, since this is where the gradient is small, and DIIS will in such cases be counterproductive.

12.3 Choice of Coordinates

Naively one may think that any set of coordinates that uniquely describes the function is equally good for optimization. This is not the case! A “good” set of coordinates may transform a divergent optimization into a convergent one, or increase the rate of convergence. We will look specifically at the problem of optimizing a geometry given an energy function depending on nuclear coordinates, but the same considerations hold equally well for other types of optimization. We will furthermore use the straight Newton–Raphson formula (12.10) to illustrate the concepts.

Given the first and second derivatives, the NR formula calculates the geometry step as the inverse of the Hessian times the gradient.

$$(\mathbf{x} - \mathbf{x}_0) = -\mathbf{H}^{-1} \mathbf{g} \quad (12.20)$$

In the coordinate system (\mathbf{x}') where the Hessian is diagonal, the step may be written as in eq. (12.21).

$$\begin{aligned} \Delta \mathbf{x}' &= (\Delta x'_1, \Delta x'_2, \Delta x'_3, \dots, \Delta x'_N)^t \\ \Delta x'_i &= -\frac{f_i}{\epsilon_i} \end{aligned} \quad (12.21)$$

Essentially all computational programs calculate the fundamental properties, the energy and derivatives, in Cartesian coordinates. The Cartesian Hessian matrix has the dimension $3N_{\text{atom}} \times 3N_{\text{atom}}$. Of these, three describe the overall translation of the molecule, and three describe the overall rotation. In the molecular coordinate system, there are only $3N_{\text{atom}} - 6$ coordinates needed for uniquely describing the nuclear positions. Moving all the atoms in say the x -direction by the same amount does not change the energy, and the corresponding gradient component (and all higher derivatives) is zero. The Hessian matrix should therefore have six eigenvalues identical to zero, and the corresponding gradient components, f_i , should also be identical zero. In actual calculations, however, these values are certainly small, but not exactly zero. Numerical inaccuracies may introduce errors of perhaps 10^{-14} – 10^{-16} , and this can have rather drastic consequences. Consider for example a case where the gradient in the x -translation direction is calculated to be 10^{-14} , while the corresponding Hessian eigenvalue is 10^{-16} , leading to an NR step in this direction of 100! This illustrates that care should be taken if redundant coordinates (i.e. more than are necessary for uniquely describing the system) are used in the optimization. In the case of Cartesian geometry optimization, the six translational and rotational degrees of freedom can be removed by

projecting these components out of the Hessian prior to formation of the NR step (Section 16.4). The calculated “steps” in the zero eigenvalue directions are then simply neglected.

Another way of removing the six translational and rotational degrees of freedom is to use a set of internal coordinates. For a simple acyclic system, these may be chosen as $N_{\text{atom}} - 1$ distances, $N_{\text{atom}} - 2$ angles and $N_{\text{atom}} - 3$ torsional angles, as illustrated in the construction of Z -matrices in Appendix D. In internal coordinates the six translational and rotational modes are automatically removed (since only $3N_{\text{atom}} - 6$ coordinates are defined), and the NR step can be formed straightforwardly. For cyclic systems, a choice of $3N_{\text{atom}} - 6$ internal variables that span the whole optimization space may be somewhat more problematic, especially if symmetry is present.

Diagonalization of the Hessian is an example of a linear transformation; the eigenvectors are just linear combinations of the original coordinates. A linear transformation does not change the convergence/divergence properties, or the rate of convergence. We can form the NR step directly in Cartesian coordinates by inverting the Hessian and multiplying it with the gradient vector (eq. (12.20)), or we can transform the coordinates to a system where the Hessian is diagonal, form the ratios $-f_i/\epsilon_i$ (eq. (12.21)) and back-transform to the original system. Both methods generate the exact same NR step (except for rounding-off errors). Since we need to give consideration to the six translational and rotational modes, however, the diagonal representation is advantageous.

The transformation from a set of Cartesian coordinates to a set of internal coordinates, which may for example be distances, angles and torsional angles, is an example of a non-linear transformation. The internal coordinates are connected with the Cartesian coordinates by means of square root and trigonometric functions, not simple linear combinations. A non-linear transformation will affect the convergence properties. This can be illustrate by considering a minimization of a Morse type function (eq. (2.5)) with $D = \alpha = 1$ and $x = \Delta R$.

$$E_{\text{Morse}}(x) = [1 - e^{-x}]^2 \quad (12.22)$$

We will consider two other variables obtained by a non-linear transformation: $y = e^{-x}$ and $z = e^x$. The minimum energy is at $x = 0$, corresponding to $y = z = 1$. Consider an NR optimization starting at $x = -0.5$, corresponding to $y = 1.6587$ and $z = 0.6065$. Table 12.1 shows that the NR procedure in the x -variable requires four iterations before x is less than 10^{-4} . In the y -variable the optimization only requires one step to reach the

Table 12.1 Convergence for different choices of variables

Iteration	x	y	z
0	-0.5000	1.6487	0.6065
1	-0.2176	1.0000	0.7401
2	-0.0541		0.8667
3	-0.0041		0.9570
4	0.0000		0.9951
5			0.9999
6			1.0000


Table 12.2 Convergence for different choices of variables

Iteration	$x_{\text{start}} = 0.30$			$x_{\text{start}} = 1.00$		
	x	y	z	x	y	z
0	0.3000	0.7408	1.3499	1.0000	0.3679	2.7183
1	-0.2381	1.0000	-0.2229	3.3922	1.0000	4.6352
2	-0.0633		-0.3020	4.4283		7.3225
3	-0.0055		-0.4110	5.4405		11.2981
4	0.0000		-0.5628	6.4449		17.2354

$y = 1$ minimum exactly! The optimization in the z -variable takes six iterations before the value is within 10^{-4} of the minimum.

Consider now the same system starting from $x = 0.30$ ($y = 0.7408$ and $z = 1.3499$) and $x = 1.00$ ($y = 0.3679$ and $z = 2.7183$). The first optimization step in the x -variable for the first case overshoots the minimum but then converges in three additional steps. With the z -variable the first step results in an “non-physical” negative value, and subsequent steps do not recover. With the second set of starting conditions, both the x - and z -variable optimizations diverge toward the $x = \infty$ limit. In both cases the y -variable optimization converges (exactly) in one step.

The reason for this behaviour is seen when plotting the three functional forms as shown in Figures 12.3.

**Figure 12.3** Morse curves as a function of x , y and z

The horizontal axis covers the same range of x -variables for all three figures. In the x -variable space the second derivative is negative beyond $x = \ln 2$ ($= 0.69$), and if the optimization is started at larger x -values, the optimization is no longer a minimization, but a maximization toward the $x = \infty$ asymptote. The function in the y -variable is a parabola, and the second-order expansion of the NR method is *exact*. *All* starting points consequently converge to the minimum in a single step. The transformation to the z -variable introduces a singularity at $z = 0$, and it can be seen from Figure 12.3 that the curve shape is much less quadratic than the original function. Using y as a variable is an example of a “good” non-linear transformation, while z is an example of a “poor” non-linear transformation.

These examples show that non-linear transformations may strongly affect the convergence properties of an optimization. The more “harmonic” the energy function is,

the faster the convergence. One should therefore try to choose a set of coordinates where the third and higher order derivatives are as small as possible. Cartesian coordinates are not particularly good in this respect but have the advantage that convergence properties are fairly uniform for different systems. A “good” set of internal coordinates may speed up the convergence but a “bad” set of coordinates may slow it down or cause divergence. For acyclic systems the abovementioned internal coordinates consisting of $N_{\text{atom}} - 1$ distances, $N_{\text{atom}} - 2$ angles and $N_{\text{atom}} - 3$ torsional angles are normally better than Cartesian coordinates. Cyclic systems, however, are notoriously difficult to choose a good set of internal coordinates for. Cyclopropane, for example, has three C–C bonds and three CCC angles, but only three independent variables (not counting the hydrogens). Choosing two distances and one angle introduces a strong coupling between the angle and distances due to the “remote” C–C bond, which is described indirectly by the other three variables. Cartesian coordinates may display better convergence characteristics in such systems.

Another problem is when very soft modes are present. A prototypical example is rotation of a methyl or hydroxy group. Near the minimum the energy changes very little as a function of the torsional angle, i.e. the corresponding Hessian eigenvalue is small. Consequently, even a small gradient may produce a large change in geometry. The potential is not very harmonic, and the result is that the optimization spends many iterations flopping from side to side. A similar problem is encountered in optimization of molecular clusters where the optimum structure is governed by weak van der Waals type interactions.

The problem of an “optimum” choice of coordinates has been addressed by Pulay and coworkers, who suggested using *Natural Internal Coordinates*.⁹ The atoms are first classified into three types: “terminal” (having only one bond), “ring” (part of a ring) or “internal”. All distances between bonded atoms are used as variables. The ring and internal atoms are assumed to have a local symmetry depending on the number of terminal atoms attached to them, i.e. an internal atom with three terminal bonds has local C_{3v} symmetry, one with two terminal bonds has C_{2v} , a ring has local D_{nh} symmetry, etc. Suitable linear combinations of bending and torsional angles are then formed such that the coupling between these coordinates is exactly zero if the local symmetry is the exact symmetry. This will usually not be the case, but the local symmetry coordinates tend to minimize the coupling, and thus the magnitude of third and higher derivatives, thereby improving the NR performance. Natural internal coordinates appear to be a good choice for optimization to minima on an energy surface, since the bonding pattern is usually well defined for stable molecules. For locating transition structures, however, it is much less clear whether natural internal coordinates offer any special advantage. The bonding pattern is not as well defined for TS’s, and a “good” set of coordinates at the starting geometry may become ill behaved during the optimization. For loosely bound complexes, it has been suggested that coordinates depending on the inverse distance or scaled by an inverse reference distance should improve the convergence.¹⁰

In the original formulation, a set of $3N_{\text{atom}} - 6$ independent natural internal coordinates was chosen. It was later discovered that the same optimization characteristics could be obtained by using all distances and bending and torsional angles between atoms within bonding distance as variables.¹¹ Such a set of coordinates will in general be redundant (i.e. the number of coordinates is larger than $3N_{\text{atom}} - 6$), and special care must be taken to handle this. More recently, it has been argued that an “optimum” set

of non-redundant coordinates may be extracted from a large set of (redundant) internal coordinates by selecting the eigenvectors corresponding to non-zero eigenvalues of the square of the matrix defining the transformation from Cartesian to internal coordinates. These linear combinations have been denoted *delocalized internal coordinates*, and are in a sense a generalization of the natural internal coordinates.¹² A major advantage is that delocalized internal coordinates can be generated automatically without any user involvement.

In summary, the efficiency of Newton–Raphson-based optimizations depends on the following factors:

- (1) Hessian quality (exact or updated).
- (2) Step control (augmented Hessian, choice of shift parameter(s)).
- (3) Coordinates (Cartesian, internal).
- (4) Trust radius update (maximum step size allowed).

A comparison of various combinations of these can be found in reference 13.

12.4 Optimizing General Functions: Finding Saddle Points (Transition Structures)

Locating minima for functions is fairly easy. If everything else fails, the steepest descent method is guaranteed to lower the function value. Finding first-order saddle points, *transition structures* (TS), is much more difficult. There are *no* general methods that are guaranteed to work! Many different strategies have been proposed, the majority of which can be divided into two general categories, those based on interpolation between two minima,¹⁴ and those using only local information.¹⁵ Interpolation methods assume that the reactant and product geometries are known, and that a TS is located somewhere “between” these two end-points. It should be noted that many of the methods in this group do not actually locate the TS, they only locate a point close to it. Local methods propagate the geometry using only information about the function and its first and possibly also second derivatives at the current point, i.e. they require no knowledge of the reactant and/or product geometries. Local methods usually require a good estimate of the TS in order to converge. Once the TS has been found, the whole reaction path may be located by tracing the intrinsic reaction coordinate (Section 12.8), which corresponds to a steepest descent path in mass-weighted coordinates, from the TS to the reactant and product.

12.4.1 One-structure interpolation methods: coordinate driving, linear and quadratic synchronous transit, and sphere optimization

The intuitively simple approach for locating a TS is to select one or a few internal “reaction” coordinates, i.e. those that describes the main difference between the reactant and product structures. A typical example is a torsional angle for describing a conformational TS, or two bond distances for a bond breaking/forming reaction. The selected coordinate(s) is (are) fixed at certain values, while the remaining variables are optimized, thereby adiabatically mapping the energy as a function of the reaction variable(s), and such methods are often called “*coordinate driving*”. The goal is to find a

geometry where the residual gradients for the fixed variables are “sufficiently” small. The success of this method depends on the ability to choose a good set of reaction variables, with a good choice being equated with large coefficients for the selected variables in the actual reaction coordinate vector at the TS (as given by the Hessian eigenvector with a negative eigenvalue). The reaction coordinate at the TS, however, is only known after the TS has actually been found, making the choice strongly user-biased, and impossible to verify *a priori*.

If only one or two variables change significantly between reactant and product, the coordinate driving usually works well, and the constrained optimized geometry with the smallest residual gradient is a good approximation to the TS. Some typical examples are rotation of a methyl group (reaction variable is the torsional angle), the HNC to HCN rearrangement (reaction variable is the HCN angle) and S_N2 reactions of the type $X + CH_3Y \rightarrow XCH_3 + Y$ (reaction variables are the XC and CY distances). Good approximations to many conformational TS's can be generated by “driving” a selected torsional angle, and this is often the basis for conformational analysis using force field energy functions. It should be stressed that the highest energy structure located in this fashion is *not* exactly the TS, but it is usually a very good approximation to it. A mapping with more than two reaction variables becomes cumbersome, and rarely leads anywhere.

If a bad choice of reaction variables has been made, “hysteresis” is often observed. This is the term used when the optimization suddenly changes the geometry drastically for a small change in the fixed variable(s). Furthermore, a series of optimizations made by *increasing* the fixed variable(s) to a given value may produce a different result than when *decreasing* the fixed variable(s) to the same point. This indicates that the chosen reaction variable(s) do not contribute strongly to the actual reaction coordinate at the TS. Some TS's have reaction vectors that are not dominated by a few internal variables, and such TS's are difficult to find by constrained optimization methods. In some cases, another set of (internal) coordinates may alleviate the problem, but finding these is part of the “black magic” involved in locating TS's.

The *Linear Synchronous Transit* (LST) method may be considered as a coordinate driving method where all (Cartesian or internal) coordinates are varied linearly between the reactant and product, and no optimization is performed.¹⁶ The assumption is that all variables change at the same rate along the reaction path, and the TS estimate is simply the highest energy structure along the interpolation line. The assumed synchronous change for all variables is rarely a good approximation, and only for simple systems does LST lead to a reasonable estimate of the TS. The *Quadratic Synchronous Transit* (QST) approximates the reaction path by a parabola instead of a straight line. After the maximum on the LST is found, the QST is generated by minimizing the energy in the directions perpendicular to the LST path, and the QST path may then be searched for an energy maximum. These methods are illustrated in Figure 12.4, where the *Intrinsic Reaction Coordinate* (IRC) represents the “true” reaction coordinate.

Bell and Crighton refined the method by performing the minimization from the LST maximum in the directions *conjugate* to the LST instead of the *orthogonal* directions as in the original formulation.¹⁷ A more recent variation of QST, called *Synchronous Transit-guided Quasi-Newton* (STQN), uses a circle arc instead of a parabola for the interpolation, and uses the tangent to the circle for guiding the search towards the TS




Figure 12.4 Illustration of the linear and quadratic synchronous transit methods; energy maxima and minima are denoted by * and •, respectively

region.¹⁸ Once the TS region is located, the optimization is switched to a quasi-Newton–Raphson (Section 12.4.6).

The *Sphere* optimization technique involves a sequence of constrained optimizations on hyperspheres with increasingly larger radii, using the reactant (or product) geometry as a constant expansion point.¹⁹ The lowest energy point on each successive hypersphere thus traces out a low energy path on the energy surface, as illustrated in Figure 12.5. The sphere method may be considered as a coordinate driving algorithm where the driving coordinate is the distance to the minimum. Ohno and Maeda have suggested a variation where the optimization is done in vibrational normal coordinates scaled by the square root of the corresponding Hessian eigenvalues.²⁰ This makes all directions equivalent in an energetic sense, and potentially allows more saddle points to be found, but at the expense of searching the full variable space rather than just the low-energy region. They have suggested that an exhaustive search along all the normal mode directions can potentially find all the TS's connected with a given minimum. Tracing the IRC from all these TS's will lead to other minima, which then can be subjected to a TS search, thereby potentially tracing out all possible reaction paths for a given system.




Figure 12.5 Illustration of the sphere method; energy minima on the hyperspheres are denoted by •, while R indicates a (local) minimum in the full variable space

Barkema and Mousseau have suggested a closely related dynamical version where the gradient at a given point is split into two components parallel and perpendicular to the vector from the minimum to the current point.²¹ The gradient component in the perpendicular directions is followed in the downhill direction, while the structure is advanced in the uphill direction along the parallel component.

It should be noted that the success or failure of LST/QST, as with all optimizations, depends on the coordinates used in the interpolation. Consider for example the HNC to HCN rearrangement. In Cartesian coordinates, the LST path preserves the linearity of the reactant and product, and thus predicts that the hydrogen moves *through* the nitrogen and carbon atoms. In internal coordinates, however, the angle changes from 0° to 180°, and the LST will in this case locate a much more reasonable point with the hydrogen moving around the C–N moiety.




Figure 12.6 LST path in Cartesian and internal coordinates

For large complex systems, the LST path, even in internal coordinates, may involve geometries where two or more atoms clash and it may be difficult or impossible to obtain a function value, for example due to an iterative (SCF) procedure failing to converge.

12.4.2 Two-structure interpolation methods: saddle, line-then-plane, ridge and step-and-slide optimizations

The methods in Section 12.4.1 all optimize one geometrical structure, and differ primarily in how they parameterize the reaction path. The methods in this section operate with two geometrical structures, which attempt to bracket the saddle point and gradually converge on the TS from the reactant and product sides.

In the *Saddle* algorithm,²² the lowest of the reactant and product minima is first identified. A trial structure is generated by displacing the geometry of the lower energy species a fraction (for example 0.05) towards the high energy minimum. The trial structure is then optimized, subject to the constraint that the distance to the high-energy minimum is constant. The lowest energy structure on the hypersphere becomes the new interpolation end-point, and the procedure is repeated. The two geometries will (hopefully) gradually converge on a low-energy structure intermediate between the original two minima, as illustrated in Figure 12.7.




Figure 12.7 Illustration of the saddle method; energy minima on the hyperspheres are denoted by •

A related idea is used in the *Line-Then-Plane* (LTP) algorithm,²³ where the constrained optimization is done in the hyperplane perpendicular to the interpolation line between the two end-points, rather than on a hypersphere.

The *Ridge* method initially locates the energy maximum along the LST path connecting the reactant and product, and defines two points on either side of the energy maximum.²⁴ These points are allowed to relax in the downhill direction a given distance, and a new energy maximum is located along the interpolation line connecting the two relaxed points, and the cycle is repeated. As the saddle point is approached, the two ridge points gradually contract on the actual TS. This method requires a careful adjustment of the magnitude of the “side” and “downhill” steps as the optimization proceeds.

The *Step-and-Slide* algorithm²⁵ is a variation where the reactant and product structures are stepped along the LST line until they have energies equal to a preset value. Both structures are then optimized with respect to minimizing the distance between them, subject to being on an isoenergetic contour surface. The energy is increased, followed by another *step-and-slide* optimization, and this sequence is continued until the distance between the two structures decreases to zero, i.e. converging on the saddle point.

12.4.3 Multi-structure interpolation methods: chain, locally updated planes, self-penalty walk, conjugate peak refinement and nudged elastic band

The methods in this section operate with multiple (more than two) structures or *images* connecting the reactant and product, and are often called *chain-of-state* methods. Relaxation of the images will in favourable cases not only lead to the saddle point, but also to an approximation of the whole reaction path. The initial distribution of structures will typically be along a straight line connecting the reactant and product (LST), but may also involve one or more intermediate geometries to guide the search in a certain direction.

The *Self-Penalty Walk* (SPW) method approximates the reaction path by minimizing the average energy along the path, given as a line integral between the reactant and product geometries (R and P).²⁶




Figure 12.8 Illustration of the SPW method; optimized path points are denoted by x

$$S(\mathbf{R}, \mathbf{P}) = \frac{1}{L} \int_{\mathbf{R}}^{\mathbf{P}} E(\mathbf{x}) d\mathbf{l}(\mathbf{x}) \quad (12.23)$$

The line element $d\mathbf{l}(\mathbf{x})$ belongs to the reaction path, which has a total length of L . In practice, the line integral is approximated as a finite sum of M points, where M typically is of the order of 10–20.

$$S(\mathbf{R}, \mathbf{P}) \approx \frac{1}{L} \sum_{i=1}^M E(\mathbf{x}_i) \Delta l_i \quad (12.24)$$

In order to avoid all points aggregating near the minima (reactant and product), constraints are imposed for keeping the distance between two neighbouring points close to the average distance. Furthermore, repulsion terms between all points are also added to keep the reaction path from forming loops. The resulting target function $T_{\text{SPW}}(\mathbf{R}, \mathbf{P})$ may then be minimized by for example a conjugate gradient method.

$$T_{\text{SPW}}(\mathbf{R}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M, \mathbf{P}) = \frac{1}{L} \sum_{i=1}^M E(\mathbf{x}_i) \Delta l_i + \gamma \sum_{i=0}^M (d_{i,i+1} - \bar{d})^2 + \rho \sum_{i>j+1}^{M+1} \exp\left(-\frac{d_{ij}}{\lambda \bar{d}}\right) \quad (12.25)$$

$$d_{ij} = |\mathbf{x}_i - \mathbf{x}_j|$$

$$\bar{d} = \sqrt{\frac{1}{M+1} \sum_{i=0}^M d_{i,i+1}^2}$$

The γ , λ and ρ parameters are suitable constants for weighting the distance and repulsion constraints relative to the average path energy. In the original version of SPW, the TS is estimated as the grid point with the highest energy after minimization of the target function, but Ayala and Schlegel have implemented a version where one of the points is optimized directly to the TS and the remaining points form an approximation to the IRC path.²⁷

The *Chain* method initially calculates the energy at a series of points placed at regular intervals (spacing of d_{max}) along a suitable reaction coordinate.²⁸ The highest energy point is allowed to relax by a maximum step size of d_{max} along a direction defined by the gradient component orthogonal to the line between by the two neighbouring points. This process is repeated with the new highest energy point until the gradient becomes tangential to the path (within a specified threshold). When this happens, the current highest energy point cannot be further relaxed, and is instead moved to a maximum *along* the path.




Figure 12.9 Illustration of the chain method; initial points along the path are denoted by x , and relaxed points are denoted by \bullet

During the relaxation the chain may form loops, in which case intermediate point(s) is (are) discarded. Similarly, it may be necessary to add points to keep the distance between neighbours below d_{\max} .

The *Locally Updated Planes* (LUP) minimization is related to the chain method, where the relaxation is done in the hyperplane perpendicular to the reaction coordinate, rather than along a line defined by the gradient.²⁹ Furthermore, all the points are moved in each iteration, rather than one at a time.

The *Conjugate Peak Refinement* (CPR) method may be considered as a dynamical version of the chain method, where points are added or removed based on a sequence of maximizations along line segments and minimizations along the conjugate directions.³⁰ The first cycle is analogous to the Bell and Crighton version of the QST: location of an energy maximum along a line between the reactant and product, followed by a sequential minimization in the conjugate directions. The corresponding point becomes a new path point, and an attempt is made to locate an LST maximum between the reactant and midpoint, and between the midpoint and product. If such a maximum is found, it is followed by a new conjugate minimization, which then defines a new intermediate point, etc. The advantage over the chain and LUP methods is that points tend to be distributed in the important region near the TS, rather than uniformly over the whole reaction path.

In practice, it may not be possible to minimize the energy in *all* the conjugate directions, since the energy surface in general is not quadratic. Once the gradient component along the LST path between two neighbouring points exceeds a suitable tolerance during the sequential line minimizations, the optimization is terminated and the geometry becomes a new interpolation point. It may also happen that one of the interpolation points has the highest energy along the path without being sufficiently close to a TS (as measured by the magnitude of the gradient), in which case the point is removed and a new interpolation is performed.

The *Nudged Elastic Band* (NEB) method defines a target function (“elastic band”) as the sum of energies of all images and adds a penalty term having the purpose of distributing the points along the path.³¹ A single spring constant k will distribute the images evenly along the path, but it may also be taken to depend on the energy in order to provide a better sampling near the saddle point.

$$T_{\text{NEB}}(\mathbf{R}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M, \mathbf{P}) = \sum_{i=1}^M E(\mathbf{x}_i) + \sum_{i=1}^{M-1} \frac{1}{2} k (\mathbf{x}_{i+1} - \mathbf{x}_i)^2 \quad (12.26)$$

A straightforward minimization of T_{NEB} gives a reaction path that has a tendency of cutting corners if the spring constant k is too large, and a problem of points sliding down towards the minima if the spring constant is too small. These problems can of course be solved by employing a large number of points, but that would render the optimization inefficient. The “corner-cutting” and “down-sliding” problems for a manageable number of points can be alleviated by “nudging” the elastic band, i.e. using only the component of the spring force *parallel* to the tangent of the path, and only the *perpendicular* component of the energy force in the optimization of T_{NEB} . Since the reaction path is represented by a discrete set of points, the tangent to the path at a given point must be estimated from the neighbouring points, and different definitions influence the optimization efficiency. Furthermore, the different projection of the two force parts and the fact that the projection direction is different for each point mean that there is not a well-defined target function to minimize, and implementation of for example conjugate gradient or Newton–Raphson optimization schemes is not straightforward. The minimization can instead be done using a Newtonian dynamics method (e.g. velocity Verlet, Section 14.2.1), where the velocity is quenched regularly, which effectively corresponds to a steepest descent algorithm with a dynamical step size. As discussed in Section 12.2.1, this is a rather inefficient optimization method, and minimization of T_{NEB} therefore often requires a large number of iterations. The magnitude of the spring constant also influences the optimization efficiency; a small value causes an erratic coverage of the reaction path, while a large value focuses the effort on distributing the points rather than on finding the reaction path, and consequently slows the convergence down.

Despite the optimization issue, the NEB method is one of the most popular interpolation methods, and several improvements and variations have been reported. In the *Climbing Image* (CI-NEB) version, one of the images is allowed to move along the elastic band to become the exact saddle point.³² The *String Method* (SM) attempts to solve the optimization problem by re-distributing the points after each optimization cycle, thereby dispensing with the requirement of using the projected spring force in the optimization.³³ An adaptive version of NEB has also been proposed, which gradually increases the number of images and concentrates the points near the important saddle point region.³⁴ The problem of generating a suitable initial path when the LST path is unsuitable has been addressed by gradually adding images from the reactant and product sides.³⁵

12.4.4 Characteristics of interpolation methods

Interpolation methods have the following characteristics:

- (1) There may not be a TS connecting two minima directly. The algorithm may then find an intermediate geometry having a gradient substantially different from zero, i.e. no nearby stationary point. This is primarily a problem for the one-structure methods in Section 12.4.1.
- (2) The TS found is not necessarily one that connects the two minima used in the interpolation. A calculation of the IRC path may reveal that it is a TS for a different reaction. This is primarily a problem for the one- and two-structure methods in sections 12.4.1 and 12.4.2.

- (3) There may be several TS's (and therefore at least one minimum) between the two selected end-points. Some algorithms may find one of these, and the two connecting minima can then be found by tracing the IRC path, or all the TS's and intermediate minima may be located. Multi-structure methods (e.g. NEB) are examples of the latter behaviour.
- (4) The reaction path formed by a sequence of points generated by constrained optimizations may be discontinuous. For methods where two points are gradually moved from the reactant and product sides (e.g. saddle and LTP), this means that the distance between end-points does not converge towards zero.
- (5) There may be more than one TS connecting two minima. As many of the interpolation methods starts off by assuming a linear reaction coordinate between the reactant and product, the user needs to guide the initial search (for example by adding various intermediate structures) to find more than one TS.
- (6) A significant advantage is that the constrained optimization can usually be carried out using only the first derivative of the energy. This avoids an explicit, and computationally expensive, calculation of the second derivative matrix.
- (7) For the one- and two-structure methods, each successive refinement of the TS estimate requires either location of an energy maximum or minimum along a one-dimensional path (typically a line), or a constrained optimization in an $N - 1$ dimensional hyperspace. A path minimization or maximization will normally involve several function evaluations, while a multi-dimensional minimization requires several gradient calculations. Geometry changes are often quite small in the endgame, but each step may still require a significant computational effort involving many function and/or gradient calculations. In such cases, it is often advantageous to switch to one of the Newton–Raphson methods described in Section 12.4.6, but the dimensionality of the problem may prevent this.
- (8) The multi-structure methods in Section 12.4.3 involve an optimization of a target function with M images each having $3N_{\text{atom}}$ coordinates, i.e. optimization of a function with $\sim 3MN_{\text{atom}}$ variables. Since the number of iterations typically increases with the number of variables, the optimization of the target function often requires a large number (several hundred or thousand) of gradients.
- (9) The multi-structure methods are quite tolerant towards the presence of many soft degrees of freedom, which often causes problems with the local optimization methods described in Sections 12.4.5–12.4.7. Multi-structure methods such as NEB are therefore well suited for large systems such as extended (periodic) systems.

12.4.5 Local methods: gradient norm minimization

Since transition structures are points where the gradient is zero, they may in principle be located by minimizing the gradient norm. This is in general not a good approach for two reasons:

- (1) There are typically many points where the gradient norm has a minimum without being zero.
- (2) Any stationary point has a gradient norm of zero, thus all types of saddle points and minima/maxima may be found, not just TS's.




Figure 12.10 An example of a function and the associated gradient norm

Figure 12.10 shows an example of a one-dimensional function and its associated gradient norm. It is clear that a gradient norm minimization will only locate one of the two stationary points if started near $x = 1$ or $x = 9$. Most other starting points will converge on the shallow part of the function near $x = 5$. The often very small convergence radius makes gradient norm minimizations impractical for routine use.

12.4.6 Local methods: Newton–Raphson

By far the most common local methods are based on the augmented Hessian Newton–Raphson approach. Sufficiently close to the TS, the standard NR formula will locate the TS rapidly. Sufficiently close means that the Hessian should have exactly one negative eigenvalue, and the eigenvector should be in the correct direction, along the “reaction coordinate”. Furthermore, the NR step should be inside the trust radius. By using augmented Hessian techniques, the convergence radius may be enlarged over the straight NR approach, and first-order saddle points may be located even when started in a region where the Hessian does not have the correct structure, as long as the lowest eigenvector is in the “correct” direction.

Near a first-order saddle point the NR step maximizes the energy in one direction (along the Hessian TS eigenvector) and minimizes the energy along all other directions. Such a step may be enforced by choosing suitable shift parameters in the augmented Hessian method, i.e. the step is parameterized as in eq. (12.12). The minimization step is similar to that described in Section 12.2.4 for locating minima, the only difference is for the unique TS mode.

In the *Partitioned Rational Function Optimization* (P-RFO), two shift parameters are employed.³⁶

$$\begin{aligned}\sum_{i \neq \text{TS}} \frac{f_i^2}{\epsilon_i - \lambda} &= \lambda \\ \frac{f_{\text{TS}}^2}{\epsilon_{\text{TS}} - \lambda_{\text{TS}}} &= \lambda_{\text{TS}}\end{aligned}\quad (12.27)$$

The λ for the minimization modes is determined as for the RFO method, eq. (12.14). The equation for λ_{TS} is quadratic and by choosing the solution that is larger than ϵ_{TS} it is guaranteed that the step component in this direction is along the gradient, i.e. a maximization. As for the RFO step, there is no guarantee that the total step length will be within the trust radius.

The *Quadratic Approximation* (QA) method uses only one shift parameter, requiring that $\lambda_{\text{TS}} = -\lambda$, and restricts the total step length to the trust radius (compare with eq. (12.15)).³⁷

$$|\Delta \mathbf{x}'|^2 = \sum_{i \neq \text{TS}} \left(\frac{f_i}{\epsilon_i - \lambda} \right)^2 + \left(\frac{f_{\text{TS}}}{\epsilon_{\text{TS}} + \lambda} \right)^2 = R^2 \quad (12.28)$$

The exact same formula may be derived using the concept of an “image potential” (obtained by inverting the sign of f_{TS} and λ_{TS}), and the QA name is often used together with the TRIM acronym, which stand for *Trust Radius Image Minimization*.³⁸

The ability of augmented Hessian methods for generating a search toward a first-order saddle point, even when started in a region where the Hessian has all positive eigenvalues, suggests that it may be possible to start directly from a minimum and “walk” to the TS by following a selected Hessian eigenvector uphill. Such mode followings, however, are only possible if the eigenvector being followed is only weakly coupled to the other eigenvectors (i.e. third and higher derivatives are small). All NR-based methods assume that one of the Hessian eigenvectors points in the general direction of the TS, but this is only strictly true when the higher order derivatives are small. If this is not the case, NR-based methods may fail to converge even when started from a “good” geometry, where the Hessian has one negative eigenvalue. Note also that the magnitude of the higher derivatives depends on the choice of coordinates, i.e. a “good” choice of coordinates may transform a divergent optimization into a convergent one.

All NR methods assume that a reasonable guess of the TS geometry is available. Generating this guess is part of the magic, but some of the interpolating schemes described in Sections 12.4.1–12.4.3 may be useful in this respect.

There are two main problems with all NR-based methods. One is the already mentioned need for a good starting geometry. The other is the requirement of a Hessian, which is quite expensive in terms of computer time for electronic structure methods. Contrary to minimizations, TS optimizations cannot start with a diagonal matrix and update it as the optimization proceeds. An NR TS search requires the definition of a direction along which to maximize the energy, the reaction vector, i.e. the start Hessian should preferably have one negative eigenvalue. Normally the Hessian needs to be calculated explicitly at the first step; at subsequent step the Hessian may be updated. An interesting alternative is to use a force field Hessian for starting the optimization, since

this effectively removes one of the more expensive steps in a TS optimization.³⁹ If the geometry changes substantially during the optimization, however, it may be necessary to recalculate the Hessian at certain intervals. Owing to the relatively high cost of calculating the energy, gradient and especially the Hessian, pseudo-Newton–Raphson methods have traditionally been the preferred algorithm with *ab initio* wave functions.

12.4.7 Local methods: the dimer method

The main problem with Newton–Raphson methods is the need for generating (calculating or updating) and manipulating (storing and diagonalizing) the Hessian matrix. The main function of the Hessian for saddle point optimization is to provide the direction along which the energy should be maximized. Sufficiently near the TS, the direction is along the eigenvector corresponding to the lowest eigenvalue. Determination of this direction, however, can be done *without* calculating the Hessian by placing two symmetrically displaced images, a *dimer*, and minimizing the sum of their energies, subject to a constant distance between them.⁴⁰ After minimization the lowest mode direction is given by the line connecting the two images, and it can be used for displacing the central structure, followed by a new dimer optimization. Since the dimer optimization can be done using only first derivatives, this alleviates the need for the Hessian matrix. There is relatively little experience with this method so far, but one would expect it to have the same requirements as Newton–Raphson-based methods, i.e. a good starting geometry is required for a stable convergence to the TS. Whether the added computational cost of optimizing each dimer configuration outweighs the saving by not having an explicit Hessian is unclear, and will in any case depend on the size of the system.

12.4.8 Coordinates for TS searches

The choice of a “good” set of coordinates is even more critical in TS optimizations than for minimizations. A good set of coordinates enlarges the convergence region, relaxing the requirement of a good starting geometry. On the other hand, a poor set of coordinates decreases the convergence radius, forcing the user to generate a starting point very close to the actual TS in order for NR methods to work. Furthermore, NR methods are best suited for relatively “stiff” systems; large flexible systems with many small eigenvalues in the Hessian are better handled by some of the interpolations methods, such as NEB.

Mapping out whole reaction pathways by locating minima and connecting TS’s is often computationally demanding. The (approximate) geometries of many of the important minima are often known in advance, and as mentioned above, energy minimizations are fairly uncomplicated. Locating TS’s is much more involved. On a multi-dimensional energy surface, there will in general not be TS’s connecting all pairs of minima. It is, however, essentially impossible to prove that a TS does not exist.

Symmetry can sometimes be used to facilitate the location of TS’s. For some reactions, especially those where the reactant and product are identical, the TS will have a symmetry different from the reactant/product. The reaction vector will belong to one of the non-totally symmetric representations in the point group. The TS can therefore

be located by constraining the geometry to a certain symmetry and *minimizing* the energy. Consider for example the S_N2 reaction of Cl^- with CH_3Cl . The reactant and product have C_{3v} symmetry, but the TS has D_{3h} symmetry. Minimizing the energy under the constraint that the geometry should have D_{3h} symmetry will produce the lowest energy structure within this symmetry, which is the TS.




Figure 12.11 The TS for an identity S_N2 reaction has a higher symmetry than the reactant/product

For non-identity reactions, it is often useful to start a search for stationary points by minimizing high-symmetry geometries. A subsequent frequency calculation on the symmetry-constrained (and minimized) structure will reveal the nature of the stationary point. If it is a minimum or TS we have already obtained useful information. If it turns out to be a higher order saddle point, the normal coordinates associated with the imaginary frequencies show how the symmetry should be lowered to produce lower energy species, which may be either minima or TS's. As calculations on highly symmetric geometries are computationally less expensive than on non-symmetric structures, it is often a quite efficient strategy to start the investigation by concentrating on structures with symmetry.

12.4.9 Characteristics of local methods

Local methods have the following characteristics:

- (1) A starting geometry close to the saddle point is needed. Especially for reactions that are not dominated by a few (internal) reaction variables, it may be difficult to generate such a guess. In many cases, the convergence radius is small, i.e. the starting geometry must be (very) close to the saddle point in order to converge.
- (2) Hessian-based methods (Section 12.4.6) require explicit calculation of the second derivative matrix, which may be computationally expensive. Furthermore, for large systems, an explicit handling of the Hessian matrix becomes problematic.
- (3) Systems with many soft vibrational modes are often problematic, as the resulting low Hessian eigenvalues interfere with the negative curvature along the reaction vector.
- (4) If a good starting geometry and Hessian is available, the convergence is rapid, often requiring only a few tens of gradient calculations.

12.4.10 Dynamic methods

The methods in Sections 12.4.1–12.4.8 focus on finding a TS connecting a reactant and product, and the resulting activation energy can provide reaction rates via the Arrhenius or Eyring formula (eqs (13.39) and (13.40)). For large complex systems,

however, the concept of a single “structure” becomes blurred. In cycloheptadecane, for example, there are hundreds of conformations within 10 kJ/mol of the global minimum, and any experimentally observed property at room temperature will be a Boltzmann average over many individual conformations. The reaction rate for a large system will similarly be a Boltzmann average over perhaps hundreds of TS’s, and a single reaction path connecting two minima via a saddle point no longer dominates the reaction rate.⁴¹ A systematic location of all minima (conformations) and corresponding TS’s followed by a Boltzmann averaging is a possibility, but this rapidly becomes unmanageable even for medium-sized systems. For large systems, one is therefore forced to perform a sampling of the TS’s in order to estimate the reaction rate, in analogy with the ensemble averaging discussed in Section 13.6 for minima. Standard molecular dynamics methods (Section 14.2.1) only sample the low-energy part of the surface, and are therefore unsuitable for the (high-energy) saddle point region. A specific part of the surface can be sampled by a biasing potential, as for example in the *umbrella sampling* technique (Section 14.2.9). Such an approach, however, requires *a priori* knowledge of the reaction path or at least the saddle point region. Alternatively, a dynamics simulation may be initiated in the transition state region and the trajectory followed in both directions.⁴²

12.5 Constrained Optimization Problems

In some cases, there are restrictions on the variables used to describe the function, as for example:

- (1) Certain geometrical constraints may be imposed. Experimental data, for example, may indicate that some atom pairs are within a certain distance of each other, or one may for analysis reasons want to impose certain geometrical restrictions on a molecular structure.
- (2) Fitting atomic charges to give a best match to a calculated electrostatic potential. The constraint is that the sum of atomic charges should equal the net charge of the molecule.
- (3) A variation of wave function coefficients is subject to constraints such as maintaining orthogonality of the MOs, and normalization of the MOs and the total wave function.
- (4) Finding conical intersections between different energy surfaces. The constraint is that two different energy functions should have the same energy for the same set of nuclear coordinates.

There are three main methods for enforcing constraints during function optimization:

- (1) Penalty functions
- (2) Lagrange undetermined multipliers
- (3) Projection methods.

The *penalty function* approach adds a term of the type $k(r - r_0)^2$ to the function to be optimized. The variable r is constrained to be near the target value r_0 , and the “force constant” k describes how important the constraint is compared with the unconstrained optimization. By making k arbitrarily large, the constraint may be fulfilled to any given accuracy. It cannot, however, make the constraint variable *exactly* equal to r_0 . This

would require the constant k to go towards infinity, and in practice cause numerical problems when it becomes sufficiently large compared with the other terms. The penalty function approach is often used for restricting geometrical variables, such as distances or angles, during geometry optimizations with force field methods. It may also be used for “driving” a selected variable (Section 12.4.1), such as a torsional angle. In certain cases, the constraint is not to limit a variable to a single value, but rather to keep it between lower and upper limits. This is typically the situation for refining a force field structure subject to constraints imposed by experimental nuclear Overhauser effect (NOE) data, and in such cases, the penalty function may be taken as a “flat bottom” potential, i.e. the penalty term is zero within the limits and rises harmonically outside the limits. The gradient for the penalty function simply has one additional term from each constraint, and the penalty function may be optimized using the methods described in Sections 12.2.1–12.2.3.

A more elegant method of enforcing constraints is the *Lagrange method*. The function to be optimized depends on a number of variables, $\mathbf{f}(x_1, x_2, \dots, x_N)$, and the constraint condition can always be written as another function, $\mathbf{g}(x_1, x_2, \dots, x_N) = c$. Define now a Lagrange function as the original function minus (or plus) a constant times the constraint condition.

$$\mathbf{L}(x_1, x_2, \dots, x_N, \lambda) = \mathbf{f}(x_1, x_2, \dots, x_N) - \lambda[\mathbf{g}(x_1, x_2, \dots, x_N) - c] \quad (12.29)$$

If there is more than one constraint, one additional multiplier term is added for each constraint. The optimization is then performed on the Lagrange function by requiring that the gradient components with respect to the x - and λ -variable(s) are equal to zero. The multiplier(s) λ can in many cases be given a physical interpretation at the end. In the variational treatment of an HF wave function (Section 3.3), the MO orthogonality constraints turn out to be MO energies, and the multiplier associated with normalization of the total CI wave function (Section 4.2) becomes the total energy.

The Lagrange method increases the number of variables by one for each constraint, which is counterintuitive since introduction of a constraint should decrease the number of variables by one. For simple objective and constraint functions, the reduction can be obtained by solving the constraint condition for one of the variables, and substituting it into the object function.

$$\begin{aligned} \mathbf{g}(x_1, x_2, \dots, x_{N-1}, x_N) = c &\Leftrightarrow x_N = \mathbf{h}(x_1, x_2, \dots, x_{N-1}, c) \\ \mathbf{f}(x_1, x_2, \dots, x_{N-1}, x_N) &\Rightarrow \mathbf{f}(x_1, x_2, \dots, x_{N-1}, \mathbf{h}(x_1, x_2, \dots, x_{N-1}, c)) \end{aligned} \quad (12.30)$$

In the large majority of cases, however, the object and constraint functions are so complicated that an analytical elimination of a variable is essentially intractable, and this is especially true when there is more than one constraint. The main exception is when the constraint equation is *linear*, in which case it can be considered as a vector in the coordinate space. Instead of eliminating a variable explicitly, the constraint condition can be fulfilled by removing the corresponding component of the object function by *projection* (Section 16.4), and performing the optimization on \mathbf{f}_p .

$$\mathbf{f}_p = \mathbf{f} - \langle \mathbf{f} | \mathbf{g} \rangle \frac{\mathbf{g}}{|\mathbf{g}|} \quad (12.31)$$

A general (non-linear) constraint condition can be approximately fulfilled by projecting out the first-order (linear) Taylor approximation to the function. Since the

optimization normally proceeds by iterative methods, the linear approximation may be sufficient in each step. Alternatively, a micro-iterate based on successive linear approximations may be performed in each optimization of the objective function.

12.6 Conformational Sampling and the Global Minimum Problem

The methods described in Section 12.2 can only locate the “nearest” minimum, which is normally a *local* minimum, when starting from a given set of variables. In some cases, the interest is in the lowest of all such minima, the *global* minimum; in other cases it is important to sample a large (preferably representative) set of local minima. Considering that the number of minima typically grows exponentially with the number of variables, the global optimization problem is an extremely difficult task for a multi-dimensional function.⁴³ It is often referred to as the *multiple minima* or *combinatorial explosion* problem in the literature.

Consider for example determining the lowest energy conformations of linear alkanes, $\text{CH}_3(\text{CH}_2)_{n+1}\text{CH}_3$, by a force field method, with three possible energy minima for rotation around each C–C bond. For butane, there are thus three conformations, one *anti* and two *gauche* (which are symmetry equivalent). These minima may be generated by starting optimizations from three torsional angles separated by 120° . In the $\text{CH}_3(\text{CH}_2)_{n+1}\text{CH}_3$ case there are n such rotatable bonds, giving a possible 3^n different conformations, and in order to find the global minimum, the energy must be calculated for all of them. Assume for the sake of argument that each conformation optimization takes one second of computer time. Table 12.3 gives the number of possible conformations, and the time required for optimizing them all.

Table 12.3 Possible conformations for linear alkanes, $\text{CH}_3(\text{CH}_2)_{n+1}\text{CH}_3$

n	Number of possible conformations (3^n)	Time (1 conformation = 1 second)
1	1	3 seconds
5	243	4 minutes
10	59049	16 hours
15	14 348 907	166 days

It is clear that if the degrees of freedom exceed ~15–20, a systematic search becomes impossible. For the linear alkanes, it is known in advance that *anti* conformations in general are favoured over *gauche*, thus we may put some restrictions on the search, such as having a maximum of three *gauche* interactions in total. For most systems, however, there are no good guidelines for such *a priori* selections. Furthermore, for some cases the sampling interval must be less than 120° ; in ring systems it may be more like 60° , increasing the potential number of conformations to 6^n . Cycloheptadecane is a frequently used test case for conformational searching, and various methods have established that there are 262 different conformations within 12 kJ/mol of the global minimum with the MM2 force field.⁴⁴ In the early 1990s, this system was close to the limit for being able to establish the global minimum, but with the increase in computer hardware performance such systems can now be treated within a few hours of

computer time. Nevertheless, the exponential increase in the number of conformations means that it is impossible to perform a complete sampling of systems with more than ~20 degrees of freedom.

The total number of conformations for a given resolution of each variable (e.g. 120° steps) can be thought of as branches in a combinatorial tree, as illustrated in Figure 12.12.




Figure 12.12 Visualizing conformations as a combinatorial tree

For a reasonable sized system, there may be certain combinations of torsional angle that always lead to high-energy structures, for example by atoms clashing. These combinations correspond to specific branches in the combinatorial tree (illustrated by dashed lines in Figure 12.12), and it is clear that these may be pruned from the search at an early stage. This allows somewhat larger systems to be treated compared with a brute force combinatorial search, but the number of possible conformations still increases rapidly with the size of the system.⁴⁵

Finding “reasonable” minima for large biomolecular systems is heavily dependent on selecting a “good” starting geometry. One way of attempting this is by “building up” the structure. A protein, for example, may be built from amino acids, which have been optimized to their global minimum, and/or smaller fragments of the whole structure may be subjected to a global minimum search. By combining such pre-optimized fragments, it is hoped that the starting geometry for the whole protein will also be “near” the global minimum for the full system.⁴⁶

The systematic, or grid, search is only possible for small molecules. For larger systems, there are methods that can be used for perturbing a geometry from one local minimum to another. Some commonly used methods for conformational sampling are:

- (1) Stochastic and Monte Carlo methods
- (2) Molecular dynamics
- (3) Simulated annealing
- (4) Genetic algorithms
- (5) Diffusion methods
- (6) Distance geometry methods.

None of these are guaranteed to find the global minimum, but they may in many cases generate a local minimum that is close in energy to the global minimum (but not necessarily close in terms of structure). A brief description of the ideas in these methods is given below. For simplicity, we assume that the optimization is of an energy as a function of atomic coordinates, but it is of course equally valid for any function depending on a set of parameters.

12.6.1 Stochastic and Monte Carlo methods

These methods start from a given geometry, which is typically a (local) minimum, and new configurations are generated by adding a random “kick” to one or more atoms. In *Monte Carlo* (MC) methods, the new geometry is accepted as a starting point for the next perturbing step if it is lower in energy than the current. Otherwise, the Boltzmann factor $e^{-\Delta E/kT}$ is calculated and compared with a random number between 0 and 1. If $e^{-\Delta E/kT}$ is less than this number, the new geometry is accepted, otherwise the next step is taken from the old geometry. This generates a sequence of configurations from which geometries may be selected for subsequent minimization. In order to have a reasonable acceptance ratio, however, the step size must be fairly small, and is often chosen to give an acceptance ratio of ~ 0.5 .

In *stochastic* methods, the random kick is somewhat larger and is usually performed on all the atoms, and a standard minimization is carried out starting at the perturbed geometry.⁴⁷ The efficiency of the optimization can be improved by first re-adjusting all bond lengths to values close to their starting values. The optimization may or may not produce a new minimum, and a database of all unique structures is gradually built up. A new perturbed geometry is then generated from one of the structures in the database and minimized, etc. There are several variations on how this is done:

- The length of the perturbing step is important; a small kick essentially always return the geometry to the starting minimum, a large kick may produce high-energy structures, which minimize to high-energy local minima.
- The perturbing step may be done directly in Cartesian coordinates or in a selected set of internal coordinates, such as torsional angles. The Cartesian procedure has the disadvantage that many of the perturbed geometries are high in energy as two (or more) atoms are moved close together by the kick, although this can be partly alleviated by an adjustment of the bond lengths prior to the optimization. The use of torsional angles as variables is highly efficient for acyclic systems but is problematic for cyclic and confined structures. Cyclic structures can be treated by opening the ring, performing a random perturbation of the torsional angles, and attempting to re-close the ring. In the majority of cases this is not possible, and this results in many trial structures being discarded.
- The perturbing step may be taken either from the last minimum found or from all the previous found minima, weighted by a probability factor such that low-energy minima are used more often than high-energy structures.

Kolossvary and Guida have proposed generating the perturbing step along the eigenvectors with small eigenvalues obtained by diagonalizing the Hessian matrix at each minimum, a method called *low-mode* search.⁴⁸ The premise is that the soft deformation modes for a given structure is likely to lead to low-energy transition structures, and consequently to other low-energy minima. The strategy is thus similar to the eigenvector-following tactic discussed in Section 12.4.6 for locating transition structures, except that that no attempt is made to find the actual TS. The interest is only in perturbing the geometry sufficiently to get “past” the TS, such that a minimization will locate a new minimum. The advantage of the low-mode search is that the search is concentrated on the low-energy part of the energy surface, and the method furthermore

essentially solves the problem of generating trial structures for ring systems. The number of acceptable trial structures generated by the open-perturb-re-close method is often very low, only a few percent, resulting in an inefficient search. Since the Hessian eigenvalues contain information about the coupling of the internal (torsional) coordinates, the low-mode technique can generate trial structures without opening and re-closing ring.

The disadvantage of the low-mode search is that it requires calculating and diagonalizing the Hessian matrix for each minimum found, which becomes problematic for systems with more than a few hundred atoms. In order to use the method for large systems, the soft Hessian modes can be calculated by an iterative procedure requiring only the gradient.⁴⁹ Although this solves the problem of calculating and diagonalizing the Hessian, the computational effort for determining the low-mode directions is still substantial. It has been suggested that for e.g. proteins, the low-mode directions for one minimum can be reused for other minima as well, thereby avoiding the expensive mode calculation.

The main problem with stochastic methods is generating trial structures. In small flexible molecules, the torsional angles form a good set of coordinates for randomly perturbing the geometry. For cyclic and confined structures, however, a perturbation of a single torsional angle will usually lead to a high-energy structure, either because the remaining (cyclic) structure becomes strained, or because of atoms clashing into each other. The low-mode technique solves this by determining the proper combinations of internal coordinates to avoid this, but the Hessian diagonalization prevents its use for systems with thousands of atoms. Stochastic methods are therefore primarily useful for searching the conformational space for flexible extended systems, but not for confined molecules such as proteins and DNA. Stochastic methods, however, have the big advantage that they can generate conformations separated by large energy barriers, since the random kick is performed without calculating any energies along the perturbing step, i.e. the conformations can “tunnel” through large energy barriers.

12.6.2 Molecular dynamics

Molecular Dynamics (MD) methods solve Newton's equation of motion for atoms on an energy surface (see Section 16.3.1). The available energy for the molecule is distributed between potential and kinetic energy, and molecules are thus able to overcome barriers separating minima if the barrier height is less than the total energy minus the potential energy. Given a high-enough energy, which is closely related to the simulation temperature, the dynamics will sample the whole surface but will also require an impractically long simulation time. Since quite small time steps must be used for integrating Newton's equation, the simulation time is short (pico- or nanoseconds). Combined with the use of “reasonable” temperatures (a few hundreds or thousands of degrees), this means that only the local area around the starting point is sampled, and that only relatively small barriers (a few tens of a kJ/mol) can be overcome. Different (local) minima may be generated by selecting configurations at suitable intervals during the simulation and subsequently minimizing these structures. MD methods use the inherent dynamics of the system to search out the low-energy deformation modes and they can be used for sampling the conformational space for large confined systems. MD methods are typically used for sampling the conformational space when

the starting geometry is derived from experimental information, such as an X-ray or NMR structure. The main disadvantage of MD is the inability to overcome barriers larger than the internal energy determined by the simulation temperature. Since this is one of the advantages of MC methods, it is no surprise that mixed MC/MD methods have been developed.⁵⁰

12.6.3 *Simulated annealing*

Both MD and MC methods employ a temperature as a guiding parameter for generating or accepting new geometries. At sufficiently high temperatures and long run times, all the conformational space is sampled. In *Simulated Annealing* (SA) techniques, the initial temperature is chosen to be high, maybe 2000–3000 K.⁵¹ An MD or MC run is then initiated, during which the temperature is slowly reduced. Initially the molecule is allowed to move over a large area, but as the temperature is decreased, it becomes trapped in a minimum. If the cooling is done infinitely slowly (implying an infinite run time), the resulting minimum is the global minimum. In practice, however, an MD or MC run is so short that only the local area is sampled. The name, simulated annealing, comes from the analogy of growing crystals. If a melt is cooled slowly, large single crystals can be formed. Such a single crystal represents the global energy minimum for a solid state. A rapid cooling produces a glass (local minimum), i.e. a disordered solid.

12.6.4 *Genetic algorithms*

Genetic Algorithms (GA) or *Evolutionary Algorithms* take their concepts and terminology from biology.⁵² The idea is to have a “population” of structures, each characterized by a set of “genes”. The “parent” structures are allowed to generate “children” having a mixture of the parent genes, allowing for “mutations” to occur in the process. The best species from a population are selected based on Darwin’s principle, survival of the fittest, and carried on to the next “generation”, while the less fit structures are discarded.

Consider for example a molecule having 20 torsional angles, which may have $\sim 10^9$ possible conformations. The species in an initial population of say 100 different conformations are characterized by their fitness, for example a low energy is equivalent to a good structure. These 100 structures are allowed to produce offspring with a probability depending on their fitness, i.e. low-energy structures are more likely to contribute to the next generation than high-energy conformations. Two child conformations can be generated by taking the first n torsional angles from one of the parents and the remaining $20 - n$ from the other (“single-point cross-over”), with the second child being the complementary. A small amount of mutation is usually allowed in the process, i.e. randomly changing angles to produce conformations outside the range contained in the current population. Having generated say 100 such children, their (minimized) energies are determined and a suitable portion of the best parent and children structures are carried over to the next generation. The population is then allowed to evolve for perhaps a few hundred generations. There are many variations on the GA method: varying the size of the population, the mutation rate, the breeding selection, the ratio of children to parents surviving to the next generation,

single- or multi-point cross-over, etc. Genetic algorithms have become popular in recent years as they are easy to implement and have proven to be robust for locating a point in parameter space close to the global minimum. If the parameters are coded into genes the sampling is pointwise, and the final structures should therefore be refined using a standard gradient optimization. Alternatively, the trial structures may be subjected to a local optimization, making the parameter space continuous.⁵³

12.6.5 Diffusion methods

In *Diffusion Methods* the energy function is changed such that it will eventually contain only one minimum.⁵⁴ The function may be changed for example by adding a contribution proportional to the local curvature of the function (second derivative). This means that minima are raised in energy, and saddle points (and maxima) are reduced in energy (negative curvature). Eventually only one minimum remains. Using the single minimum geometry of the modified potential, the process can be reversed, ending up with a minimum on the original surface that often (but not necessarily) is the global minimum. The mathematical description of this process turns out to be identical to the differential equation describing diffusion.




Figure 12.13 Illustration of the diffusion method

12.6.6 Distance geometry methods

The idea in *Distance Geometry* methods is that trial geometries can be generated from a set of lower and upper bounds on distances between all pairs of atoms.⁵⁵ The method

was originally developed for generating possible geometries based on experimental information such as NMR NOE effects, which place certain constraints on the distance between protons that may be far from each other in terms of bonding. The bonding information itself, however, also places restrictions on distances between all pairs of atoms. Once a set of upper and lower bounds for all pair distances has been assigned, many different trial sets of distances may be generated by selecting random numbers between these limit. Such a random distance matrix can then be translated into a three-dimensional structure, a procedure known as *embedding*. Distance geometry can thus be used for generating trial conformations that can be optimized using conventional methods. The main advantage of distance geometry method is the ease with which distant constraints between atoms far apart in terms of bonding can be translated into valid trial structures. Without such constraints, some of the other methods in this section are usually more efficient in searching the conformational space.

From the above it may be clear that MD, MC and stochastic methods tend to primarily sample the local area, generating a relatively large number of local minima in the process. The use of a larger step size in stochastic methods normally means that they are more efficient than MC or MD. Simulated annealing and diffusion methods, on the other hand, are primarily geared to locating the global minimum, and will in general only produce one final structure, this being the best estimate of the global minimum. Genetic algorithms also focus on the global minimum, but the final population contains a distribution of low-energy structures. Distance geometry methods are more or less random searches, where “impossible” structures are excluded. Simulated annealing normally explores a significantly smaller space than genetic algorithms, but there is currently no clear consensus on which method is better for locating the global minimum. It is most likely that the best method will depend on the problem at hand.

12.7 Molecular Docking

An important example of a global optimization problem is determining the best alignment of two molecules with respect to each other, typically trying to fit a small molecule into a large protein structure, a process called *docking*. Given an X-ray structure of an enzyme, preferably with a bound ligand to identify the active site, the ligand can be removed, and other (virtual) compounds may be docked into the active site to possibly identify new molecules with a stronger binding affinity. Since many drugs act by inhibiting specific enzymes, docking is an important element in drug design and lead optimization.

The process of docking a ligand into the active site of a rigid enzyme has six degrees of freedom, three translational and three rotational, besides those arising from the ligand conformations. The three translational degrees of freedom can be sampled on a grid, for example by placing the ligand centre of mass within a central box with grid points every 1 Å, which for even a rather small $10 \times 10 \times 10$ Å box generates ~1000 possible points. For each of these points, the overall rotational orientation of the ligand must be sampled, for example by the Euler angles, typically generating a few hundred possibilities. A specific set of intermolecular translational and rotational variables is called a *pose*, and each ligand conformation may thus have $\sim 10^5$ possible poses, even on a rather coarse grid. Even though the majority of these can be rejected based on for example atom pair distances between the ligand and receptor, the combinatorial

space is still large for even a relatively small ligand. A systematic sampling is often too demanding, and global optimization schemes such as genetic algorithms are often employed for solving the optimization problem. Since the interest is typically to dock perhaps thousands of ligands for (virtual) screening a library of compounds, this furthermore calls for a fast method for estimating the binding energy. Given that the binding energy is the Gibbs free energy, this is clearly a challenging task.

A standard force field energy-only attempt to calculate the enthalpic interaction and an estimate of the free energy by simulation methods (Section 14.5) is much too expensive computationally. Instead, the non-bonded part of a force field function can be augmented with empirical terms, hopefully capturing some of the entropy and solvent effects, and the resulting *scoring function* can be parameterized against experimental binding data. The entropy terms are typically structural descriptors, such as the number of torsional degrees of freedom and the number of hydrogen acceptors and donors, the argument being that the fixing of torsional angles by binding to the enzyme causes a rather constant loss of entropy for each entity.⁵⁶

$$\Delta G_{\text{scoring}} = a_1 \Delta E_{\text{vdw}} + a_2 \Delta E_{\text{el}} + a_3 \Delta G_{\text{rot}} + a_4 \Delta G_{\text{H-bond}} + a_5 \Delta G_{\text{solv}} + \dots \quad (12.32)$$

The a_i weighting factors can then be fitted to actual binding data for specific protein–ligand systems.

Developing scoring functions capable of accurately ranking binding energies is an active area of research but it is probably fair to say that no general scoring function has yet been developed. Some scoring functions employ only force field terms, such as the first two in eq. (12.32), others parameterize it entirely from descriptive terms, such as the last three in eq. (12.32), and some employs a mixture of these. It should be noted that the interaction of the enzyme atoms with potential ligand atoms at the grid points in the active site is the same for all ligands and their poses, and can therefore be pre-computed to save computational resources. The main purpose of the scoring function is to rapidly rank a large number of poses, from which a smaller number of promising candidates may be subjected to more refined methods for estimating the binding energies.

The main problem of docking ligands into an active site generated by removing an existing ligand from an X-ray structure is that the hole left behind naturally bears a strong resemblance to the compound removed. There will thus be a tendency of finding compounds differing only slightly from the already-known inhibitor. The fundamental problem is that the flexibility of the enzyme is neglected, i.e. the protein is able to some extent to adapt to the different shapes of the active site for different ligands. Taking the enzyme conformational degrees of freedom into consideration during the docking increases the computational problem to essentially unmanageable proportions. A heuristic proposition is to reduce the van der Waals parameters of the enzyme atoms at the surface of the active site, thereby allowing larger ligands to be docked. Subsequently, the original van der Waals parameters can be re-introduced followed by relaxation of the enzyme structure. Such methods are at the forefront of current research.

12.8 Intrinsic Reaction Coordinate Methods

The optimization methods described in Sections 12.2–12.4 concentrate on locating stationary points on an energy surface. The important points for discussing chemical reactions are minima, corresponding to reactant(s) and product(s), and saddle points,

corresponding to transition structures. Once a TS has been located, it should be verified that it indeed connects the desired minima. At the TS the vibrational normal coordinate associated with the imaginary frequency is the reaction coordinate (Section 16.2.2), and an inspection of the corresponding atomic motions may be a strong indication that it is the “correct” TS. A rigorous proof, however, requires a determination of the *Minimum Energy Path* (MEP) from the TS to the connecting minima. If the MEP is located in mass-weighted coordinates, it is called the *Intrinsic Reaction Coordinate* (IRC).⁵⁷ The IRC path is of special importance in connection with studies of reaction dynamics, since the nuclei will usually stay close to the IRC, and a model for the reaction surface may be constructed by expanding the energy to for example second order around points on the IRC (Section 14.2.7). The IRC is formally the reaction path taken in the limit of a zero temperature, and for a modest temperature the deviation from this path is usually small. For high temperatures, however, the favoured dynamical path will tend to be the shortest path, regardless of the fact that this may be significantly higher in energy than along the (longer) IRC path.

The IRC path is defined by the differential equation (12.33).

$$\frac{d\mathbf{x}(s)}{ds} = -\frac{\mathbf{g}}{|\mathbf{g}|} = \mathbf{t} \quad (12.33)$$

Here \mathbf{x} are the (mass-weighted) coordinates, s is the path length and \mathbf{t} is the (negative) normalized gradient. Determining the IRC requires solving eq. (12.33) starting from a geometry slightly displaced from the TS along the normal coordinate for the imaginary frequency.

The simplest method for integrating eq. (12.33) is the *Euler* method. A series of steps are taken in the opposite direction of the normalized gradient at the current geometry \mathbf{x}_n .

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \Delta s \mathbf{t}(\mathbf{x}_n) \quad (12.34)$$

This corresponds to a steepest descent minimization with a fixed step size Δs . As discussed in Section 12.2.1, such an approach tends to oscillate around the true path and consequently requires a small step size to follow the IRC accurately.

A more advanced method is the *Runge–Kutta* (RK) algorithm. The idea here is to generate some intermediate steps that allow a better and more stable estimate of the next geometry for a given step size. The *second-order* Runge–Kutta (RK2) method first calculates the gradient at a point corresponding to an Euler step with half the step size. The gradient at the halfway point is then used for taking the full step.

$$\begin{aligned} \mathbf{k}_1 &= \Delta s \mathbf{t}(\mathbf{x}_n) \\ \mathbf{k}_2 &= \Delta s \mathbf{t}(\mathbf{x}_n + \tfrac{1}{2} \mathbf{k}_1) \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + \mathbf{k}_2 \end{aligned} \quad (12.35)$$

The *fourth-order* Runge–Kutta (RK4) method generates four intermediate gradients, and combines the steps as follows.

$$\begin{aligned}
 \mathbf{k}_1 &= \Delta \mathbf{st}(\mathbf{x}_n) \\
 \mathbf{k}_2 &= \Delta \mathbf{st}(\mathbf{x}_n + \tfrac{1}{2} \mathbf{k}_1) \\
 \mathbf{k}_3 &= \Delta \mathbf{st}(\mathbf{x}_n + \tfrac{1}{2} \mathbf{k}_2) \\
 \mathbf{k}_4 &= \Delta \mathbf{st}(\mathbf{x}_n + \mathbf{k}_3) \\
 \mathbf{x}_{n+1} &= \mathbf{x}_n + \tfrac{1}{6} \mathbf{k}_1 + \tfrac{1}{3} \mathbf{k}_2 + \tfrac{1}{3} \mathbf{k}_3 + \tfrac{1}{6} \mathbf{k}_4
 \end{aligned}
 \tag{12.36}$$

Another method for following the IRC which does not rely on integration of the differential equation (12.33) has been developed by Gonzales and Schlegel (GS).⁵⁸ The idea is to generate points on the IRC by means of a series of constrained optimizations. The algorithm is illustrated in Figure 12.14.




Figure 12.14 Illustration of the Gonzales–Schlegel constrained optimization method for following an IRC

An expansion point is generated by taking a step along the current direction with a step size of $\frac{1}{2}\Delta s$. The energy is then minimized on a hypersphere with radius $\frac{1}{2}\Delta s$, located at the expansion point. This is an example of a constrained optimization that can be handled by means of a Lagrange multiplier (Section 12.5). The GS procedure ensures that the tangent to the IRC path is correct at each point.

Although it is clear that RK4 is more stable and accurate than the Euler method for a given step size, this does not necessarily mean that it is the most efficient method. Since the RK4 method requires four gradient calculations for each step, the simple Euler can employ a step size 4 times as small for the same computational cost. Similarly, although the Gonzales–Schlegel method appears to be quite tolerant for large step sizes, each constrained optimization may take a significant number of gradient calculations to converge, which could also be used for advancing the Euler algorithm at a slower pace. Nevertheless, the Gonzales–Schlegel method appears at present to be one of better methods for accurately following the IRC path. Which algorithm is the optimum will depend on the system at hand and the required accuracy of the IRC path. If only the nature of the two minima on each side of the TS is required, a crude IRC is sufficient, and a simple Euler algorithm may be the most cost efficient. For use in connection with reaction path methods (Section 14.2.7), however, the IRC needs to be located very accurately, and a sophisticated method and a small step size may be required.

References

1. T. Schlick, *Rev. Comp. Chem.*, **3** (1992), 1; H. B. Schlegel, *Adv. Chem. Phys.*, **67** (1987), 249; H. B. Schlegel, *Modern Electronic Structure Theory*, Part I, D. Yarkony, Ed., World Scientific, 1995, pp. 459–500; R. Fletcher, *Practical Methods of Optimization*, Wiley, 1980; M. L. McKee, M. Page, *Rev. Comp. Chem.*, **4** (1993), 35.
2. J. A. Nelder, R. Mead, *Computer J.*, **7** (1965), 308.
3. A. Banerjee, N. Adams, J. Simons, R. Shepard, *J. Phys. Chem.*, **89** (1985), 52.
4. P. Culot, G. Dive, V. H. Nguyen, J. M. Ghuysen, *Theor. Chim. Acta.*, **82** (1992), 189.
5. J. M. Anglada, J. M. Bofill, *J. Comp. Chem.*, **19** (1998), 349.
6. H. B. Schlegel, *Theor. Chim. Acta.*, **66** (1984), 333.
7. P. Csaszar, P. Pulay, *J. Mol. Struct.*, **114** (1984), 31.
8. F. Eckert, P. Pulay, H.-J. Werner, *J. Comp. Chem.*, **18** (1997), 1473.
9. G. Fogarasi, X. Zhou, P. W. Taylor, P. Pulay, *J. Am. Chem. Soc.*, **114** (1992), 8191.
10. J. Baker, P. Pulay, *J. Comp. Chem.*, **21** (2000), 69; P. E. Maslen, *J. Chem. Phys.*, **122** (2005), 014104.
11. C. Peng, P. Y. Ayala, H. B. Schlegel, M. J. Frisch, *J. Comp. Chem.*, **17** (1996), 49.
12. J. Baker, A. Kessi, B. Delley, *J. Chem. Phys.*, **105** (1996), 192.
13. F. Eckert, P. Pulay, H.-J. Werner, *J. Comp. Chem.*, **18** (1997), 1473; V. Bakken, T. Helgaker, *J. Chem. Phys.*, **117** (2002), 9160.
14. G. Henkelman, G. Jóhannesson, H. Jónsson, *Prog. Theo. Chem. Phys.*, **5** (2000), 269.
15. H. B. Schlegel, *J. Comp. Chem.*, **24** (2003), 1514.
16. T. A. Halgren, W. N. Lipscomb, *Chem. Phys. Lett.*, **49** (1977), 225.
17. S. Bell, J. S. Crighton, *J. Chem. Phys.*, **80** (1984), 2464.
18. C. Peng, H. B. Schlegel, *Isr. J. Chem.*, **33** (1993), 449.
19. Y. Abashkin, N. Russo, *J. Chem. Phys.*, **100** (1994), 4477.
20. K. Ohno, S. Maeda, *Chem. Phys. Lett.*, **384** (2004), 722; S. Maeda, K. Ohno, *J. Phys. Chem. A*, **109** (2005), 5742.
21. G. T. Barkema, N. Mousseau, *Phys. Rev. Lett.*, **77** (1996), 4358.
22. M. J. S. Dewar, E. F. Healy, J. J. P. Stewart, *J. Chem. Soc., Faraday Trans. 2*, **80** (1984), 227.
23. C. Cardenas-Lailhacar, M. C. Zerner, *Int. J. Quant. Chem.*, **55** (1995), 429.
24. I. V. Ionova, E. A. Carter, *J. Chem. Phys.*, **98** (1993), 6377.
25. R. A. Miron, K. A. Fichthorn, *J. Chem. Phys.*, **115** (2001), 8742.
26. R. Czerminski, R. Elber, *Int. J. Quant. Chem. Symp.*, **24** (1990), 167.
27. P. Y. Ayala, H. B. Schlegel, *J. Chem. Phys.*, **107** (1997), 375.
28. D. A. Liotard, *Int. J. Quant. Chem.*, **44** (1992), 723.
29. C. Choi, R. Elber, *J. Chem. Phys.*, **94** (1991), 751.
30. T. Fischer, M. Karplus, *Chem. Phys. Lett.*, **194** (1992), 252.
31. G. Henkelman, H. Jónsson, *J. Chem. Phys.*, **113** (2000), 9978.
32. G. Henkelman, B. P. Uberuaga, H. Jónsson, *J. Chem. Phys.*, **113** (2000), 9901.
33. W. E. W. Ren, E. Vanden-Eijnden, *Phys. Rev. B*, **66** (2002), 52301.
34. P. Maragakis, S. A. Andreiev, Y. Brumer, D. R. Reichman, E. Kaxiras, *J. Chem. Phys.*, **117** (2002), 4651.
35. B. Peters, A. Heyden, A. T. Bell, A. Chakraborty, *J. Chem. Phys.*, **120** (2004), 7877.
36. A. Banerjee, N. Adams, J. Simons, R. Shepard, *J. Phys. Chem.*, **89** (1985), 52.
37. P. Culot, G. Dive, V. H. Nguyen, J. M. Ghuysen, *Theor. Chim. Acta.*, **82** (1992), 189.
38. T. Helgaker, *Chem. Phys. Lett.*, **182** (1991), 503.
39. F. Jensen, *J. Chem. Phys.*, **119** (2003), 8804.
40. G. Henkelman, H. Jónsson, *J. Phys. Chem.*, **111** (1999), 7010; A. Heyden, A. T. Bell, F. J. Keil, *J. Chem. Phys.*, **123** (2005), 224101.

41. C. Dellago, P. G. Bolhuis, P. L. Geissler, *Adv. Chem. Phys.*, **123** (2002), 1.
42. J. E. Basner, S. D. Schwartz, *J. Am. Chem. Soc.*, **127** (2005), 13822.
43. A. R. Leach, *Rev. Comp. Chem.*, **2** (1991), 1; A. E. Howard, P. A. Kollman, *J. Med. Chem.*, **31** (1988), 1669; C. D. Maranus, C. A. Floudas, *J. Chem. Phys.*, **100** (1994), 1247.
44. M. Saunders, K. N. Houk, Y.-D. Wu, W. C. Still, M. Lipton, G. Chang, W. C. Guida, *J. Am. Chem. Soc.*, **112** (1990), 1419.
45. F. Villamagna, M. A. Whitehead, *J. Soc. Chem. Faraday Trans.*, **90** (1994), 47.
46. A. Smellie, R. Stanton, R. Henne, S. Teig, *J. Comp. Chem.*, **24** (2002), 10.
47. G. Chang, W. C. Guida, W. C. Still, *J. Am. Chem. Soc.*, **111** (1989), 4379; J. Chandrasekhar, M. Saunders, W. L. Jorgensen, *J. Comp. Chem.*, **22** (2001), 1646.
48. I. Kolossvary, W. C. Guida, *J. Am. Chem. Soc.*, **118** (1996), 5011.
49. I. Kolossvary, G. M. Keseru, *J. Comp. Chem.*, **22** (2001), 21.
50. S. Duane, A. Kennedy, B. J. Pendleton, D. Roweth, *Phys. Lett. B*, **195** (1987), 216; R. Faller, J. J. de Pablo, *J. Chem. Phys.*, **116** (2002), 55.
51. S. Kirkpatrick, C. D. Gelatt Jr, M. P. Vecchi, *Science*, **220** (1983), 671; S. R. Wilson, W. Cui, *Biopolymers*, **29** (1990), 225.
52. R. S. Judson, *Rev. Comp. Chem.*, **10** (1997), 1.
53. G. A. Morris, D. S. Goodsell, R. S. Halliday, R. Huey, W. E. Hart, R. K. Belew, A. J. Olson, *J. Comp. Chem.*, **19** (1998), 1639; F. Koskowski, B. Hartke, *J. Comp. Chem.*, **26** (2005), 1169.
54. J. Kostrowicki, H. A. Scheraga, *J. Phys. Chem.*, **96** (1992), 7442.
55. J. M. Blaney, J. S. Dixon, *Rev. Comp. Chem.*, **5** (1994), 299.
56. R. Wang, Y. Lu, S. Wang, *J. Med. Chem.*, **46** (2003), 2287; I. Muegge, M. Rarey, *Rev. Comp. Chem.*, **17** (2001), 1.
57. K. Fukui, *Acc. Chem. Res.*, **14** (1981), 363.
58. C. Gonzales, H. B. Schlegel, *J. Chem. Phys.*, **95** (1991), 5853.

13 Statistical Mechanics and Transition State Theory

The separation of the nuclear and electronic degrees of freedom by the Born–Oppenheimer approximation leads to a mental picture of a chemical reaction as nuclei moving on a potential energy surface. The easiest path from one minimum to another, i.e. for transforming one chemical species to another, is along the reaction path having the lowest energy. The highest energy point along this path is the transition structure, and the energy relative to the reactant completely determines the reaction rate within *Transition State Theory* (TST). Transition state theory is a semi-classical theory where the quantum nature is taken into account by means of the quantization of vibrational and rotational energy states. The connection between the properties of a single molecule and the experimental conditions employing a very large number of species is given by statistical mechanics, which provides a framework for performing the statistical averaging over a very large number of possible energy distributions. For an ideal gas the averaging can be performed in a closed analytical form within the rigid-rotor harmonic-oscillator approximation. For systems in condensed states, i.e. liquid or solid states, the averaging must be done by explicitly sampling the phase space.

13.1 Transition State Theory

Consider a chemical reaction of the type $A + B \rightarrow C + D$. The rate of reaction may be written as in eq. (13.1), with k_{rate} being the rate constant.

$$\frac{d[C]}{dt} = \frac{d[D]}{dt} = -\frac{d[A]}{dt} = -\frac{d[B]}{dt} = k_{\text{rate}}[A][B] \quad (13.1)$$

If k_{rate} is known, the concentration of the various species can be calculated at any given time from the initial concentrations. At the *microscopic* level, the rate constant is a

function of the quantum states of A, B, C and D, i.e. the electronic, translational, rotational and vibrational quantum numbers. The *macroscopic* rate constant is an average over such “microscopic” rate constants, weighted by the probability of finding a molecule with a given set of quantum numbers. For systems in equilibrium, the probability of finding a molecule in a certain state depends on its energy by means of the Boltzmann distribution, and the macroscopic rate constant thereby becomes a function of temperature.

Stable molecules correspond to minima on the potential energy surface within the Born–Oppenheimer approximation and a chemical reaction can be described as nuclei moving from one minimum to another. In the lowest level of approximation, the motion is assumed to occur along the path of least energy, and this path forms the basis for transition state theory.¹ The *Transition State* is the configuration that divides the reactant and product parts of the surface (i.e. a molecule that has reached the transition state will continue on to product), while the geometrical configuration of the energy maximum along the reaction path is called the *Transition Structure*. The transition state is thus a macroscopic ensemble with a Boltzmann energy distribution, while the transition structure refers to the microscopic system. The two terms are often used interchangeably, and share the same acronym, TS. In the multi-dimensional case, the TS is a first-order saddle point on the potential energy surface, a maximum in the reaction coordinate direction and a minimum along all other coordinates.




Figure 13.1 Schematic illustration of a reaction path

TST is a semi-classical theory where the dynamics along the reaction coordinate is treated classically, while the perpendicular directions take into account the quantization of for example the vibrational energy. It furthermore assumes an equilibrium energy distribution among all possible quantum states at all points along the reaction coordinate. The probability of finding a molecule in a given quantum state is proportional to $e^{-\Delta E/kT}$, which is a Boltzmann distribution. Assuming that the molecules at the TS are in equilibrium with the reactant, the macroscopic rate constant can be expressed as in eq. (13.2).

$$k_{\text{rate}} = \frac{kT}{h} e^{-\Delta G^\ddagger/RT} \quad (13.2)$$

$$\Delta G^\ddagger = \Delta G_{\text{TS}} - \Delta G_{\text{reactant}}$$

ΔG^\ddagger is the Gibbs free energy difference between the TS and reactant, and k is Boltzmann's constant. Actually, the TST expression only holds if all molecules that pass from the reactant over the TS go on to product. The dividing surface separating the reactant from the product is a hyperplane perpendicular to the reaction coordinate at the TS. The TST assumption is that no re-crossings occur for a given temperature, i.e. all molecules passing through the dividing surface will go on to form product. Note that this indicates that the rate constant calculated from eq. (13.2) will always be an upper limit to the true rate constant. In more refined models, the dividing surface may be located by minimizing the flux through the surface, i.e. forming the dynamical bottleneck for the reaction, for example by taking dynamics and entropy effects into account. To allow for "re-crossings", where a molecule passes over the TS but is reflected back to the reactant side, a *transmission coefficient* κ is sometimes introduced. This factor also allows for the quantum mechanical phenomenon of tunnelling, i.e. molecules that have insufficient energy to pass *over* the TS may tunnel *through* the barrier and appear on the product side. The transmission coefficient is difficult to calculate but is usually close to 1 and rarely falls outside the range 0.5–2. At low temperatures the tunnelling contribution dominates, leading to $\kappa > 1$, while the re-crossing effect is the most important at high temperatures, giving $\kappa < 1$. For the majority of reactions the calculated accuracy in ΔG^\ddagger introduces errors much larger than a factor of 2 and the transmission coefficient is usually ignored.

From the TST expression in eq. (13.2) it is clear that if the free energy of the reactant and TS can be calculated, the reaction rate follows trivially. Similarly, the equilibrium constant for a reaction can be calculated from the free energy difference between the reactant(s) and product(s).

$$K_{\text{eq}} = e^{-\Delta G_0/RT} \quad (13.3)$$

The Gibbs free energy is given in terms of the enthalpy and entropy, $G = H - TS$, and the enthalpy and entropy for a macroscopic ensemble of particles may be calculated from properties of a relatively few molecules by means of statistical mechanics, as discussed in Section 13.4.

The picture in Figure 13.1 relates to chemical reactions occurring on a single energy surface, as is typical for a thermal reaction. Photochemical reactions, on the other hand, occur on at least two and possibly more surfaces. The reaction is initiated by absorption of a photon to produce an excited state with the same nuclear coordinates as the ground state. This geometry will rarely be a stationary point on the excited surface, and the resulting nuclear movements may be explored by minimization or dynamical methods analogous to those on the ground state (Chapters 12 and 14). At some point, however, the system must return to the ground electronic surface. While this can occur by a radiative transition (fluorescence or phosphorescence), it may also occur by a non-radiative process where the excess energy is transferred to vibrational energy on the ground state surface. The probability for the latter process depends on the energy difference between the two surfaces, and therefore has a tendency of occurring at nuclear

geometries where the two surfaces “touch” each other, a point known as a *conical intersection*.² Two energy surfaces of the same symmetry cannot cross for a diatomic system, and instead make an avoided crossing, as illustrated in Figure 3.2. In a multi-dimensional system, however, there is no such restriction, and two energy surfaces may have the same energy for the same set of nuclear coordinates. Locating such conical intersections is a constrained optimization problem, involving finding a set of nuclear coordinates where two different energy functions have the same value, and also involves the non-adiabatic coupling elements between different wave functions, as discussed in Section 3.1.

A conical intersection may be thought of as a “funnel” in the $N - 2$ dimensional subspace that serves as the dynamical bottleneck for a photochemical reaction, analogously to the TS for a thermal reaction. It should be noted, however, that the transition between the excited and ground state surfaces may occur over quite a wide range of nuclear configurations, making the concept of a single “transition structure” somewhat blurred. Furthermore, for many systems it is the movement on the excited state surface to achieve the geometry of the conical intersection that limits the reaction rate, and not the actual transition between the two surfaces at the conical intersection.

13.2 Rice–Ramsperger–Kassel–Marcus Theory

The canonical TST theory in Section 13.1 assumes fast energy exchange with the surroundings, i.e. that the reacting molecule is in thermal equilibrium with the environment. For unimolecular reactions in the gas phase this assumption may not hold, especially not if the pressure is low (e.g. fragmentations in a mass spectrometer). Alternatively, TST may be formulated in terms of the total energy, also known as *micro-canonical* TST. When applied to unimolecular reactions, this is usually known as *Rice–Ramsperger–Kassel–Marcus* (RRKM) theory. The fundamental assumption here is that no re-crossing occurs for a given total energy of the molecule.

Consider a reaction where a molecule A acquires energy by collision with a molecule M (which may be the same as A) to form an energized molecule A^* , with the energy being distributed between the translation, rotation and vibrational degrees of freedom. The vibrational energy can be transferred between the different modes owing to vibrational anharmonicity, and if it is higher than the activation energy E^\ddagger , it may at some point accumulate in a specific mode to reach an activated state A^\ddagger (transition state) leading to a chemical product P.



Assuming that the decay rate k^\ddagger for the activated A^\ddagger is much faster than k_2 , the rate for production of P can be written in terms of the k_1 , k_{-1} and k_2 constants by making a steady state approximation for A^* , as shown in eq. (13.5).

$$\frac{d[P]}{dt} = k_2[A^*] = \frac{k_1 k_2 [M][A]}{k_{-1}[M] + k_2} = k_{\text{eff}}[A] \quad (13.5)$$

The effective rate constant k_{eff} is thus a function of the concentration of M, i.e. the pressure of the gas. The amount of energy transferred to A^* by M will be a variable, and the rate constants for the activation and reaction (but not the deactivation) will depend on the energy, i.e. $k_1(E)$ and $k_2(E)$. The effective rate constant in a small energy interval around E is obtained by rearranging eq. (13.5).

$$k_{\text{eff}}(E + dE) = \frac{(dk_1(E)/k_{-1})k_2(E)}{1 + k_2(E)/k_{-1}[M]} \quad (13.6)$$

The ratio k_1/k_{-1} is the equilibrium constant for the first step in eq. (13.4) and $dk_1(E)/k_{-1}$ is the probability of A^* being in a state with energy E , $P(E)$. The $k_{-1}[M]$ factor is the collision frequency for deactivation that is usually denoted by ω . The unimolecular rate constant can be obtained by integrating the effective rate constant over all energies higher than the activation energy.

$$k_{\text{uni}} = \int_{E^\ddagger}^{\infty} \frac{k_2(E)}{1 + k_2(E)/\omega} P(E) dE \quad (13.7)$$

The probability factor $P(E)$ is given by a Boltzmann distribution for the reactant, while $k_2(E)$ is determined by the number of vibrational quantum states for the activated state A^\ddagger . The details are sufficiently complex that the reader is referred to more specialized textbooks,³ but the essence of eq. (13.7) is that the rate constant can be evaluated from the geometries and vibrational frequencies of the reactant and activated complex. In the fast energy exchange limit (i.e. $\omega \rightarrow \infty$) the RRKM expression becomes equivalent to the TST expression (eq. (13.2)). RRKM calculations typically assume harmonic vibrations, which may be poor for high-barrier reactions where the vibrational anharmonicity significantly increases the state count. An exact calculation of all the anharmonic vibrational states, however, is a significant computational undertaking.

13.3 Dynamical Effects

The inherent assumption of both TST and RRKM is that the internal (vibrational) energy redistribution is significantly faster than the timescale for breaking/forming a bond. This means that the reaction rate only depends on the total *amount* of internal energy, not on *how* the energy is acquired. In other words, the reaction is independent of whether the energy is supplied by excitation of bending or stretching vibrations. In the large majority of chemical reactions, this is probably a valid assumption. For certain reactions where the reaction path involves an intermediate, however, the product distribution indicates that the energy is not completely randomized for the intermediate, i.e. the timescale for internal redistribution is comparable to that for the progression along the reaction coordinate.⁴ A specific example is shown in Figure 13.2.

The reaction in Figure 13.2 involves a biradical intermediate and if it has a sufficiently long lifetime, the thermal randomization of the energy should lead to a symmetric product distribution. Experimentally, however, the *exo* product is found to be favoured over the *endo* isomer by 4:1. Given that the potential energy surface is symmetric, this has been interpreted as a non-statistical distribution of the internal kinetic energy in the bond-forming step.⁵ The bond-breaking reaction occurs from a sample of molecules with a Boltzmann energy distribution, but the small fraction of molecules




Figure 13.2 Thermal decomposition leads to a 4:1 preference for the *exo* isomer

that in a given timeframe actually reacts must necessarily have the energy localized in the C–N bonds. The molecules passing over the first TS (breaking the C–N bonds) therefore enter the biradical minimum on the potential energy surface with the nuclear kinetic energy in a non-random fashion. In the example shown in Figure 13.2, a direct continuation of the nuclear movement arising from breaking both the C–N bonds leads to the *exo* product. The favouring of the *exo* product can thus be explained by the molecules “surfing” over the intermediate minimum on the potential energy surface, rather than being trapped and thermally randomized. Describing such effects requires an explicit simulation of the dynamics, as discussed in Section 14.2.

13.4 Statistical Mechanics

Most experiments are performed on macroscopic samples, containing perhaps $\sim 10^{20}$ particles. Calculations, on the other hand, are performed on relatively few particles, typically $1\text{--}10^3$, or up to 10^6 in special cases. The (macroscopic) result of an experimental measurement can be connected with properties of the microscopic system. The temperature, for example, is related to the average kinetic energy of the particles.

$$\langle E_{\text{kin}} \rangle = \frac{3}{2} RT \quad (13.8)$$

The connection between properties of a microscopic system and a macroscopic sample is provided by *statistical mechanics*.

At a temperature of 0 K, all molecules are in their energetic ground state but at a finite temperature there is a distribution of molecules in all possible (quantum) energy states. The relative probability P of a molecule being in a state with an energy ε at a temperature T is given by a Boltzmann factor.

$$P \propto e^{-\varepsilon/kT} \quad (13.9)$$

The exponential dependence on the energy means that there is a low (but non-zero) probability for finding a molecule in a high-energy state. This decreased probability for high-energy states is partly offset by the fact that there are many more states with high




Figure 13.3 Boltzmann energy distribution

energy than low energy. The most probable energy of a molecule in a macroscopic ensemble is therefore not necessarily the one with lowest energy, and a typical distribution is shown in Figure 13.3.

The key feature in statistical mechanics is the *partition function*.⁶ Just as the wave function is the cornerstone in quantum mechanics (from which everything else can be calculated by applying proper operators), the partition function allows calculation of all macroscopic functions in statistical mechanics.

The partition function for a single molecule is usually denoted q , and is defined as a sum of exponential terms involving all possible quantum energy *states*.

$$q = \sum_{i=\text{states}} e^{-\epsilon_i/kT} \quad (13.10)$$

The partition function can also be written as a sum over all distinct energy *levels*, multiplied with a degeneracy factor g_i that indicates how many states there are with the same energy ϵ_i .

$$q = \sum_{i=\text{levels}} g_i e^{-\epsilon_i/kT} \quad (13.11)$$

The partition function can be considered as an average excited state number-operator, since it is the probability-weighted sum of energy states, each counted with a factor of 1. It may also be viewed as the normalization factor for the Boltzmann probability distribution.

$$P(\epsilon_i) = q^{-1} e^{-\epsilon_i/kT} \quad (13.12)$$

The partition function q is for a single particle; the corresponding quantity Q for a collection of N *non-interacting* particles (ideal gas) is given in eq. (13.13).

$$\begin{aligned} Q &= q^N \quad (\text{different particles, non-interacting}) \\ Q &= \frac{q^N}{N!} \quad (\text{identical particles, non-interacting}) \end{aligned} \quad (13.13)$$

If the particles are *interacting* (liquid or solid state), the partition function Q must be calculated by summing over all energy states E_i for the whole system. Note that Q here describes the whole system consisting of N interacting particles, and the energy states E_i are consequently for all the particles.

$$Q = \sum_i e^{-E_i/kT} \quad (13.14)$$

Owing to the closely spaced energy levels, quantum effects can often be neglected and the state distribution treated as continuous. This corresponds to replacing the discrete sum over energies by an integral over all coordinates (\mathbf{r}) and momentum (\mathbf{p}), called the *phase space*.

$$Q = \int e^{-E(\mathbf{r},\mathbf{p})/kT} d\mathbf{r}d\mathbf{p} \quad (13.15)$$

More correctly, the partition function in eq. (13.15) should be written in terms of the Hamiltonian for the system, i.e. replacing E with \mathbf{H} . The kinetic and potential energy components, however, can be separated ($\mathbf{H} = \mathbf{T} + \mathbf{V}$). In the absence of potential energy, the Hamiltonian is purely kinetic energy and the system is an ideal gas. The interesting component is therefore the potential energy part in the partition function, which we denote by E . In the large majority of cases, the energy E is of the force field type described in Chapter 2.

The significance of the partition function Q is that thermodynamic functions, such as the internal energy U and Helmholtz free energy A ($A = U - TS$) can be calculated from it.

$$U = kT^2 \left(\frac{\partial \ln Q}{\partial T} \right)_V \quad (13.16)$$

$$A = -kT \ln Q$$

Macroscopic observables, such as pressure P and heat capacity at constant volume C_V , may be calculated as derivatives of thermodynamic functions.

$$P = - \left(\frac{\partial A}{\partial V} \right)_T = kT \left(\frac{\partial \ln Q}{\partial V} \right)_T \quad (13.17)$$

$$C_V = \left(\frac{\partial U}{\partial T} \right)_V = 2kT \left(\frac{\partial \ln Q}{\partial T} \right)_V + kT^2 \left(\frac{\partial^2 \ln Q}{\partial T^2} \right)_V$$

Other thermodynamic functions, such as the enthalpy H , the entropy S and Gibbs free energy G , may be constructed from these relations.

$$H = U + PV = kT^2 \left(\frac{\partial \ln Q}{\partial T} \right)_V + kTV \left(\frac{\partial \ln Q}{\partial V} \right)_T$$

$$S = \frac{U - A}{T} = kT \left(\frac{\partial \ln Q}{\partial T} \right)_V + k \ln Q \quad (13.18)$$

$$G = H - TS = kTV \left(\frac{\partial \ln Q}{\partial V} \right)_T - kT \ln Q$$

Note the difference between energetic properties such as U , P and H , which all depend on derivatives of Q , and entropic properties such as A , S and G , which depend directly

on Q . For simplicity, we will use U and A for illustrations in the following, but other quantities such as H and S can be treated completely analogously.

In order to calculate the partition function $q(Q)$, one needs to know *all* possible quantum states for the system. In principle, these can be calculated by solving the nuclear Schrödinger equation, once a suitable potential energy surface is available, for example from solving the electronic Schrödinger equation. Such a rigorous approach is only possible for di- and triatomic systems. For an isolated polyatomic molecule, the energy levels for a single conformation can be calculated within the *rigid-rotor harmonic-oscillator* (RRHO) approximation, where the electronic, vibrational and rotational degrees of freedom are assumed to be separable. Additional conformations can be included straightforwardly by simply offsetting the energy scale relative to the most stable conformation. An isolated molecule corresponds to an ideal gas state, and the partition function can be calculated *exactly* for such a system within the RRHO approximation, as discussed in Section 13.5.

For a condensed phase (liquid, solution, solid) the intermolecular interaction is comparable to or larger than a typical kinetic energy, and no separation of degrees of freedom is possible. Calculating the partition function by summing over all energy levels, or integrating over all phase space, is therefore impossible. It is, however, possible by sampling to estimate differences in Q and derivatives such as $\partial \ln Q / \partial T$ from a representative sample of the phase space, as discussed in Section 13.6.

13.5 The Ideal Gas, Rigid-Rotor Harmonic-Oscillator Approximation

For an isolated molecule, the total energy can be approximated as a sum of terms involving translational, rotational, vibrational and electronic states, and this is a good approximation for the large majority of systems. For linear, “floppy” (soft bending potential) molecules the separation of the rotational and vibrational modes may be problematic. If two energy surfaces come close together (avoided crossing), the separability of the electronic and vibrational modes may be a poor approximation (break-down of the Born–Oppenheimer approximation, Section 3.1).

There are in principle also energy levels associated with nuclear spins. In the absence of an external magnetic field, these are degenerate and consequently contribute a constant term to the partition function. As nuclear spins do not change during chemical reactions, we will ignore this contribution.

The assumption that the energy can be written as a sum of terms implies that the partition function can be written as a product of terms. As the enthalpy and entropy contributions involve taking the logarithm of q , the product of q ’s thus transforms into sums of enthalpy and entropy contributions.

$$\begin{aligned} \epsilon_{\text{tot}} &= \epsilon_{\text{trans}} + \epsilon_{\text{rot}} + \epsilon_{\text{vib}} + \epsilon_{\text{elec}} \\ q_{\text{tot}} &= q_{\text{trans}} q_{\text{rot}} q_{\text{vib}} q_{\text{elec}} \\ H_{\text{tot}} &= H_{\text{trans}} + H_{\text{rot}} + H_{\text{vib}} + H_{\text{elec}} \\ S_{\text{tot}} &= S_{\text{trans}} + S_{\text{rot}} + S_{\text{vib}} + S_{\text{elec}} \end{aligned} \tag{13.19}$$

For each of the partition functions the sum over allowed quantum states runs to infinity. However, since the energies become larger, the partition functions are finite. Let us examine each of the q factors in a little more detail.

13.5.1 Translational degrees of freedom

The translational degrees of freedom can be exactly separated from the other $3N - 3$ coordinates. The allowed quantum states for the translational energy are determined by placing the molecule in a “box”, i.e. the potential is zero inside the box but infinite outside. The only purpose of the box is to allow normalization of the translational wave function, i.e. the exact size is not important. The solutions to the Schrödinger equation for such a “particle in a box” are standing waves, cosine and sine functions. The energy levels are associated with a quantum number n , and depend only on the total molecular mass M .

$$\epsilon_n = \frac{n^2 h^2}{8\pi^2 M} \quad (13.20)$$

Although the energy levels are quantized, the energy difference between levels is so small that the distribution can be treated as continuous. The summation involved in the partition function can therefore be replaced by an integral (an integral is just a sum in the limit of infinitely small contributions).

$$q_{\text{trans}} = \sum_{n=0}^{\infty} e^{-\epsilon_n/kT} \approx \int_0^{\infty} e^{-\epsilon_n/kT} dn \quad (13.21)$$

Inserting the energy expression and performing the integration gives eq. (13.22).

$$q_{\text{trans}} = \left(\frac{2\pi M k T}{h^2} \right) V \quad (13.22)$$

The only molecular parameter that enters is the total molecular mass M . The volume depends on the number of particles. It is customary to work on a molar scale, in which case V is the volume of 1 mol of (ideal) gas.

13.5.2 Rotational degrees of freedom

In the lowest approximation, the rotation of a molecule is assumed to occur with a geometry that is independent of the rotational and vibrational quantum numbers. A more refined treatment allows the geometry to “stretch” with rotational energy, which may be described by adding a “centrifugal” correction, and such corrections are typically of the order of a few percent. The presence of vibrational anharmonicity will furthermore cause the effective geometry to depend on the vibrational quantum state. Within the *rigid-rotor* approximation these effects are neglected, i.e. the rotation of the molecule is assumed to occur with a fixed geometry.

The energy levels calculated from the Schrödinger equation for a diatomic “rigid rotor” are given in terms of a quantum number J running from zero to infinity, and the moment of inertia I .

$$\epsilon_J = J(J+1) \frac{h^2}{8\pi^2 I} \quad (13.23)$$

The moment of inertia is calculated from the atomic masses m_1 and m_2 and the distances r_1 and r_2 of the nuclei relative to the centre of mass.

$$I = m_1 r_1^2 + m_2 r_2^2 \quad (13.24)$$

For all molecules, except very light species such as H_2 and LiH , the moment of inertia is so large that the spacing between the rotational energy levels is much smaller than kT at ambient temperatures. As for q_{trans} , this means that the summation in eq. (13.10) can be replaced by an integral.

$$q_{\text{rot}} = \sum_{J=0}^{\infty} e^{-\epsilon_J/kT} \approx \int_0^{\infty} e^{-\epsilon_J/kT} dJ \quad (13.25)$$

Performing the integration yields eq. (13.26).

$$q_{\text{rot}} = \frac{8\pi^2 I kT}{h^2 \sigma} \quad (13.26)$$

The symmetry index σ is 2 for a homonuclear system and 1 for a heteronuclear diatomic molecule.

For a polyatomic molecule, the equivalent of eq. (13.24) is a 3×3 matrix.

$$\mathbf{I} = \begin{pmatrix} \sum_i m_i (y_i^2 + z_i^2) & -\sum_i m_i x_i y_i & -\sum_i m_i x_i z_i \\ -\sum_i m_i x_i y_i & \sum_i m_i (x_i^2 + z_i^2) & -\sum_i m_i y_i z_i \\ -\sum_i m_i x_i z_i & -\sum_i m_i y_i z_i & \sum_i m_i (x_i^2 + y_i^2) \end{pmatrix} \quad (13.27)$$

Here the coordinates are again relative to the centre of mass. By choosing a suitable coordinate transformation, this matrix may be diagonalized (Section 16.2), with the eigenvalues being the *moments of inertia* and the eigenvectors called *principal axes of inertia*.

For a general polyatomic molecule, the rotational energy levels cannot be written in a simple form. A good approximation, however, can be obtained from classical mechanics, resulting in the following partition function.

$$q_{\text{rot}} = \frac{\sqrt{\pi}}{\sigma} \left(\frac{8\pi^2 kT}{h^2} \right)^{3/2} \sqrt{I_1 I_2 I_3} \quad (13.28)$$

Here I_i are the three moments of inertia. The symmetry index σ is the order of the rotational subgroup in the molecular point group (i.e. the number of proper symmetry operations); for H_2O it is 2, for NH_3 it is 3, for benzene it is 12, etc. The rotational partition function requires only information about the atomic masses and positions (eq. (13.27)), i.e. the molecular geometry.

13.5.3 Vibrational degrees of freedom

In the lowest approximation, the molecular vibrations may be described as those of a harmonic oscillator. This can be derived by expanding the energy as a function of the nuclear coordinates in a Taylor series around the equilibrium geometry. For a diatomic molecule, the only relevant coordinate is the internuclear distance R .

$$E(R) = E(R_0) + \frac{dE}{dR}(R - R_0) + \frac{1}{2} \frac{d^2 E}{dR^2} (R - R_0)^2 + \frac{1}{6} \frac{d^3 E}{dR^3} (R - R_0)^3 + \dots \quad (13.29)$$

The first term may be taken as zero, since this is just the zero point for the energy. The second term (the gradient) vanishes since the expansion is around the equilibrium geometry. Keeping only the lowest non-zero term results in the harmonic approximation, where k is the force constant.

$$E(\Delta R) \cong \frac{1}{2} \frac{d^2 E}{dR^2} \Delta R^2 = \frac{1}{2} k \Delta R^2 \quad (13.30)$$

Including higher order terms leads to anharmonic corrections to the vibration, and such effects are typically of the order of a few percent. The energy levels obtained from the Schrödinger equation for a one-dimensional harmonic oscillator (diatomic system) are given in eq. (13.31).

$$\begin{aligned} \epsilon_n &= (n + \frac{1}{2}) h \nu \\ \nu &= \frac{1}{2\pi} \sqrt{\frac{k}{\mu}} \\ \mu &= \frac{m_1 m_2}{m_1 + m_2} \end{aligned} \quad (13.31)$$

Here n is a quantum number running from zero to infinity and ν is the vibrational frequency given in terms of the force constant k ($\partial^2 E / \partial R^2$) and the reduced mass μ .

In contrast to the translational and rotational energy levels, the spacing between vibrational energy levels is comparable to kT for temperatures around 300 K, and the summation for q_{vib} (eq. (13.10)) cannot be replaced by an integral. Due to the regular spacing, however, the infinite summation can be written in a closed form.

$$\begin{aligned} q_{\text{vib}} &= \sum_{n=0}^{\infty} e^{-\epsilon_n/kT} = e^{-h\nu/2kT} + e^{-3h\nu/2kT} + e^{-5h\nu/2kT} + \dots \\ q_{\text{vib}} &= e^{-h\nu/2kT} (1 + e^{-h\nu/kT} + e^{-2h\nu/kT} + \dots) \\ q_{\text{vib}} &= \frac{e^{-h\nu/2kT}}{1 - e^{-h\nu/kT}} \end{aligned} \quad (13.32)$$

In the infinite sum, each successive term is smaller than the previous by a constant factor ($e^{-h\nu/kT}$, which is <1), and can therefore be expressed in a closed form. Calculating the vibrational partition function for a harmonic oscillator thus requires the second derivative of the energy and the atomic masses.

For a polynuclear molecule, the force constant k is replaced by a $3N_{\text{atom}} \times 3N_{\text{atom}}$ matrix containing all the second derivatives of the energy with respect to the coordinates. By mass-weighting and transforming to a new coordinate system called the vibrational normal coordinates, this may be brought to a diagonal form (see Section 16.2.2 for details). In the vibrational normal coordinates, the $3N$ -dimensional Schrödinger equation can be separated into $3N$ one-dimensional equations, each having the form of a harmonic oscillator. Of these, three describe the overall translation and three (two for a linear molecule) describe the overall rotation, leaving $3N - 6(5)$ vibrations.

If the stationary point is a minimum on the energy surface, the eigenvalues of the force constant matrix are all positive. If, however, the stationary point is a TS, one (and only one) of the eigenvalues is negative. This corresponds to the energy being a

maximum in one direction and a minimum in all other directions. The “frequency” for the “vibration” along the eigenvector with a negative force constant will formally be imaginary, as it is the square root of a negative number (eq. (13.31)), and for a TS there are thus only $3N - 7$ vibrations.

Within the harmonic approximation, the vibrational degrees of freedom are decoupled in the normal coordinate system. Since the energy of the $3N - 6$ vibrations can be written as a sum, the partition function can be written as a product over $3N - 6$ vibrational partition functions.

$$E_{\text{vib}} = \sum_{i=1}^{3N_{\text{atom}}-6(7)} \left(n_i + \frac{1}{2}\right) h\nu_i \quad (13.33)$$

$$q_{\text{vib}} = \prod_{i=1}^{3N_{\text{atom}}-6(7)} \frac{e^{-h\nu_i/2kT}}{1 - e^{-h\nu_i/kT}}$$

The vibrational frequencies are needed for calculating q_{vib} , and can be obtained from the force constant matrix and atomic masses.

13.5.4 Electronic degrees of freedom

The electronic partition function involves a sum over electronic quantum states. These are the solutions to the electronic Schrödinger equation, i.e. the lowest (ground) state and all possible excited states. In almost all molecules, the energy difference between the ground and excited states is large compared with kT , which means that only the first term (the ground state energy) in the partition function summation (eq. (13.11)) is important.

$$q_{\text{elec}} = \sum_{i=0}^{\infty} g_i e^{-\varepsilon_i/kT} \approx g_0 e^{-\varepsilon_0/kT} \quad (13.34)$$

Defining the zero point for the energy as the electronic energy of the reactant, the electronic partition functions for the reactant and TS is given in eq. (13.35).

$$q_{\text{elec}}^{\text{reactant}} = g_0$$

$$q_{\text{elec}}^{\text{TS}} = g_0 e^{-\Delta E^\ddagger/kT} \quad (13.35)$$

The ΔE^\ddagger term is the difference in electronic energy between the reactant and TS, and g_0 is the electronic degeneracy of the (ground state) wave function. The degeneracy may be either in the spin part ($g_0 = 1$ for a singlet, 2 for a doublet, 3 for a triplet, etc.) or in the spatial part ($g_0 = 1$ for wave functions belonging to an A, B or Σ representation in the point group, 2 for an E, Δ or Φ representation, 3 for a T representation, etc.). The large majority of stable molecules have non-degenerate ground state wave functions, and consequently $g_0 = 1$.

13.5.5 Enthalpy and entropy contributions

Given the partition function, the enthalpy and entropy terms may be calculated by carrying out the required differentiations in eq. (13.18). For one mole of molecules, the results for a non-linear system are (R being the gas constant)

$$\begin{aligned}
H_{\text{trans}} &= \frac{5}{2}RT \\
H_{\text{rot}} &= \frac{3}{2}RT \\
H_{\text{vib}} &= R \sum_{i=1}^{3N-6(7)} \left(\frac{h\nu_i}{2k} + \frac{h\nu_i}{k} \frac{1}{e^{h\nu_i/kT} - 1} \right) \\
H_{\text{elec}}^{\text{reactant}} &= 0 \\
H_{\text{elec}}^{\text{TS}} &= \Delta E^\ddagger
\end{aligned} \tag{13.36}$$

$$\begin{aligned}
S_{\text{trans}} &= \frac{5}{2}R + R \ln \left(\frac{V}{N_A} \left(\frac{2\pi M k T}{h^2} \right)^{3/2} \right) \\
S_{\text{rot}} &= R \left(\frac{3}{2} + \ln \left(\frac{\sqrt{\pi}}{\sigma} \left(\frac{8\pi^2 k T}{h^2} \right)^{3/2} \sqrt{I_1 I_2 I_3} \right) \right) \\
S_{\text{vib}} &= R \sum_{i=1}^{3N-6(7)} \left(\frac{h\nu_i}{kT} \frac{1}{e^{h\nu_i/kT} - 1} - \ln(1 - e^{-h\nu_i/kT}) \right) \\
S_{\text{elec}}^{\text{reactant}} &= S_{\text{elec}}^{\text{TS}} = R \ln g_0
\end{aligned} \tag{13.37}$$

The rotational terms are slightly different for a linear molecule, and the vibrational terms will contain one vibrational contribution more.

$$\begin{aligned}
H_{\text{rot}}(\text{linear}) &= RT \\
S_{\text{rot}}(\text{linear}) &= R \left[1 + \ln \left(\frac{8\pi^2 I k T}{\sigma h^2} \right) \right]
\end{aligned} \tag{13.38}$$

The vibrational enthalpy consists of two parts, the first being a sum of $1/2 h\nu$ contributions giving the zero point energies. The second part depends on temperature, and is a contribution from molecules that are not in the vibrational ground state. This contribution goes toward zero as the temperature goes to zero where all molecules are in the ground state. Note also that the sum over vibrational frequencies runs over $3N - 6$ for the reactant(s), but only $3N - 7$ for the TS. At the TS, one of the normal vibrations has been transformed into the reaction coordinate, which formally has an imaginary frequency.

In order to calculate $\Delta G^\ddagger = G_{\text{TS}} - G_{\text{reactant}}$, we need ΔH^\ddagger and ΔS^\ddagger . $\Delta H_{\text{elec}}^\ddagger$ is directly the difference in electronic energy between the TS and reactant. Except for complicated reactions involving several electronic states of different degeneracy (e.g. singlet molecules reacting via a triplet TS), $\Delta S_{\text{elec}}^\ddagger$ is zero.

For *unimolecular* reactions $\Delta H_{\text{trans}}^\ddagger$, $\Delta H_{\text{rot}}^\ddagger$ and $\Delta S_{\text{trans}}^\ddagger$ are zero, while $\Delta S_{\text{rot}}^\ddagger$ may be slightly different from zero owing to a change in geometry (thereby changing the moments of inertia). The $\Delta H_{\text{vib}}^\ddagger$ contribution is usually a few kJ/mol negative, as there is one less vibration at the TS (lack of zero point energy). The TS is normally somewhat more ordered than the reactant, typically giving a slightly negative $\Delta S_{\text{vib}}^\ddagger$.

For *bimolecular* reactions (i.e. where the reactant is two separate molecules) $\Delta H_{\text{trans}}^\ddagger$ and $\Delta H_{\text{rot}}^\ddagger$ contribute a constant $-4RT$. The translational and rotational entropy changes are substantially negative, -30 to -50 J/mol·K, due to the fact that there are six translational and six rotational modes in the reactants but only three of each at the

TS. The six remaining degrees of freedom are transformed into the reaction coordinate and five new vibrations at the TS. These additional vibrations usually make $\Delta H_{\text{vib}}^\ddagger$ a few kJ/mol positive, and $\Delta S_{\text{vib}}^\ddagger$ positive by 5–10 J/mol·K. For bimolecular reactions, the entropy typically raises the free energy barrier by 40–60 kJ/mol, relative to the electronic energy alone.

Similarly, in order to calculate $\Delta G_0 = G_{\text{product}} - G_{\text{reactant}}$ we need ΔH_0 and ΔS_0 . The generalization for the electronic, translational and rotational contributions to ΔH^\ddagger and ΔS^\ddagger given above also holds for ΔH_0 and ΔS_0 . The considerations for a unimolecular reaction hold for reactions where the number of reactant and product molecules is the same, while the generalizations for a bimolecular reaction correspond to an addition where two reactants form a single product molecule (the reverse process being a fragmentation). The vibrational contribution to ΔH_0 and ΔS_0 for a “number-conserving” reaction is usually small, since there is the same number of vibrational modes in the reactant and product. For an addition reaction, the number of vibrational modes increases by six, and the contributions to ΔH_0 and ΔS_0 are again slightly positive, typically by a few kJ/mol and 5–10 J/mol·K.

Tables 13.1–13.3 give some examples of the magnitude of each term for two bimolecular reactions (Diels–Alder and S_N2 reactions, forming either one or two molecules as the product) and a unimolecular rearrangement (Claisen reaction).




Figure 13.4 The Diels–Alder, S_N2 and Claisen reactions

All values have been calculated at the MP2 level with the 6-31G(d) basis for the Diels–Alder and Claisen reactions, and the 6-31+G(d) basis for the S_N2 reaction. ΔH and $T\Delta S$ values are given in kJ/mol at a temperature of 300 K ($RT = 2.5$ kJ/mol), ΔS values are in J/mol·K.

Table 13.1 Diels–Alder reaction of butadiene and ethylene to form cyclohexene

	ΔH^\ddagger	ΔS^\ddagger	$-T\Delta S^\ddagger$	ΔH_0	ΔS_0	$-T\Delta S_0$
Electronic	75	0	0	–220	0	0
Vibrational	14	5	–7	32	2	–2
Rotational	–4	–11	14	–4	–13	17
Translational	–6	–35	44	–6	–35	44
Total	79	–41	51	–198	–46	58
Experimental	105	–41	52	–166	–45	56

Table 13.2 S_N2 reaction of OH[−] with CH₃F to form CH₃OH and F[−]

	ΔH^\ddagger	ΔS^\ddagger	$-T\Delta S^\ddagger$	ΔH_0	ΔS_0	$-T\Delta S_0$
Electronic	21	0	0	−105	0	0
Vibrational	9	7	−9	11	1	−2
Rotational	−3	0	0	−3	−4	5
Translational	−6	−27	34	0	0	0
Total	21	−21	26	−97	−3	4

Table 13.3 Claisen rearrangement of allyl vinyl ether to form 5-hexenal

	ΔH^\ddagger	ΔS^\ddagger	$-T\Delta S^\ddagger$	ΔH_0	ΔS_0	$-T\Delta S_0$
Electronic	98	0	0	−98	0	0
Vibrational	−5	−9	11	1	0	0
Rotational	0	0	0	0	0	0
Translational	0	0	0	0	0	0
Total	93	−9	11	−97	0	0
Experimental	125	−8	10			

It should be noted that the experimental activation enthalpy for the Diels–Alder reaction is 105 ± 8 kJ/mol,⁷ i.e. the MP2/6-31G(d) value is ~ 25 kJ/mol too low. Similarly, the calculated reaction energy of -198 kJ/mol is in rather poor agreement with the experimental value of -166 kJ/mol. The S_N2 reaction refers to the situation in the gas phase where the reactants initially form an ion–dipole complex, pass over the TS and form another ion–dipole complex. The energies given above are relative to the isolated reactants, which is the reason for the low activation energy. Note also that the rotational contribution to the reaction enthalpy is not zero; this is due to the fact that one of the reactants is a diatomic molecule, while one of the products is an atom (which has no rotational term). The MP2/6-31G(d) activation enthalpy for the Claisen reaction is again somewhat lower than the experimental value of 125 kJ/mol,⁸ while the calculated activation entropy is in good agreement with the experimental value.

In summary, to calculate rate and equilibrium constants we need to calculate ΔG^\ddagger and ΔG_0 . This can be done within the RRHO approximation if the geometry, energy and force constants are known for the reactant, TS and product. The translational and rotational contributions are trivial to calculate, while the vibrational frequencies require the full force constant matrix (i.e. all energy second derivatives), which may be a significant computational effort.

The above treatment has made some assumptions, such as harmonic frequencies and “sufficiently small” energy spacing between the rotational levels. If a more elaborate treatment is required, the summation for the partition functions must be carried out explicitly. An approximate account for vibrational anharmonicity can be obtained by using the harmonic form for the partition function (and resulting enthalpy and entropy

terms, eqs (13.36) and (13.37)), but using calculated anharmonic frequencies. The latter can be obtained from the third derivative and partial (diagonal components only) fourth-order derivatives of the energy with respect to the nuclear geometry.⁹

Many molecules have internal rotations around bonds with quite small barriers. In the above treatment, they are assumed to be described by simple harmonic vibrations, which may be a poor approximation. The calculated “vibrational frequency” for a low-barrier rotation is often close to zero, and inspection of eq. (13.36) shows that the enthalpy term in such cases approaches a constant factor of RT . The entropy term (eq. (13.37)), however, goes towards infinity as the frequency approaches zero. Calculating the energy levels and partition function for a hindered rotor is somewhat complicated,¹⁰ and is rarely done. If the barrier is very low, the motion may be treated as a free rotor, in which case it contributes a constant factor of RT to the enthalpy and $\frac{1}{2}R$ to the entropy. The enthalpy contribution is thus asymptotically correct when a low-barrier internal rotation is treated as a harmonic frequency, but the entropy term is not. Even minor inaccuracies in the calculated frequency may thus lead to large errors in the entropy contribution for small frequencies, and care must be taken in such cases. A specific problem arises in bimolecular addition reactions (or the reverse fragmentation reaction), where six translational and six rotational degrees of freedom in the reactants are transformed into three translational and three rotational degrees of freedom in the product, i.e. creating six new internal degrees of freedom. At the TS, several of these often correspond to low-barrier internal rotations, which may be problematic to treat as harmonic vibrations.¹¹ Although the vibrational entropy reaches a value of $\frac{1}{2}R$ already for a harmonic frequency of around 400 cm^{-1} , the difference relative to a hindered internal rotor only becomes significant for frequencies below 100 cm^{-1} and rotational barriers comparable to RT ($\sim 3\text{ kJ/mol}$ at room temperature).

It should also be noted that the thermodynamic contributions in eqs (13.36) and (13.37) are calculated using the most common atomic isotopes, while the experimental quantities of course represent an ensemble of molecules containing a statistical mixture of isotopomers. It is straightforward but tedious to construct the thermodynamic contributions corresponding to a mixture of molecules with different atomic isotopes.¹² Since the resulting changes are substantially smaller than the error due to neglect of vibrational anharmonicity, such improvements are usually not considered.

As can be seen from Tables 13.1–13.3, the electronic energy difference between the reactant/TS and reactant/product is the most important contribution to ΔG^\ddagger and ΔG_0 . The electronic energy is furthermore the most difficult to calculate accurately. Let us consider three cases.

- (1) The error in $\Delta E^\ddagger/\Delta E_0$ is $\sim 50\text{ kJ/mol}$. It is clear that spending significant amounts of computer time in order to include vibrational, rotational and translational corrections has little value.
- (2) The error in $\Delta E^\ddagger/\Delta E_0$ is $\sim 5\text{ kJ/mol}$. The corrections from vibrations, rotations and translation now become important, and should be included. However, sophisticated treatments such as anharmonic vibrations are unimportant.
- (3) The error in $\Delta E^\ddagger/\Delta E_0$ is $\sim 0.5\text{ kJ/mol}$. Corrections from vibrations, rotations and translation are clearly necessary. Explicit calculation of the partition functions for anharmonic vibrations, and internal rotations may be considered. However, at this

point other factors also become important for the activation energy. These include for example:

- (a) The position of the TS has been assumed to be at the maximum on the electronic energy surface, whereas in reality it should be at the maximum on the ΔG surface. This would include entropy effects and thus allow the position of the TS to depend on temperature. Such treatments are referred to as *Variational Transition State Theory* (VTST)¹³ and are important for reaction with small (or zero) enthalpy barriers, such as recombination of radicals or carbene additions.¹⁴
- (b) The possibility of re-crossings and tunnelling (which requires a quantum description of the nuclear motion) should be included in order to produce a transmission coefficient. Tunnelling can be estimated from the imaginary frequency at the TS, but an accurate estimate requires elaborate calculations. The re-crossing effect requires simulation of the dynamics of the reaction, again a substantial computational problem.

Calculating the electronic barrier with an accuracy of ~ 0.5 kJ/mol is only possible for very simple systems. An accuracy of ~ 5 kJ/mol is usually considered a good, but hard to get, level of accuracy. The situation is slightly better for relative energies of stable species, but a ~ 5 kJ/mol accuracy still requires significant computational effort. Thermodynamic corrections beyond the rigid-rotor/harmonic vibrations approximation are therefore rarely performed.

A prediction of $\Delta E^\ddagger/\Delta E_0$ to within ~ 0.5 kJ/mol may produce a $\Delta G^\ddagger/\Delta G_0$ accurate to maybe ~ 1 kJ/mol. This corresponds to an error of a factor of ~ 1.5 (at $T = 300$ K) in the rate/equilibrium constant, which is poor compared with what is routinely obtained by experimental techniques. Calculating $\Delta G^\ddagger/\Delta G_0$ to within ~ 5 kJ/mol is still only possible for fairly small systems. This corresponds to predicting the absolute rate constant, or the equilibrium distribution, to within a factor of 10. Theoretical calculations are therefore not very useful for predicting absolute rate or equilibrium constants. Relative rates, however, are somewhat easier. Often the interest is not in how fast a certain product is formed, but rather on what the rate difference is between two reactions. The absolute rate (only) influences how long the total reaction time will be, or how high the temperature should be. Rate differences, on the other hand, determine what the ratio between products is. When comparing calculated activation parameters for similar reactions, one can always hope for some “cancellation of errors”. Theoretical methods are most useful for predicting and rationalizing different reaction pathways, not in predicting absolute rates.

The activation enthalpies and entropies in principle depend on temperature (eqs (13.36) and (13.37)), but only weakly so, and for a limited temperature range they may be treated as constants. Obtaining these quantities experimentally is possible by measuring the reaction rate as a function of temperature, and plotting $\ln(k_{\text{rate}}/T)$ against T^{-1} .

$$k_{\text{rate}} = \frac{kT}{h} e^{-\Delta G^\ddagger/RT}$$

$$\ln\left(\frac{K_{\text{rate}}}{T}\right) = \ln\left(\frac{k}{h}\right) + \frac{\Delta S^\ddagger}{R} - \frac{\Delta H^\ddagger}{RT} \quad (13.39)$$

Such plots should produce a straight line with the slope being equal to $-\Delta H^\ddagger/R$ and the intercept equal to $\ln(k/h) + \Delta S^\ddagger/R$. As the available temperature range often is $\sim 100^\circ\text{C}$, the error in ΔH^\ddagger will typically be $0.5\text{--}2\text{ kJ/mol}$. The activation entropy is determined by extrapolating outside the data points to $T = \infty$ ($1/T = 0$), and is usually somewhat less well defined; a typical error may be $5\text{ J/mol}\cdot\text{K}$.

Experimentalists often analyze their data in terms of an Arrhenius expression instead of the TST expression eq. (13.39) by plotting $\ln(k_{\text{rate}})$ against T^{-1} .

$$k_{\text{rate}} = A e^{-\Delta E^\ddagger/RT}$$

$$\ln(k_{\text{rate}}) = \ln(A) - \frac{\Delta E^\ddagger}{RT} \quad (13.40)$$

The connection with the TST expression (13.39) may be established from the definition in eq. (13.40) of the activation energy.

$$\Delta E^\ddagger = RT^2 \left(\frac{d \ln k_{\text{rate}}}{dT} \right)_V \quad (13.41)$$

This produces the relationship shown in eq. (13.42).

$$\Delta H^\ddagger = \Delta E^\ddagger - (1 - \Delta n)RT$$

$$\Delta S^\ddagger = R \left[\ln A - \ln \left(\frac{kT}{h} \right) - (1 - \Delta n) + \Delta n(RT) \right] \quad (13.42)$$

Here Δn is the change in the number of molecules from the reactant to the TS, i.e. $\Delta n = 0$ for a unimolecular reaction, -1 for a bimolecular reaction, etc. For a solution phase reaction Δn is approximately 0.

Note that for a reaction taking place by multiple reaction paths (e.g. conformational TS's), the *observed* activation energy is obtained from the observed rate constant, which is a sum over individual rate constants.

$$k_i = \frac{kT}{h} e^{-\Delta G_i^\ddagger/RT}$$

$$k_{\text{observed}} = \sum_i k_i \quad (13.43)$$

$$e^{-\Delta G_{\text{observed}}^\ddagger/RT} = \sum_i e^{-\Delta G_i^\ddagger/RT}$$

The presence of multiple reaction paths with similar activation energies will thus result in an effective activation energy that is *lower* than the activation energy of the lowest TS.

13.6 Condensed Phases

For a single molecule in the rigid-rotor harmonic-oscillator approximation, the (quantum) energy states are sufficiently regular to allow an explicit construction of the partition function. For a collection of many interacting particles (condensed phase), the relevant energy states are those describing the vibrations, and translation and rotation of molecules relative to each other. For such systems, the energy levels are

not only numerous but also so irregularly spaced that it is impossible to derive them directly from molecular quantities. It is consequently not possible to construct the partition function explicitly. It is, however, possible to estimate *derivatives* of Q and *differences* in Q by a representative sample of the system. Condensed phases can be modelled by periodic boundary conditions, and configurations generated by either molecular dynamics or Monte Carlo procedures, as discussed in Section 14.1.

We can derive formal expressions for U and A from eqs (13.14) and (13.16) by using the fact that $\partial \ln Q / \partial T = Q^{-1} \partial Q / \partial T$.

$$\begin{aligned} U &= \frac{kT^2}{Q} \frac{\partial Q}{\partial T} = \frac{kT^2}{Q} \frac{\partial}{\partial T} \left(\sum_i e^{-E_i/kT} \right) = \sum_i E_i (Q^{-1} e^{-E_i/kT}) \\ A &= -kT \ln Q = -kT \ln \left(\sum_i e^{-E_i/kT} \right) \end{aligned} \quad (13.44)$$

The Boltzmann probability function P can be written either in a discrete energy representation or in a continuous phase space formulation.

$$\begin{aligned} P(E_i) &= Q^{-1} e^{-E_i/kT} \\ P(\mathbf{r}, \mathbf{p}) &= Q^{-1} e^{-E(\mathbf{r}, \mathbf{p})/kT} \end{aligned} \quad (13.45)$$

Here Q^{-1} is a normalization factor. The internal energy U in eq. (13.44) can thus be written as in eq. (13.46).

$$\begin{aligned} U &= \sum_i E_i P(E_i) \\ U &= \int E(\mathbf{r}, \mathbf{p}) P(\mathbf{r}, \mathbf{p}) d\mathbf{r} d\mathbf{p} \end{aligned} \quad (13.46)$$

Eq. (13.46) shows that U is simply a sum of energies weighted by the probability of being in that state, i.e. U is the average (potential) energy of the system. Since high-energy states occur with a low probability, only the low-energy region of the phase space is important for the internal energy.

A similar expression may be derived for A by substituting 1 with $e^{E/kT} e^{-E/kT}$ in eq. (13.44) and summing over all N_{states} .

$$\begin{aligned} A &= -kT \ln Q = kT \ln \left(\frac{1}{Q} \right) = kT \ln \left(\frac{\sum_i e^{-E_i/kT} e^{E_i/kT}}{N_{\text{states}} Q} \right) \\ A &= -kT \ln(N_{\text{states}}) + kT \ln \left(\sum_i e^{E_i/kT} P(E_i) \right) \end{aligned} \quad (13.47)$$

The $\ln(N_{\text{states}})$ term is constant and corresponds to a change of the zero point, and can consequently be neglected. Alternatively, A may be written as an integral over phase space.

$$A = kT \ln \int e^{E(\mathbf{r}, \mathbf{p})/kT} P(\mathbf{r}, \mathbf{p}) d\mathbf{r} d\mathbf{p} \quad (13.48)$$

In contrast to U (eq. (13.46)), the Helmholtz free energy A depends exponentially on the energy, i.e. although high-energy states occur infrequently, they contribute significantly owing to the exponential weighting factor. Alternatively stated, U depends only on the *derivative* of Q , while A depends *directly* on Q .

It is not possible to carry out the summation over *all* states, or equivalently integrate over *all* phase space in eqs (13.46)–(13.48). The U and A values could in principle be calculated by sampling the phase space in a random fashion (Monte Carlo type integration), but such an approach will suffer from an extremely slow convergence as the large majority of points will have high energies, and consequently contribute with a very small probability. If, however, a *representative* collection of configurations can be generated, the sum over all states can be approximated by an average over a finite set of configurations. Representative here means that the number of configurations with a given energy is proportional to that given by the Boltzman distribution, and that all “important” parts of the phase space are sampled. For a finite number of points M , it is possible to calculate the average value of a given property X according to eq. (13.49), where the points can be denoted either by their energies or by their positions and momentum.

$$\langle X \rangle_M = \frac{1}{M} \sum_{i=1}^M X(E_i) = \frac{1}{M} \sum_{i=1}^M X(\mathbf{r}_i, \mathbf{p}_i) \quad (13.49)$$

In a typical simulation the number of sampling points is perhaps $\sim 10^6$, which represents only an infinitesimal fraction of the $6N_{\text{atom}}$ -dimensional phase space (a rough 10 point sampling in each dimension would give 10^{6N} points). As already mentioned, however, the vast majority of the huge phase space is high in energy and is not accessible at normal temperatures. Consider for example placing 1000 water molecules at random in a box with a dimension corresponding to a density of 1 g/cm^3 . If any two water molecules have a significant overlap, there will be a large repulsive interaction, and therefore a vanishing probability of such a configuration occurring. Placing all 1000 water molecules in the box without any two molecules having an overlap is difficult, and will essentially never occur by a random placement. Starting from an energy-minimized structure and allowing the system to evolve by a molecular dynamics algorithm, however, will *only* sample those configurations where no serious molecular overlaps occur, i.e. the important low-energy region.

The “magic” in simulations is generating an ensemble that yields a good representation of the “important” phase space for the given property. A collection of configurations is called an *ensemble*, and eq. (13.49) is called an ensemble average, with the subscript indicating what is being averaged over. There are two main techniques for generating an ensemble, Monte Carlo and molecular dynamics, which are discussed in Chapter 14. These methods are based on the *ergodic hypothesis* (which can be proven rigorously only for a hard-sphere gas), which makes the assumption that the average obtained by following a small number of particles over a long time is equivalent to averaging over a large number of particles for a short time. Taken to the limit, this implies that a time average over a single particle is equivalent to an average of a large number of particles at any given time snapshot, i.e. time-averaging is equivalent to ensemble-averaging.

$$\langle X \rangle = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau X(t) dt = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M X_i \quad (13.50)$$

Alternatively stated, the ergodic hypothesis implies that no matter where a system is started, it is possible to get to any other point in phase space. For U and A , this leads to the following expressions.

$$\begin{aligned}\langle U \rangle_M &= \frac{1}{M} \sum_{i=1}^M E_i = \langle E \rangle_M \\ \langle A \rangle_M &= kT \ln \left(\frac{1}{M} \sum_{i=1}^M e^{E_i/kT} \right) = kT \ln \langle e^{E_i/kT} \rangle_M\end{aligned}\quad (13.51)$$

In general, a macroscopic observable can be calculated as an average over a corresponding microscopic quantity. The average value of for example U calculated from eq. (13.51) has a statistical uncertainty $\sigma(U)$, which is the square root of the variance σ^2 (eq. (17.2), making the approximation $M - 1 \approx M$ for large samples).

$$\sigma^2(U) = \frac{1}{M} \sum_{i=1}^M (E_i - \langle E \rangle)^2 \quad (13.52)$$

The statistical uncertainty is therefore inversely proportional to the square root to the number of sampling points M .

$$\sigma(X) \propto \frac{1}{\sqrt{M}} \quad (13.53)$$

Increasing the sample size from 1000 to 4000 thus reduces the standard deviation by a factor of 2. How well the calculated average (from eq. (13.49)) resembles the “true” value, however, depends on whether the ensemble is representative. If a large number of points are collected from a small part of the phase space, the property may be calculated with small *statistical* error, but a large *systematic* error (i.e. the value may be precise, but inaccurate). As it is difficult to establish that the phase space is adequately sampled, this can be a very misleading situation, i.e. the property *appears* to have been calculated accurately but may in fact be significantly in error. Different parts of the phase space may furthermore be important for different properties. An ensemble that gives an accurate value for one property may not necessarily be suitable for another property. Energy properties, such as U , H and C_v , depend on the derivative of Q for which the low-energy region of the phase space is important, while entropic properties, such as A , S and G , depend directly on Q , where the whole phase space is important. Since Monte Carlo and molecular dynamics techniques preferentially sample the low-energy region, it is computationally difficult to achieve a reasonable statistical error for entropic quantities.

With standard MC or MD simulations, the ensemble reflects the temperature and only configurations that are accessible at the given temperature are represented to any significant extent. This makes it impossible to calculate the absolute value of the entropy, but it is possible to calculate *differences* in entropy properties. Using the Helmholtz free energy for illustration, we can consider two systems A and B described by two different energy functions E^A and E^B . The energy difference is given in eq. (13.54) and involves a ratio of the corresponding partition functions.

$$A_A - A_B = -kT(\ln Q_A - \ln Q_B) = -kT \ln \frac{Q_A}{Q_B} \quad (13.54)$$

Analogously to eqs (13.51), the difference can be evaluated as an ensemble average.

$$\langle \Delta A \rangle_M = kT \ln \left(\frac{1}{M} \sum_{i=1}^M e^{(E_i^B - E_i^A)/kT} \right) = kT \ln \langle e^{\Delta E_{BA}/kT} \rangle_M \quad (13.55)$$

The important difference is that the exponential now involves an energy difference $E^B - E^A$. Provided that this is sufficiently small compared with kT , the ensemble average will show much better convergence than the absolute entropy of either system. If the energy difference is large compared with kT , we may introduce intermediates states between A and B that can be described in term of a coupling parameter λ ($0 \leq \lambda \leq 1$). The simplest approach involves a linear interpolation but more complicated connections can also be used.¹⁵

$$E_\lambda = \lambda E_A + (1 - \lambda) E_B \quad (13.56)$$

The sampling can then be performed for each value of λ , and all the intermediate results added together to provide the difference $S_A - S_B$. It should be noted that a system corresponding to an intermediate value of λ does not necessarily represent an actual physically realizable system. If, for example, the objective is to calculate the solvation entropy difference between acetone and propane in a solvent, a value of $\lambda = 0.5$ corresponds to a “molecule” with “half” a carbonyl oxygen and two “half” hydrogen atoms on the central carbon. Such artificial intermediate systems do not represent special problems in terms of calculation.

The preference of standard MC or MD methods for sampling the low-energy region of a surface is a result of the way one configuration is propagated to the next. In MC methods the probability for accepting a trial move depends on the ratio of the change in energy relative to the temperature, while MD methods have a velocity (direction and magnitude) depending on the temperature. Recently an alternative MC method has been proposed where the transition probability instead depends on the inverse density of states, rather than the energy.¹⁶ This in principle makes it possible to simulate an ensemble that provides a uniform coverage of the whole phase space, and therefore allows calculation of absolute values of entropies and free energies. Since the density of states is unknown *a priori*, the method requires a sequence of simulations where the density of state diagram is gradually constructed and refined.

The main problem in estimating thermodynamic quantities from simulations is the assumption that the generated set of configurations forms a representative set. In practice, this is impossible to guarantee or verify, making simulations somewhat of a “black art”. For configurations generated by molecular dynamics, a typical simulation time is of the order of nanoseconds, and it is clear that this is a much too short a timespan to adequately sample all the phase space. There is thus a real risk of a simulation being trapped in a small volume of the phase space during the whole simulation, and thereby providing a misleading sampling. In order to evaluate the sensitivity of the results, several simulations are often performed with different starting conditions.

References

1. H. Eyring, *J. Chem. Phys.*, **3** (1935), 107.
2. M. A. Robb, M. Garavelli, M. Olivucci, F. Bernardi, *Rev. Comp. Chem.*, **15** (2000), 87.

3. K. J. Laidler, *Chemical Kinetics*, Harper and Row, 1987; J. I. Steinfeld, J. S. Francisco, W. L. Hase, *Chemical Kinetics and Dynamics*, Prentice-Hall, 1989.
4. B. K. Carpenter, *Angew. Chem. Int. Ed.*, **37** (1998), 3340.
5. M. B. Reyes, B. K. Carpenter, *J. Am. Chem. Soc.*, **122** (2000), 10163.
6. I. N. Levine, *Physical Chemistry*, McGraw-Hill, 1983; K. Lucas, *Applied Statistical Thermodynamics*, Springer-Verlag, 1991.
7. K. N. Houk, R. J. Loncharich, J. F. Blake, W. L. Jorgensen, *J. Am. Chem. Soc.*, **111** (1989), 9172.
8. F. W. Schuler, G. W. Murphy, *J. Am. Chem. Soc.*, **72** (1950), 3155.
9. V. Barone, *J. Chem. Phys.*, **120** (2004), 3059.
10. Y.-Y. Chuang, D. G. Truhlar, *J. Chem. Phys.*, **112** (2000), 1221.
11. L. Masgrau, À. González-Lafont, J. M. Lluch, *J. Comp. Chem.*, **24** (2003), 701.
12. F. Jensen, *Mol. Phys.*, **101** (2003), 2315.
13. M. S. Gordon, D. G. Truhlar, *Science*, **249** (1990), 491; S.-H. Yang, I. Hristov, P. Fleurat-Lessard, T. Ziegler, *J. Phys. Chem. A*, **109** (2005) 197.
14. A. E. Keating, S. R. Merrigan, D. A. Singleton, K. N. Houk, *J. Am. Chem. Soc.*, **121** (1999), 3933.
15. D. Frenkel, B. Smith, *Understanding Molecular Simulations*, Academic Press, 1996.
16. F. Wang, D. L. Landau, *Phys. Rev. Lett.*, **86** (2001), 2050.

14 Simulation Techniques

The analysis of a potential energy surface by locating the minima and saddle points (Chapter 12) corresponds to modelling the system at a temperature of 0 K, where all molecules are in their ground electronic, vibrational and rotational states. The effects of a finite temperature can be incorporated by means of the statistical mechanics methods discussed in Chapter 13. For a system of non-interacting molecules (ideal gas), the partition function can be evaluated quite accurately by the rigid-rotor harmonic-oscillator approximation from relatively simple quantities for the isolated molecule (geometry and vibrational frequencies). Similar approaches are possible for crystalline solid states, where the translational symmetry implies that only properties for the unit cell are required for describing the whole system. For other systems, most notably liquids and solutions, the macroscopic quantities derived from the partition function must be estimated from a representative sampling of the phase space. *Simulation* refers to methods aimed at generating a representative sampling of a system at a finite temperature.¹

Electronic structure methods are typically used for solving the Schrödinger equation for a single or a few molecules, infinitely removed from all other molecules. Physically this corresponds to the situation occurring in the gas phase under low pressure (vacuum). Experimentally, however, the majority of chemical reactions are carried out in solution. Biologically relevant processes also occur in solution, aqueous systems with rather specific pH and ionic conditions. Most reactions are both qualitatively and quantitatively different under gas and solution phase conditions, especially those involving ions or polar species. Molecular properties are also sensitive to the environment. Simulations are therefore intimately related with describing solute–solvent interactions, but such effects can also be modelled with less rigorous methods.

There are two major techniques for generating an ensemble: Monte Carlo and molecular dynamics.

In *Monte Carlo* (MC) methods,² a sequence of points in phase space is generated from an initial geometry by adding a random “kick” to the coordinates of a randomly chosen particle (atom or molecule). The new configuration is accepted if the energy decreases and with a probability of $e^{-\Delta E/kT}$ if the energy increases. This *Metropolis* procedure³ ensures that the configurations in the ensemble obey a Boltzmann

distribution, and the possibility of accepting higher energy configurations allows MC methods to climb uphill and escape from a local minimum. In order to have a reasonable acceptance ratio, however, the step size must be fairly small. This effectively means that even a few millions MC steps (a typical computational limit) only explore the local phase space around the starting geometry.

Monte Carlo methods generate configurations in a random fashion, and the Metropolis selection procedure ensures that a proper ensemble is generated. The geometry perturbation in each step may be “non-physical”, which is actually an advantage since two consecutive geometries may be separated by a high energy barrier. MC methods thus have the possibility of “tunnelling” between energetically separated regions of phase space, thereby giving a better coverage. The perturbations may be carried out in both internal and Cartesian coordinates, and it is quite easy to freeze out certain degrees of freedom, as for example sampling only the torsional angle space. MC methods are inherently non-deterministic, as each configuration only depends on the previous point and a few random numbers, and two simulations starting from the same geometry will not generate the same sampling since the random numbers will be different. MC simulations require only the ability to evaluate the energy of the system, which may be advantageous if calculating the first derivative is difficult or time-consuming. Furthermore, since only a single particle is moved in each step, only the energy changes associated with this move must be calculated, not the total energy for the whole system. A disadvantage of MC methods is the lack of the time dimension and atomic velocities, and they are therefore not suitable for studying time-dependent phenomena or properties depending on momentum.

Molecular Dynamics (MD) methods generate a series of time-correlated points in phase space (a *trajectory*) by propagating a starting set of coordinates and velocities according to Newton’s second equation by a series of finite time steps. A typical time step is $\sim 10^{-15}$ s and a simulation involving 10^6 steps thus “only” covers $\sim 10^{-9}$ s. This is substantially shorter than many important phenomena, and MD methods, in analogy with MC, tend to only sample the region in phase space close to the starting condition. Furthermore, MD methods simulate the physical evolution of configurations and can easily become trapped in energy wells.

MD simulations are cumbersome to run in anything but Cartesian coordinates, and it is somewhat difficult to enforce constraints on the system. MD simulations require small time steps and tend to spend a significant effort describing relatively unimportant bond stretching and angle bending motions. Furthermore, the ability to climb over energy barriers is limited, as any uphill motion will generate a force trying to pull the system back towards the minimum. MD is in principle deterministic, and starting two simulations with the exact same initial coordinates and velocities should give the same trajectory. Even slight differences ($\sim 10^{-8}$) in the starting conditions, however, rapidly lead to uncorrelated trajectories within a few thousand time steps owing to an exponential divergence. Furthermore, the numerical errors generated in each time step will gradually add up to become significant. As different computers (and compilers) produce different round-off errors, this means that MD simulations in practice are non-deterministic and exhibit chaotic behaviour on timescales longer than ~ 50 ps. MD simulations implicitly have both atomic velocities and time dependence, and are thus suitable for modelling for example transport phenomena and diffusion. Running an MD simulation requires the ability to calculate the force (first derivative of the energy)

on all particles in the system in addition to the energy. For parameterized energy functions, such as those used in force field methods, this is not a limitation as forces can be calculated almost as easily as the energy. Furthermore, since all particles are moved in each step, the whole energy function (and gradient) must be recomputed at each step.

The inherently non-deterministic nature of MC methods and the non-deterministic behaviour of actual MD simulations might be considered as potential problems. In reality the only concern is generating a representative sample of the phase space, and chaotic behaviour may actually help in obtaining a more complete sampling. The random and chaotic elements in simulations, however, make troubleshooting and identification of programming bugs somewhat more problematic than for many other types of computer programs. Verifying that a new simulation program is valid cannot easily be done by comparing exact numbers with another program, as even running the same program on different types of machines may produce different results. Programming bugs that produce small (systematic) errors may thus be swamped by the statistical errors inherent in all simulations and so escape detection. Development of simulation packages must therefore be done with care, involving monitoring many different quantities to ensure that the implementation is valid.⁴

The result of a simulation or an experiment involves averaging over both the number of molecules and time, but usually with significantly different averaging lengths, and it is not completely obvious that the calculated quantities are directly comparable with the experiments. An IR spectrum, for example, records averages over a sample containing perhaps 10^{18} molecules over the timeframe of perhaps 10^{-14} s (the interaction time of radiation with molecules), i.e. essentially a snapshot of the quantum states for a large selection of molecules. A simulation on the other hand, may follow the molecular motions of perhaps 10^3 molecules for 10^{-9} s. The *ergodic hypothesis* makes the assumption that the average obtained by following a small number of particles over a long time is equivalent to averaging over a large number of particles for a short time. Taken to the limit, this implies that a time average over a single particle is equivalent to an average of a large number of particles at any given time snapshot, i.e. time-averaging is equivalent to ensemble-averaging.

$$\langle X \rangle = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^{\tau} X(t) dt = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M X_i \quad (14.1)$$

Alternatively stated, the ergodic hypothesis implies that no matter where a system is started, it is possible to get to any other point in phase space. MC techniques perform an ensemble average, while MD performs a time average.

A simulation can be characterized by quantities such as volume (V), pressure (P), total energy (E), temperature (T), number of particles (N), chemical potential (μ), etc., but not all of these are independent. For a constant number of particles, either the volume or the pressure can be fixed, but not both. Similarly, either the total energy or the temperature can be fixed, but not both, and a constant chemical potential is incommensurable with a constant number of particles. The ensemble is labelled according to the fixed quantities, as shown in Table 14.1, with the remainder being derived from the simulation data, and thus displaying a statistical fluctuation.

An MC simulation employs the temperature as the parameter for deciding acceptance or rejection of trial moves, and MC simulations are therefore naturally of the

Table 14.1 Constants in different ensembles, and corresponding equilibrium states

N	P	V	T	E	μ	Acronym	Equilibrium	Name
×		×	×			NVT	A has minimum	Canonical
×		×		×		NVE	S has maximum	Micro-canonical
×	×		×			NPT	G has minimum	Isothermal-isobaric
		×	×		×	$VE\mu$	(PV) has maximum	Grand canonical

N = number of particles; P = pressure; V = volume; T = temperature; E = energy; μ = chemical potential; A = Helmholtz free energy; S = entropy; G = Gibbs free energy

NVT type. An MD simulation, on the other hand, preserves energy and is therefore naturally of the NVE type, but other ensembles for both MC and MD can be generated by the techniques described in Section 14.2.2. Table 14.2 summarizes some of the differences between MC and MD.

Table 14.2 Differences between Monte Carlo and molecular dynamics methods

Property	MC	MD
Basic information needed	Energy	Gradient
Particles moved in each step	One	All
Coordinates	Any	Cartesian
Constraints	Easy	Difficult
Atomic velocities	No	Yes
Time dimension	No	Yes
Deterministic	No	(Yes)
Sampling	Non-physical	Physical
Natural ensemble	NVT	NVE

It is probably no surprise that hybrid MC/MD methods have been devised, trying to capture the best of both methods.⁵ These combined methods typically perform an MD simulation with an occasional MC step thrown in, in order to give a better coverage of the phase space. Alternatively, a trial step may be generated by an MD recipe, using a somewhat larger time step than for pure MD, and this trial step is then accepted or rejected based on an MC criterion.

14.1 Monte Carlo Methods

One of the advantages of Monte Carlo methods is the ease with which they can be implemented in computer programs. The heart of the algorithm is a random number generator, and the ability to calculate the energy of the system for a given set of coordinates. Although truly random numbers are difficult to come by, several implementations of pseudo-random number generators are available. A pseudo-random number generator indicates a computer implementation of an algorithm that produces a sequence of seemingly random numbers, but the sequence is repeated (exactly) after some period. A good pseudo-random number generator is characterized by having a

long periodicity, and within this periodicity the numbers do not show any systematic correlation with each other. As long as the simulation only uses numbers from within one period, the random aspect is fulfilled and the simulation data should be valid.

An MC simulation starts from a suitable set of coordinates for all the particles. The set of coordinates is perturbed in a random fashion and the new geometry is accepted as a starting point for the next perturbing step if it is lower in energy than the current. If the new geometry is higher in energy, the Boltzmann factor $e^{-\Delta E/kT}$ is calculated and compared with a random number between 0 and 1. If $e^{-\Delta E/kT}$ is larger than this number the new geometry is accepted, otherwise the old configuration is added to the sampling (again), and a new perturbing step is attempted.

The main variation of MC methods is how the perturbing step is done. For a system composed of spherical particles (atoms), the only variables are the centre of mass of each particle, and the trial moves are simple translations of particles. For rigid non-spherical particles, the three rotational degrees of freedom must also be sampled, while for flexible molecules, it is usually also of interest to sample the internal degrees of freedom (conformations, vibrations). The latter can be done in Cartesian coordinates, or selectively in for example only the torsional variables.

A key point in MC methods is to ensure that the chain of configurations arises from a *symmetric* probability decision. Symmetric in this context means that each step is reversible, i.e. the probability of undoing a step by the next move is equal to the probability of generating the step, sometimes also called the *detailed balance* condition. If this is not fulfilled, the properties derived from the resulting ensemble can (but do not necessarily) display systematic errors, which are usually hard to detect. Generating random moves corresponding only to translations in the positive direction, rather than both positive and negative directions, will almost be sure to lead to artefacts, although one might think that this would be acceptable in a system subjected to periodic boundary conditions. The detailed balance condition is obeyed when a single random particle is subjected to a single random perturbation in each step, but this is not the case if a random perturbing step is applied sequentially to all the particles. Nevertheless, it has been shown that the sequential update obeys a weaker balance condition and does in fact generate a proper Boltzmann distribution.⁶ From a computational point of view, a procedure where only one particle is moved in each trial step is usually more efficient than a trial step consisting of moving all particles. When only a single particle is moved, only the change in the energy related to this particle is required, not the whole energy function. This makes the evaluation of each trial move somewhat faster than a single time step in an MD simulation. Furthermore, if many (all) particles are allowed to move in each trial step, the acceptance ratio usually becomes prohibitively small, unless very small perturbations are selected.

The size of the perturbing step is an important control parameter. A small step will give a high acceptance ratio but only a slow change of configurations. A large step, on the other hand, gives a low acceptance ratio and therefore a sampling consisting of only a few relatively widely distributed points in the configuration space. The optimum step size can formally be defined as the one that gives the fastest convergence of a given property for a given amount of computer time, but this is difficult to translate into an optimum acceptance ratio. Lacking a more objective criterion, a heuristic acceptance ratio around 0.5 is usually selected, although slightly smaller ratios in many cases may give a better sampling.

The ability to generate non-physical moves means that attention must be paid to the molecular stereochemistry. A random move of atoms makes it possible to invert the configuration of a chiral atom, a process that in reality may require large amounts of energy, but one that can easily be generated by moving a single atom. A Monte Carlo procedure must therefore be able to detect such chirality changes and reject such moves.

The procedure of making random moves of a single or several (all) particles gives MC methods a drawback for describing correlated motions. Exploring the conformational space of a larger molecule such as a protein in a solvent is inefficient, since several simultaneous perturbations of torsional angles are required for generating acceptable conformational changes. Such correlated movements are difficult to generate by random perturbations in either Cartesian or internal coordinates, and almost impossible if only single particle movements are employed in each trial step. MC methods are therefore best for exploring the translational and rotational space for relatively small molecules, such as a solvent or solution, and internal degrees of freedom for small molecules.

14.1.1 *Generating non-natural ensembles*

A standard MC simulation generates an *NVT* ensemble, i.e. the pressure and energy will fluctuate. It is quite easy to generate other types of ensembles by MC methods, the most important being the *NPT* ensemble, since this is directly related to most experimental conditions. A constant pressure necessarily means that the volume must be able to change. For simulating an *NPT* ensemble, the total volume of the system is treated as an additional variable and subjected to random perturbations.⁷ The acceptance criterion for a volume change is the same as for particle moves, except that the energy change is augmented with two additional terms, i.e. $\Delta E \rightarrow \Delta E + P\Delta V - NkT\ln(1 + \Delta V/V)$.

14.2 Time-Dependent Methods

At a finite temperature, the average kinetic energy is directly related to the temperature and the molecule(s) explores a part of the surface with energies lower than the typical kinetic energy. One possible way of simulating the behaviour at a finite temperature is by allowing the system to evolve according to the relevant dynamical equation (Section 1.4). For nuclei, this is normally Newton's second law, although the (nuclear) Schrödinger equation must be used for including quantum effects, such as zero point vibrational energy and tunnelling. A dynamics simulation is also required if the interest is in studying time-dependent phenomena, such as transport, and the results of a simulation can yield information about the spectral properties, such as the IR spectrum.

A dynamics simulation requires a set of initial coordinates and velocities, and an interaction potential (energy function). For a short time step, the interaction may be considered constant, allowing a set of updated positions and velocities to be estimated, at which point the new interaction can be calculated. By taking a (large) number of (small) time steps, the time behaviour of the system can be obtained. Since the phase space is huge, and the fundamental time step is short, the simulation will only explore

the region close to the starting point, and several different simulations with different starting conditions are required for estimating the stability of the results.

14.2.1 Molecular dynamics methods

Nuclei are heavy enough that they, to a good approximation, behave as classical particles and the dynamics can thus be simulated by solving Newton's second equation, $\mathbf{F} = m\mathbf{a}$, which in differential form can be written as in eq. (14.2).

$$-\frac{dV}{d\mathbf{r}} = m \frac{d^2\mathbf{r}}{dt^2} \quad (14.2)$$

Here V is the potential energy at position \mathbf{r} . The vector \mathbf{r} contains the coordinates for all the particles, i.e. in Cartesian coordinates it is a vector of length $3N_{\text{atom}}$. The left-hand side is the negative of the energy gradient, also called the force (\mathbf{F}) on the particle(s).

Given a set of particles with positions \mathbf{r}_i , the positions a small time step Δt later are given by a Taylor expansion.

$$\begin{aligned} \mathbf{r}_{i+1} &= \mathbf{r}_i + \frac{\partial \mathbf{r}}{\partial t}(\Delta t) + \frac{1}{2} \frac{\partial^2 \mathbf{r}}{\partial t^2}(\Delta t)^2 + \frac{1}{6} \frac{\partial^3 \mathbf{r}}{\partial t^3}(\Delta t)^3 + \dots \\ \mathbf{r}_{i+1} &= \mathbf{r}_i + \mathbf{v}_i(\Delta t) + \frac{1}{2} \mathbf{a}_i(\Delta t)^2 + \frac{1}{6} \mathbf{b}_i(\Delta t)^3 + \dots \end{aligned} \quad (14.3)$$

The velocities \mathbf{v}_i are the first derivatives of the positions with respect to time ($d\mathbf{r}/dt$) at time t_i , the accelerations \mathbf{a}_i are the second derivatives ($d^2\mathbf{r}/dt^2$) at time t_i , the hyperaccelerations \mathbf{b}_i are the third derivatives, etc. The positions a small time step Δt earlier are derived from eq. (14.3) by substituting Δt with $-\Delta t$.

$$\mathbf{r}_{i-1} = \mathbf{r}_i - \mathbf{v}_i(\Delta t) + \frac{1}{2} \mathbf{a}_i(\Delta t)^2 - \frac{1}{6} \mathbf{b}_i(\Delta t)^3 + \dots \quad (14.4)$$

Addition of eqs (14.3) and (14.4) gives a recipe for predicting the position a time step Δt later from the current and previous positions, and the current acceleration. The latter can be calculated from the force, or equivalently, the potential.

$$\begin{aligned} \mathbf{r}_{i+1} &= (2\mathbf{r}_i - \mathbf{r}_{i-1}) + \mathbf{a}_i(\Delta t)^2 + \dots \\ \mathbf{a}_i &= \frac{\mathbf{F}_i}{m_i} = -\frac{1}{m_i} \frac{dV}{d\mathbf{r}_i} \end{aligned} \quad (14.5)$$

This is the *Verlet* algorithm⁸ for solving Newton's equation numerically. Note that the term involving the change in acceleration (\mathbf{b}) disappears, i.e. the equation is correct to third order in Δt . At the initial point, the previous positions are not available, but can be estimated from a first-order approximation of eq. (14.3).

$$\mathbf{r}_{-1} = \mathbf{r}_0 - \mathbf{v}_0\Delta t \quad (14.6)$$

At each time step, the acceleration must be evaluated from the forces, eq. (14.5), which then allows the atomic positions to be propagated in time and thus generate a trajectory. As the step size Δt is decreased, the trajectory becomes a better and better approximation to the "true" trajectory, until the practical problems of finite numerical

accuracy arise (e.g. the forces cannot be calculated with infinite precision). A small time step, however, means that more steps are necessary for propagating the system a given total time, i.e. the computational effort increases inversely with the size of the time step.

The Verlet algorithm has the numerical disadvantage that the new positions are obtained by adding a term proportional to Δt^2 to a difference in positions ($2\mathbf{r}_i - \mathbf{r}_{i-1}$). Since Δt is a small number and $(2\mathbf{r}_i - \mathbf{r}_{i-1})$ is a difference between two large numbers, this may lead to truncation errors due to finite precision. The Verlet algorithm furthermore has the disadvantage that velocities do not appear explicitly, which is a problem in connection with generating ensembles with constant temperature, as discussed in Section 14.2.2.

The numerical aspect and the lack of explicit velocities in the Verlet algorithm can be remedied by the *leap-frog* algorithm.⁹ Performing expansions analogous to eqs (14.3) and (14.4) with half a time step followed by subtraction gives eq. (14.7).

$$\mathbf{r}_{i+1} = \mathbf{r}_i + \mathbf{v}_{i+\frac{1}{2}}\Delta t \quad (14.7)$$

The velocity is obtained by analogous expansions to give eq. (14.8).

$$\mathbf{v}_{i+\frac{1}{2}} = \mathbf{v}_{i-\frac{1}{2}} + \mathbf{a}_i\Delta t \quad (14.8)$$

Eqs (14.7) and (14.8) define the leap-frog algorithm, and it is seen that the position and velocity updates are out of phase by half a time step. In terms of theoretical accuracy it is also of third order, as the Verlet algorithm, but the numerical accuracy is better. Furthermore, the velocities appear directly, which facilitates a coupling to an external heat bath (Section 14.2.2). The disadvantage is that the positions and velocities are not known at the same time, they are always out of phase by half a time step. The latter abnormality can be removed by the *velocity Verlet* algorithm, where the equations used to propagate the atoms are given in eq. (14.9).¹⁰

$$\begin{aligned} \mathbf{r}_{i+1} &= \mathbf{r}_i + \mathbf{v}_i\Delta t + \frac{1}{2}\mathbf{a}_i\Delta t^2 \\ \mathbf{v}_{i+1} &= \mathbf{v}_i + \frac{1}{2}\{\mathbf{a}_i + \mathbf{a}_{i+1}\}\Delta t \end{aligned} \quad (14.9)$$

The preference of Verlet and leap-frog type algorithms over for example Runge–Kutta methods (Section 12.8) in MD simulations is that they are time-reversible, which in general tend to improve the energy conservation over long simulation times.¹¹

The above solves the dynamical equation by a numerical integration of Newton's second equation. In some cases, it is useful to rewrite the equations in a more general form. Denoting a generalized coordinate with \mathbf{q} and its conjugate moment by \mathbf{p} ($\mathbf{p} = m\partial\mathbf{q}/\partial t$), eq. (14.2) becomes eq. (14.10).

$$-\frac{\partial V}{\partial \mathbf{q}} = \frac{\partial \mathbf{p}}{\partial t} \quad (14.10)$$

This can also be formulated in terms of a Lagrange function L .

$$\begin{aligned} L &= T - V \\ \frac{d}{dt} \frac{\partial L}{\partial \mathbf{p}} - \frac{\partial L}{\partial \mathbf{q}} &= 0 \end{aligned} \quad (14.11)$$

Using the fact that $T = \mathbf{p}^2/2m$, it can be seen that the Lagrange equation is completely equivalent to the Newton formulation.

Yet another formulation is given by the Hamilton function H and eq. (14.12), which again may be verified to be completely equivalent to eq. (14.2).

$$\begin{aligned}
 H &= T + V \\
 \frac{\partial H}{\partial \mathbf{q}} + \frac{d\mathbf{p}}{dt} &= 0 \\
 \frac{\partial H}{\partial \mathbf{p}} - \frac{d\mathbf{q}}{dt} &= 0
 \end{aligned}
 \tag{14.12}$$

The main advantage of the Lagrange and Hamilton formulations is that any set of non-redundant variables can be used, while the Newton formulation focuses on spatial coordinates and corresponding velocities. The main difference between the Lagrange and Hamilton formulations is that the former is a single second-order differential equation, while the latter is a coupled set of first-order differential equations. Depending on the system, one of them may be easier to solve than the other.

The time step employed is an important control parameter for a simulation. The maximum time step that can be taken is determined by the rate of the *fastest* process in the system, i.e. typically an order of magnitude smaller than the fastest process. Molecular motions (rotations and vibrations) typically occur with frequencies in the range 10^{11} – 10^{14} s⁻¹ (corresponding to wavenumbers of 3–3300 cm⁻¹), and time steps of the order of femtoseconds (10^{-15} s) or less are required to model such motions with sufficient accuracy. This means that a total simulation time of 1 nanosecond (10^{-9} s) requires $\sim 10^6$ time steps, and 1 microsecond (10^{-6} s) requires $\sim 10^9$ time steps. A million time steps is already a significant computational effort and typical simulation times are in the nano- or picosecond range. Unfortunately, many interesting phenomena occur on a substantially longer time scale: protein folding and chemical reactions, for example, occur on the order of milliseconds or seconds. Furthermore, a single trajectory may not be adequate for representing the dynamics, thus requiring that many runs must be carried out with different starting conditions (positions and velocities) and be properly averaged.

For molecules, the fastest processes are the stretching vibrations, especially those involving hydrogen. These degrees of freedom, however, have relatively little influence on many properties. It is therefore advantageous to freeze all bond lengths involving hydrogen atoms, which allow longer time steps to be taken, and consequently longer simulation times to be obtained for the same computational cost. As all atoms move individually according to Newton's equation, constraints must be applied for keeping bond lengths fixed. This is normally done by either the *SHAKE*¹² (Verlet) or *RATTLE*¹³ (velocity Verlet) algorithms, where the distance constraints are incorporated by the method of Lagrange undetermined multipliers (Section 12.5). The atoms are first allowed to move under the influence of the forces, and subsequently forced to obey the constraints by making a few sequential passes through all the variables.

Enforcement of bond length constraints typically allows the time step to be increased by a factor of 2 or 3. Angles may also be frozen by adding a distance constraint on atoms that are 1,3 relative to each other. Angle bending, however, affects calculated properties more than bond stretching and fixing them may often introduce

unacceptable errors. Angle constraints are therefore used less frequently. A simulation can also be performed using fixed molecular geometries, i.e. only the positions and relative orientations of individual molecules are allowed to change. In such cases, the natural variables to propagate in time are the centre of mass position and the three Euler angles of each molecule.

14.2.2 Generating non-natural ensembles

A standard MD simulation generates an *NVE* ensemble, i.e. the temperature and pressure will fluctuate. The total energy is a sum of the kinetic and potential energies, and can be calculated from the positions and velocities.

$$E_{\text{tot}} = \sum_{i=1}^N \frac{1}{2} m_i \mathbf{v}_i^2 + V(\mathbf{r}) \quad (14.13)$$

Owing to the finite precision with which the atomic forces are evaluated, and the finite time step used, the total energy is *not* exactly constant, but this error can be controlled by the magnitude of the time step. Indeed, preservation of the energy to within a given threshold may be used to define the maximum permissible time step.

The temperature of the system is proportional to the average kinetic energy.

$$\langle E_{\text{kin}} \rangle = \frac{1}{2} (3N_{\text{atoms}} - N_{\text{constraint}}) kT \quad (14.14)$$

The number of constraints is typically three, corresponding to conservation of linear momentum. Note that for 1 mole of particles, eq. (14.14) reduces to the familiar expression $\langle E_{\text{kin}} \rangle = 3/2 RT$. Since the kinetic energy is the difference between the total energy (almost constant) and the potential energy (depends on the positions), the kinetic energy will vary significantly, i.e. the temperature will be calculated as an average value with an associated fluctuation. Similarly, if the volume of the system is fixed, the pressure will fluctuate.

Although the *NVE* is the natural ensemble generated by an MD simulation, it is possible also to generate *NVT* or *NPT* ensembles by MD techniques by modifying the velocities or positions in each time step. As indicated in eq. (14.14) the instant value of the temperature is given by the average of the kinetic energy. If this is different from the desired temperature, all velocities may be scaled by a factor of $(T_{\text{desired}}/T_{\text{actual}})^{1/2}$ in each time step to achieve the desired temperature. Such an “instant” correction procedure actually alters the dynamics, such that the simulation no longer corresponds to a canonical (*NVT*) ensemble. Performing the scaling at larger intervals introduces some periodicity into the simulation, which is also undesirable. Alternatively the system may be coupled to a “heat bath”, which gradually adds or removes energy to/from the system with a suitable time constant, often called a *thermostat*.¹⁴ The kinetic energy of the system is again modified by scaling the velocities, but the rate of heat transfer is controlled by a coupling parameter τ .

$$\begin{aligned} \frac{dT}{dt} &= \frac{1}{\tau} (T_{\text{desired}} - T_{\text{actual}}) \\ &\Updownarrow \\ \text{velocity scale factor} &= \sqrt{1 + \frac{\Delta t}{\tau} \left(\frac{T_{\text{desired}}}{T_{\text{actual}}} - 1 \right)} \end{aligned} \quad (14.15)$$

Thermostat methods such as eq. (14.15) are widely used but again do not produce a canonical ensemble. They do produce correct averages but give incorrect fluctuations of properties. In *Nosé–Hoover* methods¹⁵ the heat bath is considered an integral part of the system and assigned fictive dynamic variables, which are evolved on an equal footing with the other variables. These methods are analogous to the extended Lagrange methods described in Section 14.2.5, and can be shown to produce true canonical ensembles.

The pressure can similarly be held (approximately) constant by coupling to a “pressure bath”. Instead of changing the velocities of the particles, the volume of the system is changed by scaling all coordinates according to eq. (14.16).

$$\begin{aligned} \frac{dP}{dt} &= \frac{1}{\tau} (P_{\text{desired}} - P_{\text{actual}}) \\ \Downarrow \\ \text{coordinate scale factor} &= \sqrt[3]{1 + \kappa \frac{\Delta t}{\tau} (P_{\text{actual}} - P_{\text{desired}})} \end{aligned} \quad (14.16)$$

Here the constant κ is the compressibility of the system. Such *barostat* methods are again widely used, both in MC and MD simulations, but do not produce strictly correct ensembles. Alternatively, the pressure may be maintained by a *Nosé–Hoover* approach in order to produce a correct ensemble.

14.2.3 Langevin methods

Molecular dynamics methods generate detailed information about all the particles in the system and are therefore well suited for calculating collective properties. In other cases, the major interest is in the dynamics of a single molecule, in which case the surrounding molecules can be modelled by only including the average interactions. This average interaction is assumed to have a friction term (with a friction coefficient ζ) proportional to the atomic velocity, and a random component ($\mathbf{F}_{\text{random}}$) that averages to zero. These terms are in addition to the normal intramolecular forces ($\mathbf{F}_{\text{intra}}$) and possibly also external forces, for example from an electric field. The random force is associated with a temperature and adds energy to the system, while the friction term removes energy. The random force is typically taken to have a Gaussian distribution with a mean value of zero.

$$m \frac{d^2 \mathbf{r}}{dt^2} = -\zeta \frac{d\mathbf{r}}{dt} + \mathbf{F}_{\text{intra}} + \mathbf{F}_{\text{random}} \quad (14.17)$$

Eq. (14.17) is called the *Langevin* equation of motion, and gives rise to *stochastic* or *Brownian dynamics*.¹⁶ The magnitude of the friction coefficient determines the importance of the intramolecular forces compared with the friction term, and large values of ζ lead to the Brownian dynamics limit.

14.2.4 Direct methods

The major computational effort in a molecular or Langevin dynamic simulation is the calculation of the forces on all particles at each time step. In principle, any type of energy function can be used: force field, semi-empirical, *ab initio* electronic structure

or DFT methods. Owing to the small time step required, and the resulting many force evaluations necessary, the large majority of simulations are performed with parameterized energy functions of the force field type. For studying macromolecules and solvation, general force fields of the type discussed in Chapter 2 are normally used. While these may be of sufficient accuracy for simulating structural properties, they are unable to describe chemical reactions or to achieve high accuracy. For such cases, a “global” energy surface may be constructed by fitting high-level *ab initio* results and experimental data to a suitable functional form.¹⁷ For sufficiently small time steps, the result of a simulation is determined entirely by the quality of the energy surface. To obtain “converged” results for the dynamics, the energy surface must be accurate to better than 1 kJ/mol, over the *whole* surface that is accessible at the given energy (temperature). Constructing such high-quality “global” energy surfaces is very demanding and has only been done for a few systems. As mentioned in Chapter 1, the sheer dimensionality prevents an adequate sampling of a surface by point calculations for more than three or four atoms, and high-level dynamics have thus been limited to systems of this size.

Even for low-dimensional surfaces (three to six atoms), it is often difficult to design a well-behaved fitting function capable of yielding a balanced description of all reaction channels. A simulation of the reaction between an oxygen atom and methane, for example, requires a balanced energy description of the following (stable) species, as well as the reaction paths connecting these:

- $\text{CH}_4 + {}^{1,3}\text{O}$
- CH_3OH
- $\text{H}_2\text{CO} + \text{H}_2$
- ${}^{1,3}\text{HCOH}$ (*cis* and *trans*) + H_2
- $\text{CH}_3\text{O} + \text{H}$
- ${}^{1,3}\text{CH}_2 + \text{H}_2\text{O}$
- $\text{CH}_3 + \text{OH}$

Note that both singlet and triplet energy surfaces are important for this reaction. Achieving a high-quality surface will require a large number of MR-CI type calculations, and designing a suitable interpolation function to reproduce the experimental energy differences between all the above exit channels is a non-trivial exercise.

The surface design and fitting process can be bypassed by performing the dynamics “directly”, i.e. by calculating the required energies and forces in each time step of a simulation. The advantage is that a fitting function is not required, there is no parameterization step, and only the part of the surface actually visited by the dynamics has to be calculated. The disadvantage is that the same (or almost the same) points may be calculated many times, and if many trajectories are required, the total amount of points calculated may be larger than required for performing a global fit. Furthermore, it is difficult to add empirical corrections to the calculated surface. In a global fit approach, deficiencies in the employed computational method can be partly alleviated by enforcing energy differences between experimentally known species in the parameterization step. In some cases, the employed electronic structure method is pre-modified in direct approaches in order to give better agreements with experimental quantities. The latter has especially been used in connection with semi-empirical methods such as AM1 and PM3, where the atomic parameters can be re-tuned to

model a specific reaction surface better than the defaults parameters, a procedure called *Specific Reaction Parameterization* (SRP).¹⁸

14.2.5 Extended Lagrange techniques (Car–Parrinello methods)

Traditionally, direct dynamics with electronic structure methods have been done using a converged wave function at each time step. In order to fulfil energy conservation over the whole simulation length, however, such *Born–Oppenheimer dynamics* require a very tight convergence of the wave function in each time step, otherwise the electrons will create an artificial frictional term on the nuclei, and this makes the procedure computationally expensive.¹⁹ In an elegant breakthrough by *Car and Parrinello* (CP),²⁰ it was shown that it is not necessary to fully converge the wave function in each time step. After having determined a converged wave function at the first point, the essence of the CP technique is to let the wave function parameters (orbitals) evolve *simultaneously* with the changes in nuclear positions. This can be achieved by including the wave function parameters as variables with fictive “masses” in the dynamics, analogous to the nuclear positions and masses. Since this involves generalized variables, the Lagrange formulation (eq. (14.11)) for the dynamical equation is convenient. The use of such *extended Lagrange functions* for describing the evolution of a system with both “real” (nuclear/electronic) and “fictive” (method parameters) is quite general, and is for example also used in force field methods incorporating fluctuating charges and/or polarization.

For the case of the CP method, the nuclear contributions are given by eq. (14.18).

$$\begin{aligned}
 L &= T_{\text{nuc}} - V_{\text{nuc}} \\
 T_{\text{nuc}} &= \frac{1}{2} \sum_a^{N_{\text{nuc}}} M_a \left(\frac{d\mathbf{R}_a}{dt} \right)^2 \\
 V_{\text{nuc}} &= V(\mathbf{R}_{\text{nuc}})
 \end{aligned} \tag{14.18}$$

We now add contributions corresponding to treating the orbital expansion coefficients as variables with fictive masses μ_i .

$$\begin{aligned}
 \phi_i &= \sum_{\alpha}^{M_{\text{basis}}} c_{i\alpha} \chi_{\alpha} \\
 T_{\text{orb}} &= \frac{1}{2} \sum_i^{N_{\text{orb}}} \mu_i \left(\frac{dc_{i\alpha}}{dt} \right)^2 \\
 V_{\text{orb}} &= E(\mathbf{c}_{\text{orb}}) \\
 L &= T_{\text{nuc}} - V_{\text{nuc}} + T_{\text{orb}} - V_{\text{orb}}
 \end{aligned} \tag{14.19}$$

The two potential energies can be combined to a single term depending on both the nuclear positions and the orbital coefficients.

$$\begin{aligned}
 L &= T_{\text{nuc}} + T_{\text{orb}} - V_{\text{tot}} \\
 V_{\text{tot}} &= E(\mathbf{R}_{\text{nuc}}, \mathbf{c}_{\text{orb}})
 \end{aligned} \tag{14.20}$$

The orbital orthogonality constraints can be included by addition of terms involving Lagrange multipliers.

$$\begin{aligned}
\langle \phi_i | \phi_j \rangle &= \delta_{ij} \\
\sigma_{ij} &= \sum_{\alpha\beta}^{M_{\text{basis}}} c_{i\alpha} c_{j\beta} \langle \chi_\alpha | \chi_\beta \rangle - \delta_{ij} = 0 \\
L &= T_{\text{nuc}} + T_{\text{orb}} - V_{\text{tot}} - \sum_{ij}^{N_{\text{orb}}} \lambda_{ij} \sigma_{ij}
\end{aligned} \tag{14.21}$$

The resulting dynamical equations then become eq. (14.22)

$$\begin{aligned}
M_a \frac{\partial^2 \mathbf{R}_a}{\partial t^2} &= -\frac{\partial E}{\partial \mathbf{R}_a} + \sum_{ij}^{N_{\text{orb}}} \lambda_{ij} \frac{\partial \sigma_{ij}}{\partial \mathbf{R}_a} \\
\mu_i \frac{\partial^2 c_{i\alpha}}{\partial t^2} &= -\frac{\partial E}{\partial c_{i\alpha}} + \sum_{ij}^{N_{\text{orb}}} \lambda_{ij} \frac{\partial \sigma_{ij}}{\partial c_{i\alpha}}
\end{aligned} \tag{14.22}$$

The constraint forces are handled iteratively, analogously to the constraint of fixed bond lengths in the SHAKE algorithm.

If the nuclear positions \mathbf{R} are kept constant and the fictive orbital kinetic energy T_{orb} is quenched, the resulting algorithm is essentially a steepest descent minimization of the electronic energy with respect to the orbital coefficients. This is done at the initial point, but at subsequent points the orbital parameters are allowed to evolve along with the nuclear position according to the dynamical equation. This means that the nuclear forces are not strictly correct since the electronic wave function is not converged in the orbital parameter space. This error, however, can be controlled by suitable choices of the fictive masses associated with the orbital parameters, i.e. small values provide results close to the “true” Born–Oppenheimer results, but also require the use of small time steps since the resulting “orbital parameter frequency” is high.²¹ Typically, the fictive masses are taken to be a few hundred atomic units, giving time steps of ~ 0.1 femtoseconds, i.e. roughly an order of magnitude smaller than for classical molecular dynamics. It should be noted that the optimum value for the fictive masses depends on the system, and for metals and semiconductors, for example, it is difficult to choose suitable values. Systems containing hydrogen are especially problematic, since the proton mass (1836 au) is only a factor of 5–10 higher than the fictive orbital parameter mass. This in some cases leads to a coupling of these degrees of freedom, but this can be partly countered by using deuterium instead of hydrogen.

In the CP approach, the total energy is conserved, and this now includes the fictive kinetic energy of the orbital parameters. The “real” system of course has no orbital kinetic energy, and this must therefore be kept small compared with the other terms in order for the CP method to provide realistic simulation results. The magnitude of the fictive masses for the orbital parameters serves as a coupling parameter between the nuclear and parameter kinetic energies. In an equilibrium condition, the temperature associated with the nuclei and orbital parameters are the same, but it is usually desirable to have the parameter kinetic energy to be significantly lower (a value of zero corresponds to the Born–Oppenheimer case). This can be obtained by continuously removing the fictive kinetic energy associated with the orbital parameters, while compensating for the energy loss by adding energy to the nuclei. In practice this can be obtained by allowing the orbital parameters to interact with a “heat bath” of a low temperature, while the nuclei interact with a heat bath of the desired simulation temperature.

The CP technique can be used both in a “static” sense, for simultaneously optimizing the wave function and the nuclear positions by periodically quenching the kinetic energies, but it can also be used in a “dynamical” sense for sampling the (nuclear) phase space. The main advantage of the CP technique is the much better error cancellation compared with a Born–Oppenheimer dynamics, i.e. even with non-converged wave function parameters, the long-term energy conservation is fulfilled to a quite high accuracy. The coupling of the real and fictive parameters builds a self-correction into the CP method, i.e. if the nuclei at some point get slightly “ahead” of the electron cloud, they will be slowed down, thus allowing the electrons to “catch up” with the nuclei. Similarly, if the electronic parameters get ahead of the nuclei, the nuclei will be accelerated owing to the Coulomb attraction.

Such *ab initio* simulations can in principle be carried out with any type of wave function but they are still significantly more expensive computationally than traditional parameterized energy functions. The CP method was originally implemented with DFT methods using plane waves as the basis set but more recently the technique has also been used with other types of methods (e.g. HF or MP2) and Gaussian type basis functions, where the density matrix elements are used as variables instead of the molecular orbital coefficients.²² The great advantage over force field type functions is that electronic structure methods are able to describe bond breaking/formation, i.e. CP methods allow a direct simulation of chemical reactions and processes such as hydrogen exchange in water. Even with the CP technique, however, the use of *ab initio* electronic structure calculations (HF, DFT) is so expensive that only picosecond simulation can be carried out, compared with nano- or microsecond simulations with parameterized energy functions.

The CP method may be considered as a semi-classical dynamics approach where the electrons are treated quantum mechanically while the nuclear motion is treated classically. The latter implies that for example zero point vibrational effects are not included, nor can nuclear tunnelling effects be described; this requires fully quantum methods, as described in the next section.

14.2.6 Quantum methods using potential energy surfaces

In order to incorporate quantum effects into the nuclear motions (vibrational effects and tunnelling), the time-dependent (nuclear) Schrödinger equation must be used in place of Newton’s equation.

$$\mathbf{H}\Psi = (\mathbf{T} + \mathbf{V})\Psi = i\frac{\partial\Psi}{\partial t} \quad (14.23)$$

Here \mathbf{T} is the kinetic energy operator and \mathbf{V} is the potential energy. The square of the wave function is the *probability* of finding a particle at a given position. Heisenberg’s uncertainty principle means that a quantum description of a nucleus must be a continuous function, not a single specific position as in classical mechanics.²³ Such a continuous function is often denoted a *wave package* and may be modelled by Gaussian functions (semi-classical methods) or numerically (quantum methods). Analogously to classical dynamics, the wave function may be propagated through a series of small, but finite, time steps.

$$\Psi_{i+1} = -i\mathbf{H}\Psi_i\Delta t = -i(\mathbf{T} + \mathbf{V})\Psi_i\Delta t \quad (14.24)$$

Each time step thus involves a calculation of the effect of the Hamiltonian operator acting on the wave function. In fully quantum methods, the wave function is often represented on a *grid* of points, these being the equivalent of basis functions for an electronic wave function. The effect of the potential energy operator is easy to evaluate, as it just involves a multiplication of the potential at each point with the value of the wave function. The kinetic energy operator, however, involves the derivative of the wave function, and a direct evaluation would require a very dense set of grid points for an accurate representation.

The kinetic energy operator is proportional to the square of the momentum, $\mathbf{T} = \mathbf{p}^2/2m$. In a *momentum* representation (i.e. using the particle momentum instead of position as variables), \mathbf{T} is a simple multiplication operator, analogous to \mathbf{V} in position space. The transformation from position to momentum space can be achieved by a Fourier transformation. A numerical solution of the time-dependent Schrödinger equation can thus be done by switching back and forth between a position and momentum representation of the wave function, evaluating the effect of \mathbf{V} in position space, and the effect of \mathbf{T} in momentum space. Analogously to the leap-frog algorithm for the classical case (eqs (14.7) and (14.8)), the update of the wave function by the potential and kinetic energy operators may be chosen to be out of phase by half a time step to improve the accuracy. The key to the popularity of this approach is the presence of highly efficient computer routines for performing Fourier transformations.

The requirement of an accurate global energy surface is even more important for a quantum mechanical treatment than for the classical case, since the wave function depends on a finite part of the surface, not just a single point. The updating of the positions and velocities are computationally inexpensive in the classical case, once the forces are available, but the requirement of two Fourier transformations in each time step makes the quantum propagation a significant computational issue. Furthermore, the representation of the wave function on a grid effectively limits the dimensionality to a maximum of three, i.e. di- and triatomic systems (one and three internal coordinates, respectively). Larger systems necessitate freezing some of the coordinates, or treating them classically.²⁴

14.2.7 Reaction path methods

The main problem in dynamical studies is the requirement of a continuous energy surface over a wide range of geometries. A simulation will normally be done with specification of an energy (or a temperature), and a surface must thus be available for all nuclear configurations that have an energy lower than the chosen simulation value. For quantum methods, the surface must also be available at higher energies as the wave function has a tail that penetrates into classically “forbidden” areas.

Traditionally, such “global” energy surfaces have been constructed by fitting a suitable functional form to energies (and possibly also first and second derivatives) calculated by *ab initio* methods at a large number (perhaps a few hundreds or thousands) of geometries.²⁵ The function may be further refined by including experimental data (such as vibrational frequencies and geometries) in the fitting. For “large” systems (i.e. more than three or four atoms) the generation of an adequate number of fitting points

is prohibitively expensive. In order to treat large systems, it is necessary to concentrate the computational effort on the “chemically important” part of the potential energy surface.

In the simplest description, a chemical reaction takes place along the lowest energy path connecting the reactant and product, passing over the transition structure (Section 13.1) as the highest point. This is the *Minimum Energy Path* (MEP), which in mass-weighted coordinates is called the *Intrinsic Reaction Coordinate* (IRC) (Section 12.8). The idea in *Reaction Path* (RP) methods²⁶ is to only consider the energy surface in the immediate vicinity of a suitable one-dimensional reaction path, which usually (but not necessarily) is taken as the IRC. The potential is typically expanded to second order along the reaction path, corresponding to modelling the perpendicular degrees of freedom as harmonic vibrational frequencies. The reaction path potential may be generated by a series of frequency calculations at points along the IRC, and the point-wise potential made continuous by interpolation. The potential may be generated prior to the reaction path calculation, or generated “on the fly” in a “direct” fashion.²⁷ Moyano and Collins have proposed a hybrid method where all the points calculated are stored and used for interpolation if the required point is sufficiently close to prior points.²⁸ This approach thus starts out as a direct type dynamics but ends up with an implicitly parameterized surface for sufficiently long simulation times. For long simulation times or for running many trajectories, the savings by interpolation can be substantial.

The reaction path method may be generalized by having two “reaction coordinates” (a *reaction surface*) treated explicitly and the remaining degrees of freedom treated approximately, or by having three “reaction coordinates” (a *reaction volume*).²⁹ These generalizations are useful for performing mixed classical–quantum dynamics, where the dynamics with the reaction coordinate(s) are treated quantum mechanically while the remaining degrees of freedom are treated classically.

The inclusion of dynamical effects allows the calculation of corrections to simple transition state theory, often described by a transmission coefficient κ to be multiplied with the TST rate constant (Section 13.1), or used in connection with variational TST (Section 13.5). Classical dynamics allow corrections due to re-crossing to be calculated, while a quantum treatment is necessary for including tunnelling effects. Owing to the stringent requirement of a highly accurate global energy surface, there are only a few systems that have been subjected to a rigorous analysis.

The tunnelling effect is sometimes approximated by inclusion of a semi-classical correction based on tunnelling through the barrier along the minimum energy path (i.e. the IRC). The *Bell* correction is based on the assumption that the (one-dimensional) energy curve near the transition state can be approximated by a parabola.³⁰ This yields a correction factor that only depends on the activation energy ΔE^\ddagger and the magnitude of the imaginary frequency ν_i , i.e. the curvature of the potential energy surface at the TS.

$$\kappa_{\text{Bell}} = \frac{\frac{1}{2}u^\ddagger}{\sin \frac{1}{2}u^\ddagger} - u^\ddagger e^{-\Delta E^\ddagger/kT} \left(\frac{e^{-\beta}}{2\pi - u^\ddagger} - \frac{e^{-2\beta}}{4\pi - u^\ddagger} - \frac{e^{-3\beta}}{6\pi - u^\ddagger} - \dots \right) \quad (14.25)$$

$$u^\ddagger = \frac{h\nu_i}{kT} \quad \beta = \frac{2\pi\Delta E^\ddagger}{h\nu_i}$$

Except for reactions with low barriers (i.e. $<40\text{ kJ/mol}$ at $T = 300\text{ K}$) or at high temperatures, the quantity $\Delta E^\ddagger/kT$ is large and the last series can be neglected. The tunnelling correction is then given completely in terms of the magnitude of the imaginary frequency. For small values of u^\ddagger the first term may be Taylor expanded to give eq. (14.26).

$$\kappa_{\text{Bell}} = 1 + \frac{1}{24}(u^\ddagger)^2 + \dots \quad (14.26)$$

The first-order term is known as the *Wigner* correction.³¹

It is possible to derive tunnelling corrections for functional forms of the energy barrier other than an inverted parabola, but these cannot be expressed in analytical form. Since any barrier can be approximated by a parabola near the TS, and since tunnelling is most important for energies just below the top, they tend to give results in qualitative agreement with the Bell formula.

The main approximation of such one-dimensional corrections is that the tunnelling is assumed to occur *along* the MEP. This may be a reasonable assumption for reactions having either early or late (close to either reactant or product) transition states. For reactions where bond breaking and formation are both significant at the TS (as is usually the case), the dominant tunnelling effect is “corner cutting” (Figure 14.1), i.e. the favoured tunnel path is not along the MEP. Although the energy increases away from the MEP, the barrier also becomes narrower on the concave side of the reaction path, which favours the tunnelling probability.




Figure 14.1 A contour plot illustration of the “corner cutting” tunnelling path

Truhlar and coworkers have developed various approximate schemes for including tunnelling in multi-dimensional systems.³² In the *Minimum Energy Path Semi-classical Adiabatic Ground state* (MEPSAC) approximation the tunnelling is assumed to occur along the MEP, analogous to the Bell approach, but for an arbitrary shape of the energy surface. The *Small Curvature Semi-classical Adiabatic Ground state* (SCSAC) approximation allows tunnelling to occur within one vibrational half-amplitude perpendicular to the reaction path, while the *Large Curvature Ground state* (LCG) approximation

allows tunnelling to occur outside this region. The SCSAC method requires a knowledge of the (generalized) frequencies along the IRC (Section 12.8), which can be obtained by calculating the force constant matrix at suitable intervals and interpolating the results. The LCG methods require additional calculations away from the IRC.

14.2.8 *Non-Born–Oppenheimer methods*

The methods in Sections 14.2.5 and 14.2.6 attempt to include nuclear quantum corrections based on the Born–Oppenheimer separation of the nuclear and electronic degrees of freedom, i.e. solving the nuclear dynamics on a potential energy surface obtained by solving the electronic Schrödinger equation. When quantum corrections such as tunnelling are large, however, it is an implicit warning that the Born–Oppenheimer approximation may also be problematic. Rather than trying to improve on the underlying model by adding correction terms, it may be both easier and better to treat the nuclei within a quantum framework from the start.

Methods that treat all of the electron and nuclear degrees of freedom within a combined quantum framework are starting to appear; so far they are mostly based on a mean-field (i.e. Hartree–Fock) approximation where the coupling of the nuclear and electron motions is included in an average fashion. Both conceptual and computational developments are required before such methods can be considered mature.³³ One clear advantage of these methods is the ability to implicitly include both tunnelling and vibrational effects, and to selectively treat some nuclei as classical, thereby allowing a simplification for large systems.

In the spirit of the Car–Parrinello approach, the whole set of variables (nuclei and wave function parameters) may also be allowed to evolve simultaneously by solving the time-dependent Schrödinger equation. Örn and coworkers have developed an *Electron–Nuclear Dynamics* (END) method,³⁴ where both the orbitals describing the electronic wave function and the nuclear degrees of freedom are described by expansion into a Gaussian basis set, which moves along with the nuclei. Such an approach in principle allows a complete solution of the combined nuclear–electron Schrödinger equation without having to invoke approximations beyond those imposed by the basis set. Inclusion of the electronic parameters in the dynamics, however, means that the fundamental time step is short, and this results in a high computational cost for even quite short simulations and simple wave functions.

14.2.9 *Constrained sampling methods*

The reaction path methods described in Section 14.2.7 focuses on the lowest energy reaction path on the electronic energy surface. The activation energy related to the experimental reaction rate, however, depends on the free energy, i.e. one would optimally like to locate the reaction path on the free energy surface. For small systems, this can be done by adding finite temperature corrections to the enthalpy and entropy in a rigid-rotor harmonic-oscillator approximation based on a second-order expansion of the energy around each point (eqs (13.36) and (13.37)). For large systems, however, the harmonic approximation is less suitable and a more complete sampling by dynamical methods is usually desired. This will also yield information about the dynamics in the perpendicular direction around the reaction path.

A straightforward sampling of the reaction path is not possible since the dynamics at ordinary temperatures only very rarely visit the high-energy region near the TS (unless the activation energy is close to zero). In order to achieve a sampling of a specific region of the energy surface with molecular dynamics or Monte Carlo methods, the sampling must be biased towards a specific volume of phase space. Analogously to the optimization of functions with constraints (Section 12.5), this can be done by two different methods, a penalty and a Lagrange type approach.

The penalty approach corresponds to augmenting the energy surface with a biasing potential U , for example a harmonic function centred at position \mathbf{r}_0 with a suitable width k_U .

$$\begin{aligned} V'(\mathbf{r}) &= V(\mathbf{r}) + U(\mathbf{r}) \\ U(\mathbf{r}) &= k_U(\mathbf{r} - \mathbf{r}_0)^2 \end{aligned} \quad (14.27)$$

By making the biasing potential sufficiently steep (large k_U), the energy of the augmented energy surface far from \mathbf{r}_0 will become so high in energy that only the region near \mathbf{r}_0 will be sampled at ambient temperatures, and this technique is called *umbrella sampling*.³⁵ The ensemble calculated with the augmented potential V' will of course be non-Boltzmann, but this can be deconvoluted as shown in eq. (14.28).

$$\langle A \rangle = \frac{\langle A(\mathbf{r}) e^{U(\mathbf{r})/kT} \rangle_{V'}}{\langle e^{U(\mathbf{r})/kT} \rangle_{V'}} \quad (14.28)$$

Here $\langle \rangle_{V'}$ indicates an average over the ensemble generated by the augmented potential. By performing a series of simulations with biasing potentials located at different positions along the reaction path, the free energy along the reaction path, often called the *Potential of Mean Force* (PMF), can be simulated.

The Lagrange approach constrains the sampling to the $(N-1)$ -dimensional subspace corresponding to a specific value of the reaction coordinate, where the constraint is fulfilled by means of an additional term in the Hamiltonian involving a Lagrange multiplier. This is related to the extended Lagrange techniques discussed in Section 14.2.5 and is usually referred to as *Blue Moon* sampling in the literature, and the Lagrange multiplier are called holonomic constraints.³⁶

The main disadvantage of the umbrella or Blue Moon sampling techniques is that the location of the biasing potential must be selected manually and an *a priori* knowledge of an approximate reaction coordinate is therefore required. Once this has been selected, the free energy along this path can be calculated. Since the sampling explores a (small) region around the selected path, the calculated PMF may deviate slightly from the initial selection. If desired, this updated PMF can then be used for a new series of simulations with biasing potentials located along the previously calculated PMF.³⁷ Such *adaptive* umbrella sampling methods should in principle converge on the true PMF but, in practice, the convergence is sensitive to the selection of a suitable initial reaction path.

14.3 Periodic Boundary Conditions

A realistic model of a solution requires at least several hundred solvent molecules. To prevent the outer solvent molecules from boiling off into space and minimizing surface effects, *periodic boundary conditions* are normally employed. The solvent molecules

are placed in a suitable box, often (but not necessarily) having a cubic geometry (it has been shown that simulation results using any of the five types of space-filling polyhedra are equivalent³⁸). This box is then duplicated in all directions, i.e. the central box is surrounded by 26 identical cubes, which are again surrounded by 98 boxes, etc.




Figure 14.2 Periodic boundary condition

If a solvent molecule leaves the central box through the right wall, its image will enter the box through the left wall from the neighbouring box. This means that the resulting solvent model becomes quasi-periodic, with a periodicity equal to the dimensions of the box.

As mentioned in Section 2.2.5, the electrostatic interaction is long-ranged and will extend beyond the boundary of a box. Truncating the interaction by using a cutoff distance of say 10 \AA gives discontinuous energies and forces, and has some rather unfortunate consequences in giving non-physical distributions of the solvent molecules near the cutoff distance and producing “hot” and “cold” spots. A *switching* function approach, where the interaction is gradually reduced to zero over a range of a few angstroms performs significantly better.³⁹ The switching function is multiplied onto the real potential and has the effect of smoothly reducing the potential from its real value to zero over a distance range from R_1 to R_2 . An example of a third-order switching function that has zero first derivatives at both limits⁴⁰ is shown in eq. (14.29).

$$S(r) = \begin{cases} 1 & r \leq R_1 \\ \frac{(R_2^2 - r^2)^2 (R_2^2 + 2r^2 - 3R_1^2)}{(R_2^2 - R_1^2)^3} & R_1 \leq r \leq R_2 \\ 0 & r \geq R_2 \end{cases} \quad (14.29)$$

An alternative form for the central part that also has vanishing second derivatives at both limits is given in eq. (14.30).

$$S(r) = 1 - 10 \left(\frac{r - R_1}{R_2 - R_1} \right)^3 + 15 \left(\frac{r - R_1}{R_2 - R_1} \right)^4 - 6 \left(\frac{r - R_1}{R_2 - R_1} \right)^5 \quad (14.30)$$

A variation of this is to use a *shifting* function, which corresponds to a switching function with $R_1 = 0$. Such functions modify the potential for all r values less than R_2 and an example is given in eq. (14.31).

$$S(r) = \begin{cases} \left(1 - \frac{r^2}{R_2^2}\right)^2 & r \leq R_2 \\ 0 & r \geq R_2 \end{cases} \quad (14.31)$$

The use of both switching and shifting functions modifies the model, since the potential and forces are changed, and therefore affects the results of the simulation. Whether these changes are significant relative to the other approximations in the model depends on the specific system and properties.

Figure 14.3 shows the energy function of two unit charges interacting with a Coulomb potential, one that has been subjected to the switching function eq. (14.29) with $R_1 = 10 \text{ \AA}$ and $R_2 = 12 \text{ \AA}$, and one that has been subjected to the shifting function eq. (14.31) with the same limits.




Figure 14.3 Difference between original, switched and shifted Coulomb potentials

Methods have also been developed where the electrostatic interaction is treated “exactly” (to within a numerical threshold), but without having to perform the N^2 summation over all atoms. *Ewald sum* methods have been developed for periodic systems (such as crystals) but can also be applied to quasi-periodic models arising by applying periodic boundary conditions. The idea in these methods is to split the interaction into a “near”- and “far”-field contribution.⁴¹ The near-field contribution is obtained by embedding each point charge in a screening potential, taken as a Gaussian function with an exactly opposing charge centred at the position of the point charge. Outside the range of the screening function, essentially given by the width of the Gaussian, the net charge is thus zero, and the interaction between these screened point charges is therefore short-ranged and can be evaluated directly. In order to recover the original point charge interaction, the effect of the screening potentials must be subtracted again. This compensating term is an interaction between Gaussian charge distributions,

which is long-ranged. Since it is a smooth charge distribution, however, it can be evaluated efficiently in reciprocal space by Fourier transform methods. The only free parameter is the width of the Gaussian potential. A narrow Gaussian function makes the direct-space part converge rapidly, but the reciprocal-space part converge slowly, and vice versa for a wide Gaussian function. The optimum width is given by the condition that the computational effort is distributed equally between the direct and reciprocal sums.




Figure 14.4 Illustration of the Ewald method

A key point in these methods is the existence of computationally efficient methods for performing Fourier transformations, which reduces the scaling from N^2 to $N^{3/2}$. A related method is the *Particle Mesh Ewald* (PME) method, which scales (only) as $N \ln(N)$.⁴²

The *Fast Multipole Moment* (FMM) method similarly splits the contribution into a near- and far-field, and calculates the near-field exactly.⁴³ The far-field energy is calculated by dividing the physical space into boxes, and the interaction between all molecules in one box with all molecules in another is approximated as interactions between multipoles located at the centres of the boxes. The further away from each other two boxes are, the larger the boxes can be for a given accuracy, thereby reducing the formal N^2 scaling into something that approaches linear scaling. The prefactor, however, is rather large and when properly implemented it appears that the cross-over point, where FMM becomes faster than PME, is around 10^5 particles. FMM furthermore works best when the particles are relatively uniformly distributed; for a non-uniform distribution of particles, the multipole order must be significantly increased in order to achieve a given accuracy. A disadvantage of FMM is that the maximum error (relative to an exact calculation) is significantly larger than for Ewald type methods, i.e. there are certain particle pairs for which the error is larger than the average error by perhaps a factor of 10. FMM, in contrast to Ewald-based methods, however, does not have the requirement of periodicity, i.e. it is capable of modelling large non-periodic systems.

The original FMM has been refined by also adjusting the accuracy of the multipole expansion as a function of the distance between boxes, producing the *very Fast Multipole Moment* (vFMM) method.⁴⁴

The exact calculation of the electrostatic interaction, albeit by treating the system as being pseudo-crystalline, has been shown to give significantly different results than a simple truncation scheme⁴⁵ and also different from a switching function approach.⁴⁶ Given the existence of computationally efficient methods for performing for example PME, there seems to be little reason for employing a non-physical truncation of the electrostatic interaction.




Figure 14.5 Illustration of the fast multipole moment method

14.4 Extracting Information from Simulations

A necessary (but not sufficient) requirement for producing a representative sampling is that the system is in equilibrium. The starting configuration may be generated by completely random positions (and velocities for MD), but is more often taken either from a previous simulation or by placing the particles at or near the lattice points of a suitable crystal. The system is then equilibrated by running perhaps 10^4 – 10^5 MC or MD steps, followed by perhaps 10^5 – 10^7 production steps. Various quantities, such as the average potential energy or correlation functions, can be monitored to validate whether equilibrium has been achieved.

The averaging in eq. (14.1) should be over configurations that are uncorrelated, and this is not the case for nearby points in an MD trajectory or sequence of MC steps. The whole set of points should therefore be divided into blocks with a length that is sufficiently long to make equivalent points in two neighbouring blocks uncorrelated, but preferably also with a length that is sufficiently short so that no information is lost. Flyvbjerg and Pedersen have shown how to determine the optimum block length by a sequence of statistical analyses.⁴⁷ For the original data set the mean and variance are calculated according to eq. (14.32).

$$\begin{aligned}\bar{x} &= \frac{1}{N} \sum_{i=1}^N x_i \\ \sigma^2 &= \frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2\end{aligned}\tag{14.32}$$

The variance calculated from eq. (14.32) is only valid for uncorrelated data, which is not the case for the original data. In order to get a realistic estimate of the true variance, we must perform a data compression to filter out the dependence, i.e. find the block size for producing uncorrelated data and calculate the variance using this blocking. The method of Flyvbjerg and Pedersen consists of performing a sequence of data compressions by averaging two neighbouring points, thereby reducing the data size by a factor of 2, and calculating the corresponding variance (the mean is unchanged). The variance divided by the number of data points at a given level, $\sigma^2/(N' - 1)$, will

initially increase and then level off to a constant value as the data within two consecutive blocks become uncorrelated. The point where the value becomes constant is the optimum block size for the given property and the $\sigma^2/(N' - 1)$ quantity can be taken as the estimate of the true variance of the property. The distance between (uncorrelated) data in MD methods has the dimension of time and is called the correlation time. It is important to recognize that different properties may have different correlation times, and for some properties it may be comparable to or exceed the total length of the simulation. A clear advantage of the above procedure for determining the optimum block size is that the statistical error bars associated with the variance can also be calculated, i.e. the standard deviation of the variance (eq. (14.33)).

$$\sigma^2(x) \approx \frac{\sigma^2(x')}{N' - 1} \left(1 \pm \sqrt{\frac{2}{N' - 1}} \right) \quad (14.33)$$

Here the prime notation indicates the data set at a given compression level. Eq. (14.33) clearly illustrates that the estimate of the variance becomes increasingly uncertain as the number of data blocks decreases, i.e. when the data has been compressed into only two blocks, the (relative) standard deviation is $\sqrt{2}$. In order to determine whether it is actually possible to obtain uncorrelated data from the simulation, a plot of $\sigma^2/(N' - 1)$ against compression level should therefore include the associated statistical error from eq. (14.33). If the statistical errors impinge on a conclusion as to whether a constant plateau has been reached, this is an indication that the simulation length is insufficient for obtaining valid estimates of the given quantity.

Ensembles generated by MC techniques are naturally of the constant NVT type, while MD methods naturally generate a constant NVE ensemble. Both MC and MD methods, however, may be modified to simulate other ensembles, as described in Sections 14.1.1 and 14.2.2. Of special importance is the constant NPT condition, which directly relates to most experimental conditions. The primary advantage of MD methods is that time appears explicitly, i.e. such methods are natural for simulating time-dependent properties, such as correlation functions, and for calculating properties that depend on particle velocities. Furthermore, if the relaxation time for a given process is (approximately) known, the required simulation time can be estimated beforehand (i.e. it must be at least several multiples of the relaxation time).

In order to reduce the statistical error, the averaging in eq. (14.1) is typically performed on 10^3 – 10^5 points in phase space. The requirement of calculating this many points and associated energies for a model consisting of several hundred particles means that the use of *ab initio* methods is extremely demanding, even for small systems and simple wave functions. Semi-empirical electronic structure methods may be used for small systems, implicitly accepting the low accuracy of these methods, but the large number of calculations necessary still makes this computationally intensive. The large majority of simulations are therefore carried out with an energy surface generated by a parameterized function of the force field type.

The expressions derived from statistical mechanics (Section 13.4) are often rewritten into computationally more suitable forms that may be evaluated from the basic descriptors: positions \mathbf{r} , velocities \mathbf{v} or momenta \mathbf{p} and energies E . The temperature is related to the average kinetic energy.

$$\frac{1}{2}(3N_{\text{atom}} - N_{\text{constraint}})kT = \left\langle \sum_i^{N_{\text{atom}}} \frac{1}{2} m_i \mathbf{v}_i^2 \right\rangle_M \quad (14.34)$$

In a standard MC simulation the temperature is fixed (*NVT* conditions), while it is a derived quantity in a standard MD simulation (*NVE* conditions).

The pressure is related to the product of positions and forces (for pairwise potentials).

$$PV = N_{\text{atom}}kT + \frac{1}{3} \left\langle \sum_{i<j}^{N_{\text{atom}}} r_{ij} f_{ij} \right\rangle_M \quad (14.35)$$

Here the first part is for an ideal gas.

The internal (potential) energy is directly a sum of energies, which is normally given as a sum over pairwise interactions (i.e. van der Waals and electrostatic contribution in a force field description).

$$U = \sum_i^{N_{\text{atom}}} E_i = \sum_{i<j}^{N_{\text{atom}}} E_{ij} \quad (14.36)$$

The internal energy will fluctuate around a mean value that may be calculated by averaging over the number of configurations, $\langle U \rangle_M$.

The heat capacity at constant volume is the derivative of the energy with respect to temperature at constant volume (eq. (13.17)). There are several ways of calculating such response properties. The most accurate is to perform a series of simulations under *NVT* conditions and thereby determine the behaviour of $\langle U \rangle_M$ as a function of T (for example by fitting to a suitable function). Subsequently this function may be differentiated to give the heat capacity. This approach has the disadvantage that several simulations at different temperatures are required. Alternatively, the heat capacity can be calculated from the *fluctuation* of the energy around its mean value.

$$C_V = \frac{1}{kT^2} \left\langle (U - \langle U \rangle_M)^2 \right\rangle_M = \frac{1}{kT^2} (\langle U^2 \rangle_M - \langle U \rangle_M^2) \quad (14.37)$$

This approach requires only a single simulation. Since the fluctuation has a longer relaxation time than the energy itself, the ensemble average in eq. (14.37) must be over a larger number of points than for $\langle U \rangle_M$ to achieve a similar statistical error, i.e. the efficiency obtained by avoiding multiple simulations is partly lost owing to a longer simulation time required. Another disadvantage is that eq. (14.37) involves taking differences between large numbers, which is susceptible to round-off errors.

Distribution functions measure the (average) value of a property as a function of an independent variable. A typical example is the radial distribution function $g(r)$ that measures the probability of finding a particle as a function of distance from a “typical” particle relative to that expected from a completely uniform distribution (i.e. an ideal gas with density N/V). The radial distribution function is defined in eq. (14.38).

$$g(r, \Delta r) = \frac{V}{N^2} \frac{\langle N(r, \Delta r) \rangle_M}{4\pi r^2 \Delta r} \quad (14.38)$$




Figure 14.6 A typical radial distribution function

Here $N(r, \Delta r)$ is the number of molecules between r and $r + \Delta r$ from another particle, and $4\pi r^2 \Delta r$ is the volume of a spherical shell with thickness Δr .

For a solution, the radial distribution function will typically have a structure as shown in Figure 14.6 for a simulation of a benzene radical anion in water.⁴⁸

Figure 14.6 displays the radial distribution function of hydrogen relative to the centre of mass of the benzene radical anion. At short distances, the probability is zero due to van der Waals repulsion. The distribution function then rises sharply to a value of ~ 1.7 for a distance of $\sim 1.8 \text{ \AA}$, indicating that it is 1.7 times more likely to find particles with this separation than in an ideal gas. This corresponds to water molecules that are located above or below the molecular plane. A second peak occurs at $\sim 3.2 \text{ \AA}$, which corresponds to water molecules located around the edge of the benzene molecule. The integral under a peak gives the number of solvent molecules of a given type. At long range the distribution function levels off to a value of 1, i.e. the particles no longer sense each other and behave as in an ideal gas.

For molecules, the radial distribution function can be extended with orientational degrees of freedom to characterize the angular distribution.

Correlation functions measure the relationship between two variables, x and y . A common definition is given in eq. (14.39).

$$C_{xy} = \frac{\langle (x - \langle x \rangle_M)(y - \langle y \rangle_M) \rangle_M}{\sqrt{\langle (x - \langle x \rangle_M)^2 \rangle_M \langle (y - \langle y \rangle_M)^2 \rangle_M}} \quad (14.39)$$

The correlation function is a number between -1 and 1 , where 1 indicates that the two quantities are completely correlated, -1 that they are (completely) anti-correlated, and 0 means that they are independent (uncorrelated).

Often such correlation functions are time dependent and measure how the correlation between two quantities changes over time. They may be normalized by the corresponding static (i.e. $t = t_0$) limit.

$$C_{xy}(t) = \frac{\langle x(t_0)y(t) \rangle_{N,t_0}}{\langle x(t_0)y(t_0) \rangle_{N,t_0}} \quad (14.40)$$

Notice that the averaging is done over the number of particles N and t_0 , but not the number of configurations M . Since an MD simulation produces a set of time-connected configurations, the number of a given configuration is directly related to the simulation time.

In the case where x and y are the same, $C_{xx}(t)$ is called an *autocorrelation* function; if they are different, it is called a *cross-correlation* function. For an autocorrelation function, the initial value at $t = t_0$ is 1, and it approaches 0 as $t \rightarrow \infty$. How fast it approaches 0 is measured by the *relaxation* time. The Fourier transforms of such correlation functions are often related to experimentally observed spectra; the far IR spectrum of a solvent, for example, is the Fourier transform of the dipole autocorrelation function.⁴⁹

$$I(\omega) = \int_{-\infty}^{+\infty} \langle \mu(t)\mu(0) \rangle e^{i\omega t} dt \quad (14.41)$$

14.5 Free Energy Methods

As noted in Section 13.6, it is difficult to calculate entropic quantities with any reasonable accuracy within a finite simulation time. It is, however, possible to calculate *differences* in such quantities.⁵⁰ Of special importance is the Gibbs free energy, since it is the natural thermodynamic quantity under normal experimental conditions (constant temperature and pressure, Table 14.1), but we will illustrate the principle with Helmholtz free energy instead (constant temperature and volume). As indicated in eq. (13.6) the fundamental problem is the same. There are two commonly used methods for calculating differences in free energy: *Thermodynamic Perturbation* and *Thermodynamic Integration*.⁵¹

14.5.1 Thermodynamic perturbation methods

The difference in entropy properties between two systems A and B can be calculated by an ensemble average, as discussed in Section 13.6.

$$\langle \Delta A_{A \rightarrow B} \rangle_M = -kT \ln \langle e^{-(E_B - E_A)/kT} \rangle_M \quad (14.42)$$

Since the energy difference must be small compared with kT , the transformation from A to B must usually be broken into several intermediate steps described by a λ parameter, and the total free energy change is given as the sum of changes in each step.

$$E_\lambda = \lambda E_A + (1 - \lambda) E_B \quad (14.43)$$

To test the quality of the averaging, the perturbation is usually run in both directions (i.e. $A \rightarrow B$ and $B \rightarrow A$), and the difference is taken as a measure of how well ΔA is

statistically converged. It should be noted that (too) short simulation times may lead to forward- and backward-calculated values that are in good agreement, without the energy difference being calculated accurately. Establishing a reliable estimate of the statistical error requires running several independent simulations and carefully analyzing the size of the perturbation steps and the correlation times for the various processes occurring in the system.⁵² Calculation of free energy differences by means of eq. (14.42) is often called *Thermodynamic Perturbation*⁵³ or *Free Energy Perturbation (FEP)*.

Instead of performing a series of simulations with a fixed energy function as in eq. (14.42), it may also be allowed to change *continuously* during a single simulation by changing λ slightly in each time step. This is called the *Slow Growth* method and requires that the increase in λ is slow enough that the system essentially remains at equilibrium at all times. This is difficult to ensure in practice,⁵⁴ and the slow growth method is therefore less commonly used.

14.5.2 Thermodynamic integration methods

Given an energy function as in eq. (14.43), the partition function, and thereby also the free energy, is a function of λ .

$$A(\lambda) = -kT \ln Q(\lambda) \quad (14.44)$$

Differentiating this expressions yields eq. (14.45).

$$\frac{\partial A}{\partial \lambda} = -\frac{kT}{Q} \frac{\partial Q}{\partial \lambda} = \frac{\partial E}{\partial \lambda} \quad (14.45)$$

Here the definition of Q (eq. (13.4)) has been used. Replacing the right-hand side by an ensemble average and integrating over λ gives eq. (14.46).

$$A(1) - A(0) = \int_0^1 \left\langle \frac{\partial E(\lambda)}{\partial \lambda} \right\rangle_M d\lambda \quad (14.46)$$

The left-hand side is the desired free energy difference, and the right-hand side may be approximated by a discrete sum.

$$\Delta A = \sum_i \left\langle \frac{\partial E(\lambda)}{\partial \lambda} \right\rangle_M \Delta \lambda_i \quad (14.47)$$

The use of eq. (14.47) for calculating ΔA is normally called *Thermodynamic Integration (TI)*.⁵⁵ The difference between eqs (14.42) and (14.47) is that the former averages over finite differences in energy functions, while the latter averages over a *differentiated* energy function. For parameterized energy functions, it is fairly easy to form the energy derivative with respect to the coupling parameter analytically, and the averaging in (14.47) is therefore no more complicated than averaging over energy differences as in eq. (14.42). Furthermore, it should be noted that the computational cost of performing the averaging is negligible compared with the cost of generating the ensemble, and the same ensemble can therefore be used to calculate the free energy difference by either eq. (14.42) or (14.47). This allows a measure of the reliability of the calculated value to be obtained.

Free energy calculations are often combined with *thermodynamic cycles* to calculate properties that would otherwise require impossible long simulation times.⁵⁶ A direct calculation of for example solvating acetone in water would require simulating the transfer of an acetone molecule from the gas phase (vacuum) to an aqueous phase, followed by solvent reorganization. If we wish to calculate the solvation energy of acetone relative to propane, this would require a second (impossibly long) simulation of transferring a propane molecule into the aqueous phase. Alternatively, the *difference* in solvation may be calculated by means of the thermodynamic cycle shown in Figure 14.7.




Figure 14.7 An example of a thermodynamic cycle for calculating differences in solvation energies

Since G is a state function, the difference in solvation energy, $\Delta G_{\text{solv}, A} - \Delta G_{\text{solv}, B}$, which is difficult to calculate, may instead be obtained as $\Delta\Delta G_{\text{solv}} - \Delta\Delta G_{\text{gas}}$. If **A** and **B** are different molecules, such as acetone and propane, the $\Delta\Delta G$ values correspond to non-physical transformations. Theoretically, however, it is quite easy to transform an oxygen atom into two hydrogens. The $\Delta\Delta G_{\text{gas}}$ value corresponds to differences in the internal (translational, rotational and vibrational) degrees of freedom, which can be calculated as discussed in Section 13.5. This difference also is part of $\Delta\Delta G_{\text{solv}}$, but if the internal energy levels are assumed to be independent of solvent, the solvent part of $\Delta\Delta G_{\text{solv}}$ is directly the difference in solvation.

In the acetone/propane example, the **A** to **B** change means that the oxygen atom gradually disappears, and two hydrogens gradually appear at the appropriate positions. In the force field energy expression, this corresponds to reducing or increasing van der Waals parameters and atomic charges, as well as changing all other parameters that are affected by the change in atom types. For $\lambda = 0.5$, the **A/B** “molecule” thus consists of a $\text{CH}_3\text{—C—CH}_3$ framework, with the central carbon having “half” a carbonyl oxygen and two “half” hydrogens attached. Absolute values of solvation energies may be calculated by transforming a solvent molecule to the solute, but if they are structurally very different it may require long simulation times to ensure that equilibrium is attained.

The technique of thermodynamic cycles may be used for calculating relative free energies for a variety of other cases. Differential binding of two ligands to an enzyme, for example, requires transforming one ligand into the other in a pure solution, and when bound to the enzyme. The strength of free energy methods is that differences in free energies may be obtained with a statistical accuracy of a few kJ/mol, at quite reasonable computational costs. Whether the calculated values agree with experimental results depends on the quality of the force field, but there are models for many solvents that are capable of providing an accuracy of better than a few kJ/mol in terms of absolute values.

The basic requirement of free energy perturbation or thermodynamic integration methods is that the non-physical transformation is carried out in sufficiently small steps that the sampling of the phase space at two successive points overlaps. Even for quite similar systems, this often means that the transformation must be broken into 10–20 steps, with each step requiring extensive sampling, and this makes such methods computationally intensive. In the *Linear Interaction Energy* (LIE) method, only the physical end-points for the thermodynamic cycle in Figure 14.7 are subjected to a simulation. The difference in binding free energy is then parameterized as a linear combination of the difference in the non-polar (van der Waal) and polar (electrostatic) interactions between the ligands and surroundings (enzyme or solvent).⁵⁷

$$\Delta G_{\text{bind}} = \alpha \Delta \langle E_{\text{vdw}} \rangle + \beta \Delta \langle E_{\text{el}} \rangle + \gamma \quad (14.48)$$

The β constant is expected to have a value of 0.5 from theoretical arguments. Optimization of the three parameters against experimental binding energies for the P450cam system confirms that the optimum β value is close to 0.5, while α and γ have values of ~ 0.18 and ~ -4.5 .⁵⁸ It is unclear to what extent these parameter values will depend on the specific system but the LIE offers a computational saving of an order of magnitude or more compared with FEP or TI methods.

14.6 Solvation Models

An important aspect of computational chemistry is to evaluate the effect of the environment, such as a solvent. Methods for evaluating the solvent effect may broadly be divided into two types: those describing the individual solvent molecules and those that treat the solvent as a continuous medium.⁵⁹ Combinations are also possible, for example by explicitly considering the first solvation shell and treating the rest by a continuum model. Each of these may be subdivided according to whether they use a classical or quantum mechanical description. By far the most important solvent is water, and since it is also one of the most difficult systems to model, the majority of methods have been focused on water, and we will use this for exemplification in the following.

The effects of solvation can be partitioned into two main groups:

- Non-specific (long-range) solvation
 - Polarization
 - Dipole orientation
- Specific (short-range) solvation
 - Hydrogen bonds
 - van der Waals interaction
 - Solvent shell structure
 - Solvent–solute dynamics
 - Charge transfer effects
 - Hydrophobic effects (entropy effects).

The non-specific effects are primarily solvent polarization and orientation of the solvent electric multipole moments by the solute, where the dipole interaction is usually the most important. These effects cause a screening of charge interactions, leading to the (macroscopic) dielectric constant being larger than 1. The microscopic interactions are primarily located in the first solvation shell, although the second

solvation shell may also be important for multiple-charged ions. The microscopic interactions depend on the specific nature of the solvent molecule, such as the shape and the ability to form hydrogen bonds.

A molecular description involves periodic boundary conditions and sampling the phase space by simulation methods. Such methods are in principle capable of accounting for all of the above solvent effects but the quality of the results will of course depend on how realistically the solvent–solute and solvent–solvent interactions are described. The requirement of many (hundreds or thousands of) solvent molecules to form a realistic model means that force field methods are often the primary choice from computational considerations. Since polarizable force fields are not yet in common use, this means that a major part of the non-specific solvation is lacking. Car–Parrinello methods using density functional theory for describing the interaction are significantly more expensive and can therefore only give a limited sampling of the phase space. They can account for the polarization but usually have a poor description of the van der Waals interaction. Semi-empirical electronic structure methods (AM1, PM3) are in general not sufficiently accurate for calculating intermolecular potentials. Mixed QM/MM methods, where the solute is described by a (quantum) electronic structure method and the solvent by a (classical) force field can account for the polarization of the solute, but the back-polarization again requires a polarizable force field.

Methods involving an explicit description of the solvent molecules require, analogously with other many-body methods, a sampling of the phase space. Since this is computationally expensive, there is a strong interest in developing methods where the solvent is modelled in a less rigorous fashion. The solvent–solute dynamics can be taken into account in an average fashion by the Langevin dynamics method (Section 14.2.3). The non-specific effects of solvation can be modelled by considering the solvent as a homogeneous medium with a dielectric constant, as will be discussed in more detail in the next section.

14.7 Continuum Solvation Models

Continuum models⁶⁰ consider the solvent as a uniform polarizable medium with a dielectric constant of ϵ , and with the solute **M** placed in a suitably shaped hole in the medium (Figure 14.8).⁶¹




Figure 14.8 Reaction field model

Creation of a hole in the medium costs energy, i.e. this is a destabilization, while dispersion interactions between the solvent and solute add a stabilization (this is roughly the van der Waals energy between solvent and solute). In principle, there may also be a repulsive component, thus the dispersion term is sometimes denoted dispersion/repulsion. The electric charge distribution of **M** will polarize the medium (induce charge moments), which in turn acts back on the molecule, thereby producing an electrostatic stabilization. The solvation (free) energy may thus be written as in eq. (14.49).

$$\Delta G_{\text{solvation}} = \Delta G_{\text{cavity}} + \Delta G_{\text{dispersion}} + \Delta G_{\text{elec}} \quad (14.49)$$

Reaction field models differ in five aspects:

- (1) How the size and shape of the hole are defined.
- (2) How the cavity/dispersion contribution is calculated.
- (3) How the charge distribution of **M** is represented.
- (4) How the solute **M** is described, either classical (force field) or quantum (semi-empirical or *ab initio*).
- (5) How the dielectric medium is described.

The dielectric medium is normally taken to have a constant value of ϵ , but may for some purposes also be taken to depend for example on the distance from **M**. For dynamical phenomena it can also be allowed to be frequency dependent,⁶² i.e. the response of the solvent is different for a “fast” reaction, such as an electronic transition, and a “slow” reaction, such as a molecular reorientation. It should be noted that ϵ is the *only* parameter characterizing the solvent, and solvents having the same ϵ value (such as acetone, $\epsilon = 20.7$, and 1-propanol, $\epsilon = 20.1$, or benzene, $\epsilon = 2.28$, and carbon tetrachloride, $\epsilon = 2.24$) are thus treated equally. The hydrogen bonding capability of 1-propanol compared with acetone will in reality most likely make a difference, and the solvent dynamics of an almost spherical CCl_4 will be different from the planar benzene molecule.

The simplest shape for the hole is a sphere or an ellipsoid. This has the advantage that the electrostatic interaction between **M** and the dielectric medium may be calculated analytically. More realistic models employ molecular shaped holes, generated for example by interlocking spheres located on each nucleus. Taking the atomic radius as a suitable factor (a typical value is 1.2) times a van der Waals radius defines a *van der Waals surface*. Such a surface may have small “pockets” where no solvent molecules can enter and a more appropriate descriptor may be defined as the surface traced out by a spherical particle of a given radius (a typical radius of 1.4 Å to model a water molecule) rolling on the van der Waals surface. This is denoted the *Solvent Accessible Surface* (SAS) and is illustrated in Figure 14.9.

Since an SAS is computationally more expensive to generate than a van der Waals surface, and since the difference is often small, a van der Waals surface is often used in practice. Furthermore, a very small displacement of an atom may alter the SAS in a discontinuous fashion, as a cavity suddenly becomes too small to allow a solvent molecule to enter. Alternatively, the cavity may be calculated directly from the wave function, for example by taking a surface corresponding to an electron density of 0.001.⁶³ It is generally found that the shape of the hole is important, and that molecular shaped cavities are necessary to be able to obtain good agreement with experimental data (such as solvation energies). It should be emphasized, however, that reaction field




Figure 14.9 On a surface generated by overlapping van der Waals spheres there will be areas (hatched) that are inaccessible to a solvent molecule (dotted sphere)

models are incapable of modelling specific (short-range) solvation effects, i.e. those occurring within the first solvation sphere.

The energy required to create the cavity (entropy factors and loss of solvent–solvent van der Waals interactions), and the stabilization due to van der Waals interactions between the solute and solvent (which may also contain a small repulsive component), is usually assumed to be proportional to the surface area. The corresponding energy terms may be taken simply as being proportional to the total SAS area (a single proportionality constant), or parameterized by having a constant ξ specific for each atom type (analogous to van der Waals parameters in force field methods), with the ξ parameters being determined by fitting to experimental solvation data.

$$G_{\text{cavity}} + \Delta G_{\text{dispersion}} = \gamma \text{SAS} + \beta \quad (14.50)$$

$$G_{\text{cavity}} + \Delta G_{\text{dispersion}} = \sum_i^{\text{atoms}} \xi_i S_i \quad (14.51)$$

For solvent models where the cavity/dispersion interaction is parameterized by fitting to experimental solvation energies, the use of a few explicit solvent molecules for the first solvation sphere is *not* recommended, as the parameterization represents a best fit to experimental data *without* any explicit solvent present. The electrostatic component of eq. (14.49) can be described at several different levels of approximation, as discussed in the following sections.

14.7.1 Poisson–Boltzmann methods

The *Poisson* equation is a second-order differential equation describing the connection between the electrostatic potential ϕ , the charge distribution ρ and the dielectric constant ϵ .⁶⁴

$$\nabla \cdot (\epsilon(\mathbf{r}) \nabla \phi(\mathbf{r})) = -4\pi \rho(\mathbf{r}) \quad (14.52)$$

Note that the dielectric “constant” may depend on the position. When it is independent of the position (i.e. truly a constant), eq. (14.52) becomes eq. (14.53).

$$\nabla^2 \phi(\mathbf{r}) = -\frac{4\pi}{\epsilon} \rho(\mathbf{r}) \quad (14.53)$$

If the charge distribution is a point charge, the solution of eq. (14.53) reduces to the Coulomb interaction. Eq. (14.52) can be used for describing for example the solvation of a protein in water, where the protein region is taken to have a low dielectric constant ($2 < \epsilon < 5$) while the solvent has a high dielectric constant ($\epsilon = 78$). The boundary between the two regions is typically taken as the SAS.

The Poisson equation can be modified by taking into account a (thermal) Boltzmann distribution of ions in the solvent. The negative ions will accumulate where the potential is positive, and vice versa, subject to a thermal fluctuation. The charge densities from a collection of ions with charges q and $-q$ and concentration c are given by eq. (14.54).

$$\begin{aligned}\rho_+ &= qce^{-q\phi/kT} \\ \rho_- &= -qce^{-q\phi/kT}\end{aligned}\quad (14.54)$$

Addition of these contributions to eq. (14.52) leads to the *Poisson–Boltzmann Equation* (PBE).

$$\begin{aligned}\nabla \cdot (\epsilon(\mathbf{r})\nabla\phi(\mathbf{r})) - \kappa^2 \left(\frac{kT}{q} \right) \sinh\left(\frac{q\phi(\mathbf{r})}{kT} \right) &= -4\pi\rho(\mathbf{r}) \\ \kappa^2 &= \frac{8\pi q^2 I}{kT}\end{aligned}\quad (14.55)$$

Here I is the ion strength of the solution, and the κ^2 factor is inversely related to the Debye–Hückel length, measuring how far the electrostatic effects extend into the solution. The $\sinh(q\phi(\mathbf{r})/kT)$ term only applies for the region corresponding to the solvent, i.e. for \mathbf{r} outside the cavity. Since $q\phi/kT$ is dimensionless, the PBE is often written in terms of a reduced potential u instead.

$$\nabla \cdot (\epsilon(\mathbf{r})\nabla u(\mathbf{r})) - \kappa^2 \sinh(u(\mathbf{r})) = -4\pi\rho(\mathbf{r}) \quad (14.56)$$

If the potential is sufficiently small (i.e. the solute is not strongly charged), the $\sinh(x)$ function can be expanded in a Taylor series, $\sinh(x) \approx x + x^3/6 + \dots$. Keeping only the first term gives the *Linearised Poisson–Boltzmann Equation* (LPBE).

$$\nabla \cdot (\epsilon(\mathbf{r})\nabla u(\mathbf{r})) - \kappa^2 u(\mathbf{r}) = -4\pi\rho(\mathbf{r}) \quad (14.57)$$

All of these equations ((14.52)–(14.57)) are differential equations that must be solved numerically, typically by a grid representation, and the results give information about the electrostatic potential at any point in space. It can be mapped onto the surface of the solute where it may suggest regions for interaction with other polar molecules. It can also be used for generating the reaction field, defined as the difference between the potential in the presence of a solvent ($\epsilon = 78$) and in vacuum ($\epsilon = 1$), i.e. $\phi_{\text{reac}} = \phi_{\text{solv}} - \phi_{\text{vac}}$. Multiplication of the reaction field with the solute charges in either a continuous (ρ) or partial charge (Q) description gives the electrostatic component of the free energy.

$$\begin{aligned}\Delta G_{\text{elec}} &= \frac{1}{2} \int \rho(\mathbf{r})\phi_{\text{reac}}(\mathbf{r})\mathbf{d}\mathbf{r} \\ \Delta G_{\text{elec}} &= \frac{1}{2} \sum_i Q(\mathbf{r}_i)\phi_{\text{reac}}(\mathbf{r}_i)\end{aligned}\quad (14.58)$$

14.7.2 Born/Onsager/Kirkwood models

The numerical aspects of solving the Poisson or Poisson–Boltzmann equations make them too demanding for use in connections with for example geometry optimizations of macromolecules. For certain special cases, however, the Poisson equation (14.52) can be solved analytical, and this forms the basis for many approximate models for estimating the electronic component in eq. (14.49).

The simplest reaction field model is a spherical cavity, where only the lowest order electric moment of the molecule is taken into account. For a net charge q in a cavity of radius a , the difference in energy between a vacuum and a medium with a dielectric constant of ϵ is given by the *Born* model.⁶⁵

$$\Delta G_{\text{elec}}(q) = -\left(1 - \frac{1}{\epsilon}\right) \frac{q^2}{2a} \quad (14.59)$$

It can be noted that the Born model predicts equal solvation energies for positive and negative ions of the same size, which is not the observed behaviour in solvents such as water. Furthermore, the reciprocal dependence on the dielectric constant means that the calculated solvent effect is sensitive to the variation of ϵ in the low dielectric limit but virtually unaffected by large differences in the high dielectric limit. Changing ϵ from 1 to 2 gives a factor of $1/2$ in eq. (14.59) but there is virtually no difference between a solvent with a dielectric constant of 30 (e.g. acetonitrile) and one with a dielectric constant of 78 (e.g. water), although in actual experiments there may be a significant difference.

Using partial atomic charges in eq. (14.59) is often called the *generalized Born* model, which has been used especially in connection with force field methods in the *Generalized Born/Surface Area* (GB/SA) model.⁶⁶ In this case, the Coulomb interaction between the partial charges (eq. (2.20)) is combined with the Born formula by means of a function f_{ij} depending on the internuclear distance and Born radii for each of the two atoms, a_i and a_j .

$$\begin{aligned} G_{\text{elec}}(Q_i, Q_j) &= -\left(1 - \frac{1}{\epsilon}\right) \frac{Q_i Q_j}{f_{ij}} \\ f_{ij} &= \sqrt{r_{ij}^2 + a_{ij}^2} e^{-D} \\ a_{ij}^2 &= a_i a_j \\ D &= \frac{r_{ij}^2}{4a_{ij}^2} \end{aligned} \quad (14.60)$$

The effective Born radius for a given atom depends on the nature and position of all the atoms. In practice, the dependence on the other atoms is relatively weak, and updates of the a_i parameters can be done at suitable intervals, for example when updating the non-bonded list in an optimization or simulation. The boundary between the solute and solvent is usually taken as a modified van der Waals surface generated from the unification of atomic van der Waals radii scaled by a suitable factor. The cavity/dispersion terms are parameterized according to the SAS, as in eq. (14.51). The GB/SA model provides a very fast method of incorporating solvent effects, and it is furthermore relatively easy to formulate gradients of the energy function, making it possible

to perform optimization and simulations. It has been shown to reproduce the results from Poisson–Boltzmann calculations rather accurately, but it should be noted that the results are somewhat sensitive to the magnitude of the partial charges.

The dipole in a spherical cavity is known as the *Onsager* model,⁶⁷ which for a dipole moment of μ leads to an energy stabilization given by eq. (14.61).

$$\Delta G_{\text{elec}}(\mu) = -\frac{\epsilon - 1}{2\epsilon + 1} \frac{\mu^2}{a^3} \quad (14.61)$$

The *Kirkwood* model⁶⁸ refers to a general multipole expansion in a spherical cavity, while the *Kirkwood–Westheimer* model arises for an ellipsoidal cavity.⁶⁹

The charge distribution of the molecule can be represented either as atom-centred partial charges or as a multipole expansion. For a neutral molecule, the lowest order approximation considers only the dipole moment. This may be a quite poor approximation, and fails completely for symmetric molecules that do not have a dipole moment. For obtaining converged results, it is often necessarily to extend the expansion up to order six or more, i.e. including dipole, quadrupole, octupole, etc., moments. Furthermore, only for small and symmetric molecules can the approximation of a spherical or ellipsoidal cavity be considered realistic. The use of the Born/Onsager/Kirkwood models should therefore only be considered as a rough estimate of the solvent effects, and quantitative results can rarely be obtained.

14.7.3 Self-consistent reaction field models

A classical description of the molecule **M** in Figure 14.9 can be a force field with (partial) atomic charges, while a quantum description involves calculation of the electronic wave function. The latter may be either a semi-empirical model, such as AM1 or PM3, or more sophisticated electronic structure methods, i.e. HF, DFT, MCSCF, MP2, CCSD, etc. When a quantum description of **M** is employed, the calculated electric moments induce charges in the dielectric medium, which in turn acts back on the molecule, causing the wave function to respond and thereby changing the electric moments, etc. The interaction with the solvent model must thus be calculated by an iterative procedure, leading to various *Self-Consistent Reaction Field* (SCRF) models.

The interaction of a fixed dipole moment with a polarizable medium is given by eq. (14.61). This, however, is not an SCRF model, as the dipole moment and stabilization are not calculated in a self-consistent way. When the back-polarization of the medium is taken into account, the dipole moment changes, depending on how polarizable the molecule is. Taking only the first-order effect into account, the stabilization is given by eq. (14.62).

$$\Delta G_{\text{elec}}(\mu) = -\frac{\epsilon - 1}{2\epsilon + 1} \frac{\mu^2}{a^3} \left[1 - \frac{\epsilon - 1}{2\epsilon + 1} \frac{2\alpha}{a^3} \right]^{-1} \quad (14.62)$$

Here α is the molecular polarizability, i.e. the first-order change in the dipole moment with respect to an electric field. In the SCRF model the full polarization is taken into account, i.e. the initial dipole moment generates a polarization of the medium, which changes the dipole moment, which in turn generates a slightly different polarization, etc.

For spherical or ellipsoidal cavities the Poisson equation can be solved analytically, but for molecular shaped surfaces it must be done numerically. This is typically done by reformulating it in terms of a surface integral over surface charges, and solving this numerically by dividing the surface into smaller fractions each having an associated charge $\sigma(\mathbf{r}_s)$. The surface charges are related to the electric field \mathbf{F} (the derivative of the potential ϕ) perpendicular to the surface by eq. (14.63).

$$4\pi\epsilon\sigma(\mathbf{r}_s) = (\epsilon - 1)\mathbf{F}(\mathbf{r}_s) \quad (14.63)$$

Once $\sigma(\mathbf{r}_s)$ is determined the associated potential is added as an extra term to the Hamiltonian operator.

$$\begin{aligned} \mathbf{H} &= \mathbf{H}_0 + \phi_\sigma \\ \phi_\sigma(\mathbf{r}) &= \int \frac{\sigma(\mathbf{r}_s)}{|\mathbf{r} - \mathbf{r}_s|} d\mathbf{r}_s \end{aligned} \quad (14.64)$$

The potential ϕ_σ from the surface charge is given by the molecular charge distribution (eq. (14.64)), but also enters the Hamiltonian and thus influences the molecular wave function. The procedure is therefore iterative.

For the case of the Onsager model (spherical cavity, dipole moment only) the term added to the molecular Hamiltonian operator is given by eq. (14.65).

$$\phi_\sigma = -\mathbf{r} \cdot \mathbf{R} \quad (14.65)$$

Here \mathbf{r} is the dipole moment *operator* (i.e. the position vector), and \mathbf{R} is proportional to the molecular dipole moment, with the proportionality constant depending on the radius of the cavity and the dielectric constant.

$$\begin{aligned} \mathbf{R} &= g\boldsymbol{\mu} \\ g &= \frac{2(\epsilon - 1)}{(2\epsilon + 1)a^3} \end{aligned} \quad (14.66)$$

At the HF level of theory, the ϕ_σ operator corresponds to the addition of an extra term to the Fock matrix elements (Section 3.5).

$$F_{\alpha\beta} = \langle \chi_\alpha | \mathbf{F} | \chi_\beta \rangle - g\boldsymbol{\mu} \langle \chi_\alpha | \mathbf{r} | \chi_\beta \rangle \quad (14.67)$$

The additional integrals are just expectation values of x , y and z coordinates, and their inclusion requires very little additional computational effort. Generalization to higher order multipoles is straightforward.

In connection with electronic structure methods (i.e. a quantum description of \mathbf{M}), the term SCRF is quite generic and it does not by itself indicate a specific model. Typically, however, the term is used for models where the cavity is either spherical or ellipsoidal, the charge distribution is represented as a multipole expansion, often terminated at quite low orders (for example only including the charge and dipole terms), and the cavity/dispersion contributions are neglected. Such a treatment can only be used for a qualitative estimate of the solvent effect, although relative values may be reasonably accurate if the molecules are polar (dominance of the dipole electrostatic term) and sufficiently similar in size and shape (cancellation of the cavity/dispersion terms).

The cavity size in the Born/Onsager/Kirkwood models strongly influences the calculated stabilization. Unfortunately, there is no consensus on how to choose the cavity radius. In some cases, the molecular volume is calculated from the experimental density of the solvent and the cavity radius is defined by equating the cavity volume to the molecular volume. Alternatively, the cavity size may be derived from the (experimental) dielectric constant and the calculated dipole moment and polarizability.⁷⁰ In any case, the underlying assumption of these models is that the molecule is roughly spherical or ellipsoidal, which is only generally true for small compact molecules.

More sophisticated models employ molecular shaped cavities, but there is again no consensus on the exact procedure. The cavity is often defined based on van der Waals radii of the atoms in the molecule multiplied with an empirical scale factor. Alternatively, the molecular volume may be calculated directly from the electronic wave function, for example by using a contour surface corresponding to an electron density of 0.001.

The *Polarizable Continuum Model* (PCM) employs a van der Waals cavity formed by interlocking atomic van der Waals radii scaled by an empirical factor, a detailed description of the electrostatic potential, and parameterizes the cavity/dispersion contributions based on the surface area.⁷¹ The *Conductor-like Screening Model* (COSMO) also employs molecular shaped cavities, and represents the electrostatic potential by partial atomic charges. COSMO was originally implemented for semi-empirical methods but has also been used in connection with *ab initio* methods.⁷² It may be considered as a limiting case of the PCM model, where the dielectric constant is set to infinity. The *Solvation Models* (SM_x, where $x = 1-5$) developed by Cramer and Truhlar are generalized Born type models used in connection with the semi-empirical AM1 and PM3 methods.⁷³ The partial atomic charges are calculated from the wavefunction, and the dispersion/cavity terms in eq. (14.49) are parameterized based on the solvent exposed surface area, eq. (14.51). The version number of these models reflects increasingly sophisticated parameterizations.

The “mixed” solvent models, where the first solvation shell is accounted for by including a number of solvent molecules, implicitly include the solute–solvent cavity/dispersion terms, although the corresponding terms between the solvent molecules and the continuum are usually neglected. Once discrete solvent molecules are included, however, the problem of configurations sampling arises. Furthermore, a parameterization of the continuum model against experimental data must be done by explicitly taking the first solvation shell into account. Nevertheless, in many cases, the first solvation shell is by far the most important, and mixed models may yield substantially better results than pure continuum models, at the price of an increased computational cost.

Given the diversity of the various SCRF models, and the fact that solvation energies in water may range from a few kJ/mol for say ethane to perhaps several hundred kJ/mol for an ion, it is difficult to evaluate just how accurately continuum methods may in principle be able to represent solvation. It seems clear, however, that molecular shaped cavities must be employed, the electrostatic polarization needs a description either in terms of atomic charges or quite high order multipoles, and cavity and dispersion terms must be included. Properly parameterized, such models appear to be able to give absolute values with an accuracy of a few kJ/mol.⁷⁴ Comparison with results obtained by explicit solvent modelling, however, suggests that the electrostatic

component is underestimated by continuum models by roughly a factor of 2, while the non-bonded part is essentially uncorrelated with the surface area.⁷⁵

Inclusion of solvent effects may change the geometry, charge distribution and conformational preferences. Employing a PCM type solvation water model in connection with the B3LYP/aug-cc-pVTZ method for example leads to an increase of the C=O bond length in acetamide by 0.015 Å, while the C—N bond is reduced by a similar amount. The calculated dipole moment correspondingly changes from 3.9 to 5.2 debye. Since solvation preferentially stabilizes the more polar systems, it may also change the conformational preference of molecules. Using the above computational model for example changes the energy difference between the *anti* (no dipole moment by symmetry) and *gauche* (gas phase dipole moment of 2.8 debye) conformations of 1,2-dichloroethane from 6.7 kJ/mol in the gas phase to 1.1 kJ/mol in solution. Molecular properties are in many cases also sensitive to the environment, but a detailed discussion of this is outside the scope of this book.²³

References

1. J. M. Haile, *Molecular Dynamics Simulation*, Wiley, 1991; M. P. Allan, D. J. Tildesley, *Computer Simulations of Liquids*, Clarendon Press, Oxford, 1987; W. F. van Gunsteren, H. J. C. Berendsen, *Ang. Chem. Int. Ed.*, **29** (1990), 992; A. R. Leach, *Molecular Modelling. Principles and Applications*, Longman, 1996; D. Frenkel, B. Smith, *Understanding Molecular Simulations*, 2nd edition, Academic Press, 2002.
2. W. L. Jorgensen, *Adv. Chem. Phys.*, **70** (1988), 469.
3. N. A. Metropolis, A. W. Rosenbluth, A. H. Teller, E. Teller, *J. Chem. Phys.*, **21** (1953), 1087.
4. W. F. van Gunsteren, A. E. Mark, *J. Chem. Phys.*, **108** (1998), 6109.
5. S. Duane, A. Kennedy, B. J. Pendleton, D. Roweth, *Phys. Lett. B*, **195** (1987), 216; R. Faller, J. J. de Pablo, *J. Chem. Phys.*, **116** (2002), 55.
6. V. I. Maniushchuk, M. W. Deam, *J. Chem. Phys.*, **110** (1999), 2753.
7. I. R. McDonald, *Mol. Phys.*, **23** (1972), 41.
8. L. Verlet, *Phys. Rev.*, **159** (1967), 98.
9. M. P. Allen, D. J. Tildesley, *Computer Simulation of Liquids*, Clarendon, 1989.
10. H. C. Andersen, *J. Chem. Phys.*, **72** (1980), 2384.
11. S. Toxværd, *Phys. Rev. E*, **47** (1993), 343.
12. J. P. Ryckaert, G. Ciccotti, H. J. C. Berendsen, *J. Comput. Phys.*, **23** (1977), 327; D. J. Tobias, C. L. Brooks III, *J. Chem. Phys.*, **89** (1988), 5115.
13. H. C. Andersen, *J. Comput. Phys.*, **52** (1983), 24.
14. H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, J. R. Haak, *J. Chem. Phys.*, **81** (1984), 3684.
15. S. Nosé, *Mol. Phys.*, **52** (1984), 255; W. G. Hoover, *Phys. Rev. A*, **31** (1985), 1695.
16. M. P. Allen, D. J. Tildesley, *Computer Simulations of Liquids*, Clarendon Press, 1987; A. R. Leach, *Molecular Modelling. Principles and Applications*, Longman, 1996.
17. D. G. Truhlar, R. Steckler, M. S. Gordon, *Chem. Rev.*, **87** (1987), 217; L. Sun, W. L. Hase, *Rev. Comp. Chem.*, **19** (2003), 17.
18. A. Gonzalez-Lafont, T. N. Truong, D. G. Truhlar, *J. Phys. Chem.*, **95** (1991), 4618; Y. Chuang, M. L. Radhakrishnan, P. L. Fast, C. J. Cramer, D. G. Truhlar, *J. Phys. Chem.*, **103** (1999), 4893; G. H. Peslherbe, W. L. Hase, *J. Chem. Phys.*, **104** (1996), 7882.
19. X. Li, J. M. Millam, H. B. Schlegel, *J. Chem. Phys.*, **113** (2000), 10062.
20. R. Car, M. Parrinello, *Phys. Rev. Lett.*, **55** (1985), 2471; D. K. Remler, P. A. Madden, *Mol. Phys.*, **70** (1990), 921; J. S. Tse, *Ann. Rev. Phys. Chem.*, **53** (2002), 249.

21. P. Tangney, S. Scandolo, *J. Chem. Phys.*, **116** (2002), 14.
22. H. B. Schlegel, S. S. Iyengar, X. Li, J. M. Millam, G. A. Voth, G. E. Scuseria, M. J. Frisch, *J. Chem. Phys.*, **117** (2002), 8694; J. M. Herbert, M. Head-Gordon, *J. Chem. Phys.*, **121** (2004), 11542.
23. G. D. Billing, K. V. Mikkelsen, *Introduction to Molecular Dynamics and Chemical Kinetics*, Wiley, 1996; G. D. Billing, K. V. Mikkelsen, *Advanced Molecular Dynamics and Chemical Kinetics*, Wiley, 1997.
24. G. D. Billing, *Phys. Chem. Chem. Phys.*, **4** (2002), 2865.
25. D. G. Truhlar, R. Steckler, M. S. Gordon, *Chem. Rev.*, **87** (1987), 217; T.-S. Ho, T. Hollebeek, H. Rabitz, L. B. Harding, G. C. Schatz, *J. Chem. Phys.*, **105** (1996), 10462.
26. M. A. Collins, *Adv. Chem. Phys.*, **93** (1996), 389.
27. Y.-P. Liu, D.-H. Lu, A. Gonzales-Lafont, D. G. Truhlar, B. C. Garrett, *J. Am. Chem. Soc.*, **115** (1993), 7806.
28. G. E. Moyano, M. Collins, *J. Chem. Phys.*, **121** (2004), 9769.
29. G. D. Billing, *Mol. Phys.*, **89** (1996), 355.
30. R. P. Bell, *The Tunnel Effect in Chemistry*, Chapman and Hall, 1980.
31. E. Wigner, *Z. Phys. Chem. B*, **19** (1932), 203.
32. Y.-P. Liu, D.-H. Lu, A. Gonzales-Lafont, D. G. Truhlar, B. C. Garrett, *J. Am. Chem. Soc.*, **115** (1993), 7806; J. C. Corchado, J. Espinosa-Garcia, W.-P. Hu, I. Rossi, D. G. Truhlar, *J. Phys. Chem.*, **99** (1997), 687.
33. A. D. Bochevarov, E. F. Valeev, D. C. Sherrill, *Mol. Phys.*, **102** (2004), 111.
34. E. Deumens, Y. Öhrn, *J. Phys. Chem. A*, **105** (2001), 2660.
35. G. M. Torrie, J. P. Valleau, *Chem. Phys. Lett.*, **28** (1974), 278.
36. E. A. Carter, G. Ciccotti, J. A. Hynes, R. Kapral, *Chem. Phys. Lett.*, **156** (1989), 472; G. Ciccotti M. Ferrario, *Mol. Sim.*, **30** (2004), 787.
37. R. Rajamani, K. J. Naidoo, J. Gao, *J. Comp. Chem.*, **24** (2003), 1775.
38. H. Bekker, *J. Comp. Chem.*, **18** (1997), 1930.
39. C. L. Brooks III, B. M. Pettitt, M. Karplus, *J. Chem. Phys.*, **85** (1985), 5897; M. Patra, M. Karttunen, M. T. Hyvönen, E. Falck, I. Vattulainen, *Condensed Matter*, Los Alamos National Laboratory, Preprint Archive, 2004; M. Bergdorf, C. Peter, P. H. Hunenberger, *J. Chem. Phys.*, **119** (2003), 9129.
40. P. J. Steinback, B. R. Brooks, *J. Comp. Chem.*, **15** (1994), 667.
41. P. P. Ewald, *Ann. Phys.*, **64** (1921), 253; V. Natoli, D. M. Ceperley, *J. Comp. Phys.*, **117** (1995), 171; Z.-M. Chen, T. Cagin, W. A. Goddard III, *J. Comp. Chem.*, **18** (1997), 1365.
42. H. G. Petersen, *J. Chem. Phys.*, **103** (1995), 3668.
43. L. Greengard, V. Rokhlin, *J. Comput. Phys.*, **73** (1987), 325.
44. H. G. Petersen, D. Solovason, J. W. Perram, E. R. Smith, *J. Chem. Phys.*, **101** (1994), 8870.
45. J. S. Bader, D. Chandler, *J. Phys. Chem.*, **96** (1992), 6424.
46. P. E. Smith, B. M. Pettitt, *J. Chem. Phys.*, **95** (1991), 8430.
47. H. Flyvbjerg, H. G. Pedersen, *J. Chem. Phys.*, **91** (1998), 461.
48. K. V. Mikkelsen, P. Linse, P.-O. Åstrand, K. Karlström, *J. Phys. Chem.*, **98** (1994), 8209.
49. B. Guillot, *J. Chem. Phys.*, **95** (1991), 1543.
50. T. P. Straatsma, *Rev. Comp. Chem.*, **9** (1996), 81.
51. C. Chipot, D. A. Perlman, *Mol. Simul.*, **28** (2002), 1.
52. M. R. Shirts, J. W. Pitera, W. C. Swope, V. S. Pande, *J. Chem. Phys.*, **119** (2003), 5740.
53. R. W. Zwanzig, *J. Chem. Phys.*, **22** (1954), 1420.
54. D. A. Pearlman, P. A. Kollman, *J. Chem. Phys.*, **91** (1989), 7831.
55. J. G. Kirkwood, *J. Chem. Phys.*, **3** (1935), 300.
56. P. A. Kollman, K. M. Merz Jr, *Acc. Chem. Res.*, **23** (1990), 246; P. Kollman, *Chem. Rev.*, **93** (1993), 2395.
57. J. Åqvist, C. Medina, J. E. Samuelsson, *Prot. Eng.*, **7** (1994), 385.

58. M. Almlöf, B. O. Brandsdal, J. Åqvist, *J. Comp. Chem.*, **25** (2004), 1242.
59. K. V. Mikkelsen, H. Ågren, *J. Mol. Struct. (Theochem)*, 234 (1991), 425; C. J. Cramer, D. G. Truhlar, *Chem. Rev.*, **99** (1999), 2161; P. E. Smith, B. M. Pettitt, *J. Phys. Chem.*, **98** (1994), 9700.
60. B. Roux, T. Simonson, *Biophys. Chem.*, **78** (1999), 1.
61. J. Tomasi, M. Persico, *Chem. Rev.*, **94** (1994), 2027.
62. K. V. Mikkelsen, K. O. Sylvester-Hvid, *J. Phys. Chem.*, **100** (1996), 9116.
63. J. B. Foresman, T. A. Keith, K. B. Wiberg, J. Snoonian, M. J. Frisch, *J. Phys. Chem.*, **100** (1996), 16098.
64. G. Lamm, *Rev. Comp. Chem.*, **19** (2003), 147; N. A. Baker, *Rev. Comp. Chem.*, **21** (2005), 349.
65. M. Born, *Z. Physik*, **1** (1920), 45.
66. W. C. Still, A. Tempczyk, R. C. Hawley, T. Hendrickson, *J. Am. Chem. Soc.*, **112** (1990), 6127; D. Bashford, D. A. Case, *Ann. Rev. Phys. Chem.*, **51** (2000), 129.
67. L. Onsager, *J. Am. Chem. Soc.*, **58** (1936), 1486.
68. J. G. Kirkwood, *J. Chem. Phys.*, **2** (1934), 351.
69. J. G. Kirkwood, F. H. Westheimer, *J. Chem. Phys.*, **6** (1936), 506.
70. Y. Luo, H. Ågren, K. V. Mikkelsen, *Chem. Phys. Lett.*, **275** (1997), 145.
71. M. Cossi, V. Barone, R. Cammi, J. Tomasi, *Chem. Phys. Lett.*, **255** (1996), 327.
72. A. Klamt, *J. Phys. Chem.*, **99** (1995), 2224; T. Truong, E. V. Stefanovich, *Chem. Phys. Lett.*, **240** (1995), 253.
73. C. J. Cramer, D. G. Truhlar, *Rev. Comp. Chem.*, **6** (1995), 1.
74. V. Barone, M. Cossi, J. Tomasi, *J. Chem. Phys.*, **107** (1997), 3210.
75. J. Wagoner, N. A. Baker, *J. Comp. Chem.*, **25** (2004), 123; *J. Comp. Chem.*, **26** (2004), 863.

15 Qualitative Theories

Although sophisticated electronic structure methods may be able to accurately predict a molecular structure or the outcome of a chemical reaction, the results are often hard to rationalize. Generalizing the results to other similar systems therefore becomes difficult. Qualitative theories, on the other hand, are unable to provide accurate results but they may be useful for gaining *insight*, for example why a certain reaction is favoured over another. They also provide a link to many concepts used by experimentalists. Frontier molecular orbital theory considers the interaction of the orbitals of the reactants and attempts to predict relative reactivities by second-order perturbation theory. It may also be considered as a simplified version of the Fukui function, which considered how easily the total electron density can be distorted. The Woodward–Hoffmann rules allow a rationalization of the stereochemistry of certain types of reactions, while the more general qualitative orbital interaction model can often rationalize the preference for certain molecular structures over other possible arrangements.

15.1 Frontier Molecular Orbital Theory

Frontier Molecular Orbital (FMO) theory attempts to predict relative reactivity based on properties of the reactants. It is commonly formulated in term of perturbation theory, where the energy change in the initial stage of a reaction is estimated and “extrapolated” to the transition state.¹ For a reaction where two different modes of reaction are possible, this may be illustrated as shown in Figure 15.1.

The reaction mode that involves the *least* energy change in the initial stage is assumed also to have the lowest activation energy. FMO theory uses a low-order perturbation expansion with the reactants as the unperturbed reference, and it is clear that such a treatment can only be used to follow the reaction a short part of the whole reaction pathway.

The change in the energy can be derived from second-order perturbation theory (Section 4.8) and is given in equation (15.1).²




Figure 15.1 FMO region of a reaction profile

$$\Delta E = - \sum_{A,B}^{\text{atoms}} (\rho_A + \rho_B) \langle \chi_A | \mathbf{V} | \chi_B \rangle \langle \chi_A | \chi_B \rangle + \sum_{A,B}^{\text{atoms}} \frac{Q_A Q_B}{R_{AB}} + \left(\sum_{i \in A}^{\text{occ. MO}} \sum_{a \in B}^{\text{vir. MO}} + \sum_{i \in B}^{\text{occ. MO}} \sum_{a \in A}^{\text{vir. MO}} \right) \frac{2 \left(\sum_{\alpha}^{\text{AO}} c_{\alpha i} c_{\alpha a} \langle \chi_{\alpha i} | \mathbf{V} | \chi_{\alpha a} \rangle \right)^2}{\epsilon_i - \epsilon_a} \quad (15.1)$$

Here A and B denote atoms in each of the two interacting molecules. The \mathbf{V} operator contains all the potential energy operators from both molecules, and the $\langle \chi_A | \mathbf{V} | \chi_B \rangle$ integral is a “resonance” type integral between two atomic orbitals, one from each molecule. The ρ_A is the electron density on atom A, and the first term in (15.1) represents a repulsion ($\langle \chi_A | \mathbf{V} | \chi_B \rangle$ is a negative quantity) between occupied MOs (steric repulsion). This will usually lead to a net energy barrier for a reaction. The second term represents an attraction or repulsion between charged parts of the molecules, Q_A being the (net) charge on atom A. The last term is a stabilizing interaction ($\epsilon_i - \epsilon_a < 0$) due to mixing of occupied MOs on one molecule with unoccupied MOs on the other, $c_{\alpha i}/c_{\alpha a}$ being MO coefficients and ϵ_i/ϵ_a MO energies. The summation is over all pairs of occupied/unoccupied MOs.

If we are comparing reactions that have approximately the same steric requirements, the first term is roughly constant. If the species are very polar the second term will dominate, and the reaction is *charge controlled*. This means for example that an electrophilic attack is likely to occur at the most negative atom or, in a more general sense, along a path where the electrostatic potential is most negative. If the molecules are non-polar, the third term in eq. (15.1) will dominate and the reaction is said to be *orbital controlled*. This means that the reaction will occur where the molecular orbital coefficients are largest.

All other things being equal, the largest contribution to the double summation over orbital pairs in the third term will arise when the denominator is smallest. This corresponds to the *Highest Occupied Molecular Orbital* (HOMO) and *Lowest Unoccupied Molecular Orbital* (LUMO) pair of orbitals. FMO theory considers only this one contribution in the whole summation. From a purely numerical consideration this is certainly not a good approximation: the contributions from all the other pairs are much

larger than the single HOMO–LUMO term. Nevertheless, it is possible to rationalize many trends in terms of FMO theory and thus the result justifies the means. If we furthermore consider a matrix element $\langle \chi_{ai} | \mathbf{V} | \chi_{aj} \rangle$ to be non-zero only between atoms where new bonds are being formed (where it is furthermore assumed to be roughly constant), the deciding factor becomes a sum over products of MO coefficients from the HOMO on one fragment with LUMO coefficients on the other. A few examples should help clarify this.

The reaction of a nucleophile involves the addition of electrons to the reactant, i.e. interaction of the HOMO of the nucleophile with the LUMO of the reactant. If there is more than one possible centre of attack, the preferred reaction mode is predicted to occur on the atom having the largest LUMO coefficient. Figure 15.2 shows that the orbital component shows preference for addition to the 4-position of acrolein (as a model for unsaturated carbonyl compounds in general), with the second most reactive position being C₂. The net charges, however, prefer position 2, as it is the most positive carbon. Experimentally, it is found that attack at the 4-position is usually favoured (especially with “soft” nucleophiles such as organocuprates), but addition at the 2-position is also observed (and may dominate with “hard” nucleophiles such as organolithium compounds).³ This is consistent with the reaction switching from being orbital controlled to charge controlled as the nucleophile becomes more ionic.




Figure 15.2 AM1 LUMO coefficients for acrolein with net charges in parenthesis

Similarly, the reaction of an electrophile will involve the HOMO of the reactant, i.e. the reaction should occur preferentially on the atom having the largest HOMO coefficient. The coefficients for furan shown in Figure 15.3 indicate that electrophilic substitution should preferentially occur at the 2-position, again in agreement with experimental results.⁴




Figure 15.3 AM1 HOMO coefficients for furan

Consider now the reaction between butadiene and ethylene, where both 2+2 and 4+2 reaction modes are possible. The qualitative appearances of the butadiene HOMO and ethylene LUMO are given in Figure 15.4. The MO coefficients are given as a , b and c , where $a > b > c$.




Figure 15.4 FMO theory favours the 4+2 over the 2+2 reaction

For the 2+2 pathway the FMO sum becomes $(ab - ac)^2 = a^2(b - c)^2$ while for the 4+2 reaction it is $(ab + ab)^2 = a^2(2b)^2$. As $(2b)^2 > (b - c)^2$, it is clear that the 4+2 reaction has the largest stabilization, and therefore increases *least* in energy in the initial stages of the reaction (eq. (15.1), remembering that the steric repulsion will cause a net increase in energy). The 4+2 reaction should consequently have the lowest activation energy, and therefore occur more easily than the 2+2. This is indeed what is observed: the Diels–Alder reaction occurs readily but cyclobutane formation is not observed between non-polar dienes and dieneophiles.

The appearance of the difference in MO energies in the denominator in eq. (15.1) suggests that a smaller gap between the diene HOMO and dieneophile LUMO in a Diels–Alder reaction should lower the activation energy. If the diene is made more electron-rich (electron-donating substituents), or the dieneophile more electron-deficient (electron-withdrawing substituents), the reaction should proceed faster. This is indeed the observed trend. For the reaction between cyclopentadiene and cyanoethylenes (mono-, di-, tri- and tetra-substituted), the correlation is reasonably quantitative, as shown in Figure 15.5.⁵

This is of course a rather extreme example, as the reaction rates differ by $\sim 10^7$, and rate differences of over a factor of 100 are observed for quite similar HOMO–LUMO differences. For a more varied set of compounds where the reaction rates are more similar, the correlation is often quite poor.

FMO theory can also be used for explaining the stereochemistry of the Diels–Alder reaction, as can be illustrated by the reaction between 2-methylbutadiene and cyanoethylene. These may react to give two different products, the “*para*” and/or “*meta*” isomer.

The MO coefficients for the p-orbitals on the butadiene HOMO and ethylene LUMO (taken from AM1 calculations) are given in Figure 15.6. The FMO sum for the




Figure 15.5 Correlation between reaction rates and FMO energy differences (AM1 calculations)




Figure 15.6 FMO rationalizes the stereochemistry of substituted Diels-Alder reactions

“*para*” isomer is $(0.594 \times 0.682 + 0.517 \times 0.552)^2 = 0.690$, while the sum for the “*meta*” isomer is $(0.594 \times 0.552 + 0.517 \times 0.682)^2 = 0.680$. FMO theory thus predicts that the “*para*” isomer should dominate, as is indeed observed (experimental ratio 70:30). If cyanoethylene is replaced by 1,1-dicyanoethylene, the LUMO coefficients change to 0.708 and -0.511 . The corresponding “*para*” and “*meta*” FMO sums change to 0.685 and 0.670, i.e. a larger difference between the two isomers. This is again reflected in the experimental data, where the ratio is 91:9. The regiochemistry is thus determined by matching the two largest sets of coefficients and the two smaller sets, rather than making two sets of large/small.

FMO theory was developed at a time when detailed calculations of reaction paths were unfeasible. As many sophisticated computational models, and methods for actually locating the transition state, have become widespread, the use of FMO arguments for predicting reactivity has declined. The primary goal of computational chemistry, however, is not to provide *numbers*, but to provide *understanding*. As such, FMO

theory still forms a conceptual model that can be used for rationalizing trends without having to perform time-consuming calculations.

15.2 Concepts from Density Functional Theory

The success of FMO theory is not because the neglected terms in the second-order perturbation expansion (eq. (15.1)) are especially small; an actual calculation will reveal that they completely swamp the HOMO–LUMO contribution. The deeper reason is that the shapes of the HOMO and LUMO resemble features in the total electron density, which determines the reactivity. There are also other quantities derived from density functional theory that directly relate to the properties and reactivity of molecules, and these are discussed in this section.⁶

A reaction will in general involve a change in the electron density, which may be quantified in terms of the *Fukui function*.⁷

$$f(\mathbf{r}) = \frac{\partial \rho(\mathbf{r})}{\partial N_{\text{elec}}} \quad (15.2)$$

The Fukui function indicates the change in the electron density at a given position when the number of electrons is changed. We may define two finite difference versions of the function, corresponding to addition or removal of an electron.

$$\begin{aligned} f_+(\mathbf{r}) &= \rho_{N+1}(\mathbf{r}) - \rho_N(\mathbf{r}) \\ f_-(\mathbf{r}) &= \rho_N(\mathbf{r}) - \rho_{N-1}(\mathbf{r}) \end{aligned} \quad (15.3)$$

The f_+ function is expected to reflect the initial part of a nucleophilic reaction, and the f_- function an electrophilic reaction, i.e. the reaction will typically occur where the f_{\pm} function is large.⁸ For radical reactions the appropriate function is an average of f_+ and f_- .

$$f_0(\mathbf{r}) = \frac{1}{2}(f_+(\mathbf{r}) + f_-(\mathbf{r})) = \frac{1}{2}(\rho_{N+1}(\mathbf{r}) - \rho_{N-1}(\mathbf{r})) \quad (15.4)$$

The change in the electron density for each atomic site can be quantified by using the change in the atomic charges, although this of course suffers from the usual problems of defining atomic charges, as discussed in Chapter 9. The f_{\pm} functions may also be written in terms of orbital contributions.

$$\begin{aligned} f_+(\mathbf{r}) &= \phi_{\text{LUMO}}^2(\mathbf{r}) + \sum_{i=1}^{N_{\text{elec}}} \frac{\partial \phi_i^2(\mathbf{r})}{\partial n_i} \\ f_-(\mathbf{r}) &= \phi_{\text{HOMO}}^2(\mathbf{r}) + \sum_{i=1}^{N_{\text{elec}}-1} \frac{\partial \phi_i^2(\mathbf{r})}{\partial n_i} \end{aligned} \quad (15.5)$$

In the frozen MO approximation the last terms are zero and the Fukui functions are given directly by the contributions from the HOMO and LUMO. The preferred site of attack is therefore at the atom(s) with the largest MO coefficients in the HOMO/LUMO, in exact agreement with FMO theory. The Fukui function(s) may be considered as the equivalent (or generalization) of FMO methods within density functional theory (Chapter 6).

In the *Atoms In Molecules* approach (Section 9.3), the Laplacian ∇^2 (trace of the second-derivative matrix with respect to the coordinates) of the electron density measures the local increase or decrease of electrons. Specifically, if $\nabla^2\rho$ is negative, it marks an area where the electron density is locally concentrated, and therefore susceptible to attack by an electrophile. Similarly, if $\nabla^2\rho$ is positive, it marks an area where the electron density is locally depleted, and therefore susceptible to attack by a nucleophile. It has in general been found that a map of negative values of $\nabla^2\rho$ resembles the shape of the HOMO, and a map of positive values of $\nabla^2\rho$ resembles the shape of the LUMO.

The fact that features in the total electron density are closely related to the shapes of the HOMO and LUMO provides a much better rationale than the perturbation derivation as to why FMO theory works as well as it does. It should be noted, however, that improvements in the wave function do not lead to better performance of the FMO method. Indeed, the use of MOs from semi-empirical methods usually works better than data from *ab initio* wave functions. Furthermore, it should be kept in mind that only the HOMO orbital converges to a specific shape and energy as the basis set is improved in an *ab initio* calculation; the LUMO is normally determined by the most diffuse functions in the basis. The Fukui functions, on the other hand, can be calculated for any type of wave function.

Besides the already mentioned Fukui function, there are a couple of other commonly used concepts that can be connected with density functional theory (Chapter 6).⁹ The *electronic chemical potential* μ is given as the first derivative of the energy with respect to the number of electrons, which in a finite difference version is given as the average of the ionization potential (IP) and electron affinity (EA). Except for a difference in sign, this is also the Mulliken definition of *electronegativity* χ .¹⁰

$$-\mu = \chi = \frac{\partial E}{\partial N_{\text{elec}}} \approx \frac{1}{2}(\text{IP} + \text{EA}) \quad (15.6)$$

It should be noted that there are several other definitions of electronegativity, which do not necessarily agree on the ordering of the elements.¹¹

The second derivative of the energy with respect to the number of electrons is the *hardness* η (the inverse quantity η^{-1} is called the *softness*), which again may be approximated in terms of the ionization potential and electron affinity.

$$\eta = \frac{1}{2} \frac{\partial^2 E}{\partial N_{\text{elec}}^2} \approx \frac{1}{2}(\text{IP} - \text{EA}) \quad (15.7)$$

The electrophilicity, which measures the total ability to attract electrons, is defined as in eq. (15.8).¹²

$$\omega = \frac{\mu^2}{2\eta} \approx \frac{(\text{IP} + \text{EA})^2}{4(\text{IP} - \text{EA})} \quad (15.8)$$

A local version of the electrophilicity can be obtained by multiplying ω with the relevant Fukui function. These concepts play an important role in the *Hard and Soft Acid and Base* (HSAB) principle, which states that hard acids prefer to react with hard bases, and vice versa.¹³ By means of Koopmans' theorem (Section 3.4) the hardness is

related to the HOMO–LUMO energy difference. A “hard” molecule thus has a large HOMO–LUMO gap, and is expected to be chemically unreactive, i.e. hardness is related to chemical stability. A small HOMO–LUMO gap, on the other hand, indicates a “soft” molecule, and from second-order perturbation theory it also follows that a small gap between occupied and unoccupied orbitals will give a large contribution to the polarizability (Section 10.6.1), i.e. softness is a measure of how easily the electron density can be distorted by external fields, for example generated by another molecule. In terms of the perturbation equation (15.1), a hard–hard interaction is primarily charge controlled, while a soft–soft interaction is orbital controlled. Both FMO and HSAB theories may be considered as being limiting cases of chemical reactivity described by the Fukui function.⁸

15.3 Qualitative Molecular Orbital Theory

Frontier molecular orbital theory is closely related to various schemes of qualitative orbital theory where interactions between fragment MOs are considered.¹⁴ Ligand field theory, as commonly used in systems involving coordination to metal atoms, can be considered as a special case where only the d-orbitals on the metal and selected orbitals of the ligands are considered.

Two interacting orbitals will in general produce two new orbitals, having lower and higher energies than the non-interacting orbitals. The magnitude of the changes is determined by the orbital energy difference $\epsilon_a - \epsilon_b$ and the overlap S_{ab} . The overlap depends on the symmetries of the orbitals (orbitals of different symmetry have zero overlap), and the distance between them (the shorter the distance, the larger the overlap). The energies of the new orbitals can be calculated from the variational principle, and the qualitative result is shown in Figure 15.7.

$$\Delta \propto \frac{|S_{ab}|}{|\epsilon_a - \epsilon_b|} \quad (15.9)$$

There are two important features. The change in orbital energies is dependent on the magnitude of the overlap, $|S_{ab}|$, and inversely proportional to the energy difference of




Figure 15.7 Linear combination of two orbitals leads to two new orbitals with different energies

the original orbitals, $|\epsilon_a - \epsilon_b|$. Furthermore, the effect is largest for the highest energy orbital (antibonding combination), i.e. $\Delta_1 > \Delta_2$

If the two initial orbitals contain one, two or three electrons, the interaction will lead to a lower energy, with the stabilization being largest for the case of two electrons (e.g. a filled orbital interacting with an empty orbital). If both initial orbitals are fully occupied, the interaction will be destabilizing, i.e. a steric type repulsion. By adapting a set of HOMO and LUMO orbitals for atomic or molecular fragments, the favourable interactions may be identified based on overlap and energy considerations. Qualitative MO theory may thus be considered as an intramolecular form of FMO theory, with suitably chosen fragments.

Consider for example the two conformations for propene shown in Figure 15.8: which should be the more stable?




Figure 15.8 Possible propene conformations

By “chemical intuition”, the most important interaction is likely to be between the (filled) hydrogen s-orbitals and the (empty) π -orbital. The CH_3 group as a fragment has C_{3v} symmetry, and the three proper (symmetry adapted) linear combinations of the s-orbitals, together with the antibonding π -orbital, are given in Figure 15.9. The ϕ_1 orbital is lowest in energy, while the ϕ_2 and ϕ_3 orbitals are degenerate in perfect C_{3v} symmetry.

The ϕ_1 and ϕ_2 orbitals have a different symmetry than the π^* -orbital and can consequently not interact ($S = 0$). The interaction of the ϕ_3 orbital with the π^* system in the two conformations is shown in Figure 15.10.




Figure 15.9 Fragment orbitals for propene




Figure 15.10 Fragment orbital interaction

The overlap between the nearest carbon p-orbital and ϕ_3 is the largest contribution, but it is the same in the two conformations. The overlap with the distant carbon p-orbital is of opposite sign, and largest in conformation 2, since the distance is shorter. The total overlap between the ϕ_3 and π^* -orbitals is thus largest for conformation 1, which implies a larger stabilizing interaction, and it should consequently be lowest in energy. Indeed, conformation 2 is a transition structure for interconverting equivalent conformations corresponding to 1.

It is important to realize that whenever qualitative or frontier molecular orbital theory is invoked, the description is within the orbital (Hartree–Fock or density functional) model for the electronic wave function. In other words, rationalizing a trend in computational results by qualitative MO theory is only valid if the effect is present at the HF or DFT level. If the majority of the variation is due to electron correlation, an explanation in terms of interacting orbitals is not appropriate.

The interacting fragment orbital analysis can be put on more quantitative terms by performing explicit energy decomposition analysis of HF or DFT wave functions. The *extended transition state* (ETS) approach decomposes the energy change into four terms.¹⁵

$$\Delta E = \Delta E_{\text{prep}} + \Delta E_{\text{elstat}} + \Delta E_{\text{Pauli}} + \Delta E_{\text{orb}} \quad (15.10)$$

The energy change can for example be formation of a bond, and the analysis can be performed at various points along the reaction path. The preparation energy ΔE_{prep} describes the cost for perturbing the nuclear geometry from the optimum for the fragment to that of the species of interest. The electrostatic term ΔE_{elstat} describes the Coulomb interaction between the two fragment charge densities, while ΔE_{Pauli} describes the repulsion due to antisymmetrization and re-normalization of the two fragment wave functions when they are combined into one. Finally, ΔE_{orb} describes the stabilizing interaction due to mixing of occupied and unoccupied orbitals of the two fragments. The two central terms, ΔE_{elstat} and ΔE_{Pauli} , may loosely be associated with the steric repulsion.

An alternative decomposition, due to K. Morokuma, partitions the interaction energy into five terms.¹⁶

$$\Delta E = \Delta E_{\text{elstat}} + \Delta E_{\text{pol}} + \Delta E_{\text{CT}} + \Delta E_{\text{exchange}} + \Delta E_{\text{mix}} \quad (15.11)$$

The electrostatic term ΔE_{elstat} is the Coulomb interaction between the electron densities of the fragments, analogous to the corresponding ETS quantity. The polarization

term ΔE_{pol} describes the stabilization due to induced electric moments, while the charge transfer term ΔE_{CT} is a stabilization due to transfer of charge between the two fragments. The exchange term $\Delta E_{\text{exchange}}$ is analogous to the Pauli term in eq. (15.10), describing the repulsion due to the exchange energy arising from the antisymmetrization of the fragment wave functions. Finally, the mix term ΔE_{mix} contains the residual interaction not accounted for by the first four terms. The Morokuma energy decomposition is most useful for analyzing weak interactions; for strong interactions the mixing term often accounts for a significant part of the total interaction, thus obscuring the decomposition. The Morokuma energy decomposition has been combined with the NBO analysis (Section 9.6), which partly alleviates some of the instabilities in the original method.¹⁷

15.4 Woodward–Hoffmann Rules

The *Woodward–Hoffmann* (W–H) rules are qualitative statements regarding relative activation energies for two possible modes of reaction, which may have different stereochemical outcomes.¹⁸ For simple systems, the rules may be derived from a conservation of orbital symmetry, but they may also be generalized by an FMO treatment with conservation of bonding. Let us illustrate the Woodward–Hoffmann rules with a couple of examples, the preference of the 4 + 2 over the 2 + 2 product for the reaction of butadiene with ethylene, and the ring-closure of butadiene to cyclobutene.

A face-to-face reaction of two π -orbitals to form a cyclobutane involves the formation of two new C—C σ -bonds. The reaction may be imagined to occur under the preservation of symmetry, in this case C_{2v} , i.e. *concerted* (one-step, no intermediates) and *synchronous* (both bonds are formed at the same rate).




Figure 15.11 Reaction of two ethylenes to form cyclobutane under C_{2v} symmetry

Both the reactant and product orbitals may be classified according to their behaviour with respect to the two mirror planes present, being either *Symmetric* (no change of sign) or *Antisymmetric* (change of sign). The energetic ordering of the orbitals follows from a straightforward consideration of the bonding/antibonding properties. Since orbitals of different symmetries cannot mix, conservation of orbital symmetry establishes the correlation between the reactant and product sides.

The *orbital correlation diagram* shown in Figure 15.12 indicates that an initial electron configuration of $(\pi_1 + \pi_2)^2(\pi_1 - \pi_2)^2$ (ground state for the reactant) will end up as a doubly excited configuration $(\sigma_1 + \sigma_2)^2(\sigma_1^* + \sigma_2^*)^2$ for the cyclobutane product. This by itself indicates that the reaction should be substantially uphill in terms of energy. It may be put on a more sound theoretical footing by looking at the *state correlation diagram* in Figure 15.13.




Figure 15.12 Orbital correlation diagram for cyclobutane formation

The ground state wave function for the whole system (all four active and the remaining core and valence electrons) is symmetric with respect to both mirror planes, while the first excited state is antisymmetric. The intended correlation is indicated with dashed lines, the lowest energy configuration for the reactant correlates with a doubly excited configuration of the product, and vice versa. Since these configurations have the same symmetry (SS), an avoided crossing is introduced, leading to a significant barrier for the reaction. The presence of a reaction barrier due to symmetry conservation for the orbitals makes this a Woodward–Hoffmann *forbidden* reaction. The




Figure 15.13 State correlation diagram for cyclobutane formation

reaction for the excited state, however, does not encounter a barrier and is therefore denoted an *allowed* reaction.

The same conclusion may be reached directly from a consideration of the frontier orbitals. Formation of two new σ -bonds requires interaction of the HOMO of one fragment with the LUMO on the other. When the interaction is between orbital lobes on the same side (*Suprafacial*) of each fragment ($2s + 2s$), this leads to the picture shown in Figure 15.14.




Figure 15.14 $2s + 2s$ HOMO-LUMO interaction leading to two new σ -bonds

It is clearly seen that the HOMO-LUMO interaction leads to the formation of one bonding and one antibonding orbital, i.e. this is not a favourable interaction. The FMO approach also suggests that the $2 + 2$ reaction may be possible if it could occur with bond formation on opposite sides (*Antarafacial*) for one of the fragments.

Although the $2s + 2a$ reaction is Woodward-Hoffmann allowed, it is sterically so hindered that thermal $2 + 2$ reactions in general are not observed. Photochemical $2 + 2$ reactions, however, are well known.¹⁹

The $4s + 2s$ reaction of a diene with a double bond can in a concerted and synchronous reaction be envisioned to occur with the preservation of C_s symmetry.

The corresponding orbital correlation diagram is shown in Figure 15.17.

In this case the orbital correlation diagram shows that the lowest energy electron configuration in the reactant, $(\pi_1)^2(\pi_2)^2(\pi_3)^2$, correlates directly with the lowest energy




Figure 15.15 $2s + 2a$ HOMO–LUMO interaction leading to two new σ -bonds




Figure 15.16 Reaction of butadiene and ethylene to form cyclohexene under C_s symmetry

electron configuration in the product, $(\sigma_1)^2(\sigma_2)^2(\pi_1)^2$. This is also shown by the corresponding state correlation diagram, Figure 15.18.

In this case, there is no energetic barrier due to unfavourable orbital correlation, although other factors lead to an activation energy larger than zero. The direct correlation of ground state configurations for the reactant and product indicates a (relatively) easy reaction, and is an allowed reaction. The lowest excited state for the reactant, however, does not correlate with the lowest excited product state, and the photochemical reaction is consequently forbidden.

The FMO approach again indicates that the $4s + 2s$ interaction should lead directly to formation of two new bonding σ -bonds, i.e. this is an allowed reaction.

The preference for a concerted $4s + 2s$ reaction is experimentally supported by observations that show that the stereochemistry of the diene and dieneophile is carried over to the product, for example a *trans,trans*-1,4-disubstituted diene results in the two substituents ending up in a *cis* configuration in the cyclohexene product.²⁰

The ring-closure of a diene to a cyclobutene can occur with rotation of the two termini in the same (*Conrotatory*) or opposite (*Disrotatory*) directions. For suitably substituted compounds, these two reaction modes lead to products with different stereochemistry.

The disrotatory path has C_s symmetry during the whole reaction, while the conrotatory mode preserve C_2 symmetry. The orbital correlation diagrams for the two possible paths are shown as Figures 15.21 and 15.22.

It is seen that only the conrotatory path directly connects the reactant and product ground state configurations. Taking into account also the excited states leads to the state correlation diagram in Figure 15.23.

The conrotatory path is Woodward–Hoffmann allowed for a thermal reaction, while the corresponding photochemical reaction is predicted to occur in a disrotatory fashion.




Figure 15.17 Orbital correlation diagram for cyclohexene formation




Figure 15.18 State correlation diagram for cyclohexene formation




Figure 15.19 $4s + 2s$ HOMO-LUMO interaction leading to two new σ -bonds




Figure 15.20 Two possible modes of closing a diene to cyclobutene




Figure 15.21 Orbital correlation diagram for the disrotatory ring-closure of butadiene




Figure 15.22 Orbital correlation diagram for the conrotatory ring-closure of butadiene




Figure 15.23 State correlation diagram for the dis- and conrotatory ring-closure of butadiene

The same conclusion may again be reached by considering only the HOMO orbital. For the conrotatory path the orbital interaction leads directly to a bonding orbital, while the orbital phases for the disrotatory motion lead to an antibonding orbital.




Figure 15.24 HOMO orbital for the ring-closure of butadiene




Figure 15.25 FMO interactions for the [1,5]-hydrogen shift in 1,3-pentadiene

While the orbital and state diagrams can only be rigorously justified in the simple parent system where symmetry is present, the addition of substituents normally only alters the shape of the relevant orbitals slightly. The nodal structure of the orbitals is preserved for a large range of substituted systems, and the “preservation of bonding” displayed by the FMO type diagrams consequently have a substantially wider predictive range. It may be used for analyzing reactions where there is no symmetry element present under the whole reaction, as in for example the [1,5]-hydrogen shift in 1,3-pentadiene.

In the suprafacial migration the interaction of the pentadienyl radical singly occupied orbital with the hydrogen s-orbital is seen to involve breaking and making bonds where the orbital phases match. For the antarafacial path, however, the orbital in the product ends up being antibonding, i.e. a [1,5]-hydrogen migration is predicted to occur suprafacially, in agreement with experiments.²¹




Figure 15.26 FMO interactions for allowed modes of the [1,5]-methyl shift in 1,3-hexadiene

In the general case, the transferring group may migrate with either *retention* or *inversion* of its stereochemistry. A [1,5]-CH₃ migration, for example, is thermally allowed if it occurs suprafacial with retention of the CH₃ configuration, or if it occurs antarafacial with inversion of the methyl group.

The Woodward–Hoffmann allowed reactions can be classified according to how many electron are involved, and whether the reaction occurs thermally or photochemically, as shown in Table 15.1.

Table 15.1 Woodward–Hoffmann allowed reactions

Reaction type	Number of electrons	Thermally allowed	Photochemically allowed
Ring-closure	$4n$	Conrotatory	Disrotatory
	$4n + 2$	Disrotatory	Conrotatory
Cycloadditions	$4n$	Supra–antara <i>or</i> antara–supra	Supra–supra <i>or</i> antara–antara
	$4n + 2$	Supra–supra <i>or</i> antara–antara	Supra–antara <i>or</i> antara–supra
Migrations	$4n$	Antara–retention <i>or</i> supra–inversion	Supra–retention <i>or</i> antara–inversion
	$4n + 2$	Supra–retention <i>or</i> antara–inversion	Antara–retention <i>or</i> supra–inversion

The state correlation diagrams give an indication of the minimum theoretical level necessary for describing a reaction. For allowed reactions, the reactant configuration smoothly transforms into the product configuration by a continuous change of the orbitals, and they are consequently reasonably described by a single-determinant wave function along the whole reaction path. Forbidden reactions, on the other hand, necessarily involve at least two configurations since there is no continuous orbital transformation that connects the reactant and product ground states. Such reactions therefore require MCSCF type wave functions for a qualitative correct description.

While the state correlation diagram for the 2s + 2s reaction (Figure 15.13) indicates that the photochemical reaction should be allowed (and cyclobutanes are indeed observed as one of the products from such reactions), the implication that the product ends up in an excited state is not correct. Although the reaction starts out on the excited surface, it will at some point along the reaction path return to the lowest energy surface, and the product is formed in its ground state. The transition from the upper to the lower energy surface will normally occur at a geometry where the two surfaces “touch” each other, i.e. they have the same energy, and this is known as a *conical intersection*.²² Achieving the proper geometry for a transition between the two surfaces is often the dynamical bottleneck, and a conical intersection may be considered the equivalent of a TS for a photochemical reaction. As conical intersections involve two energy surfaces, MCSCF-based methods are required and non-adiabatic coupling elements (Section 3.1) are important. Locating a geometry corresponding to a conical intersection for a multi-dimensional system may be done using constrained optimization techniques (Section 12.5).

In some cases, the product of a pericyclic reaction will itself be subject to a further pericyclic rearrangement. Such cascade reactions are often synthetically useful, as they

may form complicated products in a single step. Depending on the exact system, two pericyclic reactions may occur with an intermediate, or occur in a single kinetic step, but with a very asynchronous bond breaking/formations. M. T. Reetz has coined the name *dyotropic* reaction for two sigmatropic shifts occurring in tandem,²³ while the term *bispericyclic* has been used in the more general case.²⁴

15.5 The Bell–Evans–Polanyi Principle/Hammond Postulate/Marcus Theory

The simpler the idea, the more names, could be the theme of this section.²⁵ The overriding idea is simple: for similar reactions, the more exothermic (endothermic) reaction will have the lower (higher) activation energy. This was formulated independently by *Bell*, *Evans* and *Polanyi* (BEP) in the 1930s, and is commonly known as the *BEP principle*.²⁶ The *Hammond postulate* relates the position of the transition structure to the exothermicity: for similar reactions, the more exothermic (endothermic) reaction will have the earlier (later) TS.²⁷ Compared with FMO theory, which tries to estimate relative activation energies from the *reactant* properties, the BEP principle tries to estimate relative activation energies from *product* properties (reaction energies). The above qualitative statements have been put on a more quantitative footing by the *Marcus equation*. This equation was originally derived for electron transfer reactions,²⁸ but it has since been shown that the same equation can be derived from a number of different assumptions, three of which will be illustrated below.

Let us assume a reaction coordinate x running from 0 (reactant) to 1 (product). The energy of the reactant as a function of x is taken as a simple parabola with a “force constant” of a . The energy of the product is also taken as a parabola with the same force constant, but offset by the reaction energy ΔE_0 . The position of the TS (x^\ddagger) is taken as the point where the two parabolas intersect, as shown in Figure 15.27.




Figure 15.27 Transition structure as the intersection of two parabolas

The TS position is calculated by equating the two energy expressions.

$$\begin{aligned}
 E_{\text{reactant}} &= a(x)^2 \\
 E_{\text{product}} &= a(x-1)^2 + \Delta E_0 \\
 a(x^\ddagger)^2 &= a(x^\ddagger-1)^2 + \Delta E_0 \\
 x^\ddagger &= \frac{1}{2} + \frac{\Delta E_0}{2a}
 \end{aligned}
 \tag{15.12}$$

For a thermoneutral reaction ($\Delta E_0 = 0$) the TS is exactly halfway between the reactant and product (as expected), while it becomes earlier and earlier as the reaction becomes more and more exothermic (ΔE_0 negative). The activation energy is given in eq. (15.13).

$$\begin{aligned}
 \Delta E^\ddagger &= E(x^\ddagger) = a\left(\frac{1}{2} + \frac{\Delta E_0}{2a}\right)^2 \\
 \Delta E^\ddagger &= \frac{a}{4} + \frac{\Delta E_0}{2} + \frac{\Delta E_0^2}{4a}
 \end{aligned}
 \tag{15.13}$$

Let us define the activation energy for a (possible hypothetical) thermoneutral reaction as the *intrinsic activation energy*, ΔE_0^\ddagger . As seen from eq. (15.13), $a = 4\Delta E_0^\ddagger$. The TS position and activation energy expressed in terms of ΔE_0^\ddagger are given in eq. (15.14).

$$\begin{aligned}
 x^\ddagger &= \frac{1}{2} + \frac{\Delta E_0}{8\Delta E_0^\ddagger} \\
 \Delta E^\ddagger &= \Delta E_0^\ddagger + \frac{\Delta E_0}{2} + \frac{\Delta E_0^2}{16\Delta E_0^\ddagger}
 \end{aligned}
 \tag{15.14}$$

The latter is, except for a couple of terms related to solvent reorganization, the Marcus equation. It should be noted that such curve-crossing models have been used in connection with VB methods to rationalize chemical reactivity and selectivity in a more general sense.²⁹

The central idea in the Marcus treatment is that the activation energy can be decomposed into a component characteristic of the reaction type, the intrinsic activation energy, and a correction due to the reaction energy being different from zero. Similar reactions should have similar intrinsic activation energies, and the Marcus equation obeys both the BEP principle and the Hammond postulate. Except for very exo- or endothermic reactions (or a very small ΔE_0^\ddagger), the last term in the Marcus equation is small, and it is seen that roughly half the reaction energy enters the activation energy. Note, however, that the activation energy is a parabolic function of the reaction energy. Thus for sufficiently exothermic reactions the equation predicts that the activation energy should *increase* as the reaction becomes *more* exothermic. The turnover occurs when $\Delta E_0 = -4\Delta E_0^\ddagger$. Much research has gone into proving such an “inverted” region, but experiments with very exothermic reactions are difficult to perform.³⁰

An alternative way of deriving the Marcus equation is again to assume a reaction coordinate running from 0 to 1. The intrinsic activation energy is taken as a parabola centred at $x = 1/2$. The reaction energy is taken as progressing linearly along the reaction coordinate. Adding these two contributions, and evaluating the position of the TS




Figure 15.28 Decomposition of a reaction barrier into a parabola and a linear term

and the activation energy in terms of ΔE_0 and ΔE_0^\ddagger , again leads to the Marcus equation.

Actually, the assumptions can be made even more general. The energy as a function of the reaction coordinate can always be decomposed into an “intrinsic” term, which is symmetric with respect to $x = 1/2$, and a “thermodynamic” contribution, which is anti-symmetric. Denoting these two energy functions h_2 and h_1 , it can be shown that the Marcus equation can be derived from the “square” condition, $h_2 = h_1^2$.³¹ The intrinsic and thermodynamic parts do not have to be parabolas and linear functions, as in Figure 15.28, they can be *any* type of functions. As long as the intrinsic part is the square of the thermodynamic part, the Marcus equation is recovered.

The idea can be taken one step further. The h_2 function can always be expanded in a power series of even powers of h_1 , i.e. $h_2 = c_2 h_1^2 + c_4 h_1^4 + \dots$. The exact values of the c -coefficients only influence the appearance of the last term in the resulting Marcus-like equation (eq. (15.14)). As already mentioned, this is usually a small correction anyway. For reactions where the reaction energy is less than or similar to the activation energy, there is thus a quite general theoretical background for the following statement: *For similar reactions, the difference in activation energy is roughly half the difference in reaction energy*. The trouble here is the word “similar”. How similar should reactions be in order for the intrinsic activation energy to be constant? And how do we calculate or estimate the intrinsic activation energy? We will return to the latter question shortly.

The Marcus equation provides a nice conceptual tool for understanding trends in reactivity.³² Consider for example the degenerate Cope rearrangement of 1,5-hexadiene and the ring-opening of Dewar benzene (bicyclo-[2,2,0]hexa-2,5-diene) to benzene.

The experimentally observed activation energies are 142 and 96 kJ/mol, respectively.³³ The Cope reaction is an example of a Woodward–Hoffmann allowed reaction ([3,3]-sigmatropic shift) while the ring-opening of Dewar benzene is a Woodward–Hoffmann forbidden reaction (the cyclobutene ring-opening must necessarily be disrotatory, otherwise the benzene product ends up with a *trans* double bond). Why does a forbidden reaction have a lower activation energy than an allowed reaction? This is readily explained by the Marcus equation. The Cope reaction is thermoneutral (reactant and product are identical) and the activation energy is purely




Figure 15.29 The Cope rearrangement and Dewar benzene ring-opening reaction

intrinsic, while the ring-opening is exothermic by 297 kJ/mol, and therefore has an intrinsic barrier of 218 kJ/mol. The “forbidden” reaction occurs only because it has a huge driving force in terms of a much more stable product, while the allowed reaction occurs even without a net energy gain.

The goal of understanding chemical reactivity is to be able to predict how the activation energy depends on properties of the reactant and product. Decomposing the activation energy into two terms, an intrinsic and a thermodynamic contribution, does not solve the problem. The reaction energy is relatively easy to obtain, from experiments, various theoretical methods or estimates based on additivity. But how does one estimate the intrinsic activation energy? It is purely a theoretical concept – the activation energy for a thermoneutral reaction. But most reactions are *not* thermoneutral, and there is no way of measuring such an intrinsic activation energy. For a series of “closely related” reactions it may be assumed to be constant, but the question then becomes: how closely related should reactions be?

Alternatively, it may be assumed that the intrinsic component can be taken as an average of the two corresponding *identity* reactions. Consider for example the S_N2 reaction of OH^- with CH_3Cl . The two identity reactions are $\text{OH}^- + \text{CH}_3\text{OH}$ and $\text{Cl}^- + \text{CH}_3\text{Cl}$. These two reactions are thermoneutral and their activation energies, which are purely intrinsic, can in principle be measured by isotopic substitution (for example $^{35}\text{Cl}^- + \text{CH}_3^{37}\text{Cl} \rightarrow \text{CH}_3^{35}\text{Cl} + ^{37}\text{Cl}^-$). From the reaction energy for the $\text{OH}^- + \text{CH}_3\text{Cl}$ reaction, and the assumption that the intrinsic barrier is the average of the two identity reactions, the activation energy can be calculated. An example of the accuracy of this procedure for the series of S_N2 reactions $\text{OH}^- + \text{CH}_3\text{X}$ is given in Table 15.2.

Table 15.2 Comparing experimental activation barriers (kJ/mol) to those calculated by the Marcus equation for the reaction $\text{OH}^- + \text{CH}_3\text{X}$

	ΔG^\ddagger (identity)	ΔG_0 (exp.)	ΔG^\ddagger (exp.)	ΔG^\ddagger (Marcus)
OH^-	170			
F^-	133	−94	109	108
Cl^-	111	−92	103	98
Br^-	99	−98	95	90
I^-	97	−89	97	93

Again this averaging procedure can only be expected to work when the reactions are sufficiently “similar”, which is difficult to quantify *a priori*. The Marcus equation is

therefore more a conceptual tool for explaining trends than for deriving quantitative result.

15.6 More O'Ferrall–Jencks Diagrams

The BEP/Hammond/Marcus treatment only considers changes due to energy differences between the reactant and product, i.e. changes in the TS position *along* the reaction coordinate. It is often useful also to include changes that may occur in a direction *perpendicular* to the reaction coordinate. Such two-dimensional diagrams are associated with the names of *More O'Ferrall* and *Jencks* (MOJ diagrams).³⁴

Consider for example the Cope rearrangement of 1,5-hexadiene. Since the reaction is degenerate the TS will have D_{2h} symmetry (the lowest energy TS has a conformation resembling a chair-like cyclohexane). It is, however, not clear how strong the forming and breaking C—C bonds are at the TS. If they both are essentially full C—C bonds, the reaction may be described as bond formation followed by bond breaking. The TS therefore has the character of being a 1,4-biradical, as illustrated by path **B** in Figure 15.30. Alternatively, the C—C bonds may be very weak at the TS, corresponding to a situation where bond breaking occurs before bond formation, and the TS can be described as two weakly interacting allyl radicals (path **C**). The intermediate situation, where both bonds are roughly half formed/broken can be described as having a delocalized structure similar to benzene, i.e. an “aromatic” type TS (path **A**).




Figure 15.30 MOJ diagram for the Cope rearrangement of 1,5-hexadiene

In such MOJ diagrams the x - and y -coordinates are normally taken to be bond orders (Section 9.1) or $(1 - \text{bond order})$ for the breaking and forming bonds, such that the coordinates run from 0 to 1. A third axis corresponding to the energy is implied, but rarely drawn.

At the TS, the energy along the reaction path is a maximum, while it is a minimum in the perpendicular direction(s). A one-dimensional cut through the (0,0) and (1,1)

corners for path **A** in Figure 15.30 thus corresponds to Figure 15.28. A similar cut through the (0,1) and (1,0) corners will display a normal (as opposed to inverted) parabolic behaviour, with the TS being at the minimum on the curve. The whole energy surface corresponding to Figure 15.29 will have the qualitative appearance as shown in Figure 15.31.




Figure 15.31 MOJ diagram corresponding to Figure 15.30 with the energy as the vertical axis

There is good evidence that the Cope reaction in the parent 1,5-hexadiene has an “aromatic” type TS, corresponding to path **A** in Figure 15.30, i.e. a “central” or “diagonal” reaction path. The importance of MOJ diagrams is that they allow a qualitative prediction of changes in the TS structure for a series of similar reactions. The addition of substituents that stabilize the product relative to the reactant corresponds to a lowering of the (1,1) corner, thereby moving the TS closer to the (0,0) corner, i.e. towards the reactant. The one-dimensional BEP/Hammond/Marcus treatment thus corresponds to changes along the (0,0)–(1,1) diagonal.

Substituents that do not change the overall reaction energy may still have an influence on the TS geometry. Consider for example 2,5-diphenyl-1,5-hexadiene. The reaction is still thermoneutral but the phenyl groups will preferentially stabilize the 1,4-biradical structure, i.e. lower the energy of the (1,0) corner. From Figure 15.31 it is clear that this will lead to a TS that is shifted towards this corner, i.e. moving the reaction from path **A** towards **B** in Figure 15.30. Similarly, substituents that preferentially stabilize the bis-allyl radical structure (such as 1,4-diphenyl-1,5-hexadiene) will perturb the reaction towards path **C**, since the (0,1) corner is lowered in energy relative to the other corners.

From such MOJ diagrams it can be inferred that changes in the system that alter the relative energy *along* the reaction diagonal (lower-left to upper-right) imply changes in the TS in the *opposite* direction. Changes that alter the relative energy *perpendicular* to the reaction diagonal (upper-left to lower-right) imply changes in the TS in the *same* direction as the perturbation.

The structures in the (1,0) and (0,1) corners are not necessarily stable species; they may correspond to hypothetical structures. In the Cope rearrangement, it appears that the reaction only involves a single TS, independent of the number and nature of substituents. The reaction path may change from $\mathbf{B} \rightarrow \mathbf{A} \rightarrow \mathbf{C}$ depending on the system, but there are no intermediates along the reaction coordinate.

In other cases, one or both of the perpendicular corners may correspond to a minimum on the potential energy surface, and the reaction mechanism can change from being a one-step reaction to two-step. An example of this would be elimination reactions. The x -axis in this case corresponds to the breaking bond between carbon and hydrogen, while the y -axis is the breaking bond between the other carbon and the leaving group.




Figure 15.32 MOJ diagram for elimination reactions

An E2 type reaction has simultaneous breaking of the C—H and C—L bonds while forming the B—H bond, and corresponds to the diagonal path **A** in Figure 15.32. Path **C** involves initial loss of the leaving group to form a carbocation (upper-left corner), followed by loss of H^+ (which is picked up by the base), i.e. this corresponds to an E1 type mechanism involving two TS's and an intermediate. Path **B**, on the other hand, involves formation of a carbanion, followed by elimination of the leaving group in a second step, i.e. an E1_{cb} mechanism. Substituents that stabilize the carbocation thus shift the reaction from an E2 to an E1 type mechanism, while anionic stabilizing substituents will shift the reaction towards an E1_{cb} path.

In principle MOJ diagrams can be extended to more dimensions, for example by also including the B—H bond order in the above elimination reaction, but this is rarely done, not least because of the problems of illustrating more than two dimensions.

References

1. I. Fleming, *Frontier Orbitals and Organic Chemical Reactions*, Wiley, 1976.
2. A. Devaquet, *Mol. Phys.*, **18** (1970), 233.

3. A. Alexakis, C. Chuit, M. Commercon-Bourgain, J. P. Foulon, N. Jabri, P. Mangeney, J. F. Normant, *Pure Appl. Chem.*, **56** (1984), 91.
4. G. Marino, *Adv. Heterocycl. Chem.*, **13** (1971), 235.
5. J. Sauer, *Angew. Chem. Int. Ed.*, **6** (1967), 16.
6. P. Geerlings, F. De Proft, W. Langenaeker, *Chem. Rev.*, **103** (2003), 1793.
7. R. G. Parr, W. Yang, *J. Am. Chem. Soc.*, **106** (1984), 4049.
8. Y. Li, J. N. S. Evans, *J. Am. Chem. Soc.*, **117** (1995), 7756; CR 103.
9. F. D. Proft, S. Liu, R. G. Parr, *J. Chem. Phys.*, **107** (1997), 3000.
10. R. S. Mulliken, *J. Chem. Phys.*, **2** (1934), 782.
11. K. T. Gijo, F. D. Proft, P. Geerlings, *J. Phys. Chem. A*, **109** (2005) 2925.
12. R. G. Parr, L. v. Szentpály, S. Liu, *J. Am. Chem. Soc.*, **121** (1999) 1922.
13. R. G. Pearson, *J. Am. Chem. Soc.*, **85** (1963), 3533.
14. A. Rauk, *Orbital Interaction Theory of Organic Chemistry*, Wiley, 1994; T. A. Albright, J. K. Burdett, M. -H. Whangbo, *Orbital Interactions in Chemistry*, Wiley, 1985.
15. T. Ziegler, A. Rauk, *Theor. Chim. Acta*, **46** (1977), 1; F. M. Bickelhaupt, E. J. Baerends, *Rev. Comp. Chem.*, **15** (2000), 1; F. M. Bickelhaupt, *J. Comp. Chem.*, **25** (1999), 114.
16. K. Morokuma, *Acc. Chem. Res.*, **10** (1977), 94.
17. E. D. Glendening, A. Streitwieser, *J. Chem. Phys.*, **100** (1994), 2900.
18. R. B. Woodward, R. Hoffmann, *The Conservation of Orbital Symmetry*, Academic Press, 1970; R. B. Woodward, R. Hoffmann, *Angew. Chem. Int. Ed.*, **8** (1969), 781.
19. N. Turro, *Modern Molecular Photochemistry*, The Benjamin/Cummings Publishing Co., 1978.
20. J. G. Martin, R. K. Hill, *Chem. Rev.*, **61** (1961), 537.
21. W. R. Roth, J. König, K. Stein, *Chem. Ber.*, **103** (1970), 426.
22. F. Bernardi, M. Olivucci, M. A. Robb, *Chem. Soc. Rev.*, **25** (1996), 321.
23. M. T. Reetz, *Adv. Organomet. Chem.*, **16** (1977), 33.
24. J. Limanto, K. S. Khuong, K. N. Houk, M. L. Snapper, *J. Am. Chem. Soc.*, **125** (2003), 16310; D. H. Nouri, D. J. Tantillo, *J. Org. Chem.*, **71** (2006), 3686.
25. W. P. Jencks, *Chem. Rev.*, **85** (1985), 511.
26. R. P. Bell, *Proc. R. Soc. London, Ser. A*, **154** (1936), 414; M. G. Evans, M. Polanyi, *J. Chem. Soc., Faraday Trans.*, **32** (1936), 1340.
27. G. S. Hammond, *J. Am. Chem. Soc.*, **77** (1955), 334.
28. R. A. Marcus, *J. Phys. Chem.*, **72** (1968), 891.
29. S. Shaik, P. C. Hiberty, *Rev. Comp. Chem.*, **20** (2004), 1; A. Pross, *Theoretical and Physical Principles of Organic Reactivity*, Wiley, 1995.
30. D. M. Guldi, K. -D. Asmus, *J. Am. Chem. Soc.*, **119** (1997), 5744.
31. J. Donnelly, J. R. Murdoch, *J. Am. Chem. Soc.*, **106** (1984), 4724.
32. See for example V. Aviyente, H. Y. Yoo, K. N. Houk, *J. Org. Chem.*, **62** (1997), 6121.
33. W. von Doering, V. G. Toscano, G. H. Beasley, *Tetrahedron*, **27** (1971), 5299; M. J. Goldstein, R. S. Leight, *J. Am. Chem. Soc.*, **99** (1977), 8112.
34. R. A. More O'Ferrall, *J. Chem. Soc. B*, (1970), 274; W. P. Jencks, *Chem. Rev.*, **72** (1972), 705; S. S. Shaik, H. B. Schlegel, S. Wolfe, *Theoretical Aspects of Physical Organic Chemistry. The S_N2 Mechanism*, Wiley, 1992.

16 Mathematical Methods

Computational chemistry relies on computers to solve the complicated mathematical equations describing the physics behind the model. The language for deriving and describing these models is mathematics, and this chapter summarizes some of the commonly used mathematical concepts and techniques used in computational chemistry.

16.1 Numbers, Vectors, Matrices and Tensors

Some physical quantities, such as the total molecular mass or charge, can be specified by a single number, referring to the magnitude of the quantity in a given set of units. The mathematical term for such a number is a *scalar*. Other quantities require a *set* of numbers, as for example three scalars for specifying the position of a particle in a coordinate system. A coordinate system is defined by the *origin* (“zero point”), the *directions* of the coordinate axes, and the *units* along the axes. Two common examples are Cartesian $\{x, y, z\}$ and spherical polar $\{r, \theta, \varphi\}$ systems. The same point in space can be specified either by the Cartesian x, y, z coordinates $\{0.500, 0.866, 1.000\}$ or by the spherical polar r, θ, φ coordinates $\{2, 30, 60\}$ with angles measured in degrees.




Figure 16.1 Cartesian and spherical polar coordinate systems

The direction from the origin to the point specified by the three coordinates represents a *vector*, having a length and a direction. Another example of a 3-vector is the velocity of a particle $\{v_x, v_y, v_z\}$, or alternatively $\{\partial x/\partial t, \partial y/\partial t, \partial z/\partial t\}$. For a system with N particles, the positions or velocities of all particles can be specified by a vector of length $3N$, i.e. $\{x_1, y_1, z_1, x_2, y_2, \dots, y_N, z_N\}$ or $\{v_{x1}, v_{y1}, v_{z1}, v_{x2}, v_{y2}, \dots, v_{yN}, v_{zN}\}$.

$$\{v_{x1} \quad v_{y1} \quad v_{z1} \quad v_{x2} \quad \dots\} = \left\{ \frac{\partial x_1}{\partial t} \quad \frac{\partial y_1}{\partial t} \quad \frac{\partial z_1}{\partial t} \quad \frac{\partial x_2}{\partial t} \quad \dots \right\} \quad (16.1)$$

The notation for such vectors is often generalized to simply $\mathbf{x} = \{x_1, x_2, x_3, \dots, x_{N-1}, x_N\}$, where N now refers to the total number of elements, i.e. equal to $3N$ in the above notation.

A *complex number* z can be interpreted as a 2-vector in an xy -coordinate system, $z = x + iy$, where i is the symbol for $\sqrt{-1}$ and x and y are real numbers. Here x and y are referred to as the real and imaginary parts of z . Alternatively, the complex number can be associated with polar coordinates, i.e. the distance r from the origin and the angle θ relative to the real axis.




Figure 16.2 Imaginary numbers interpreted as a point in a two-dimensional coordinate system

The *complex conjugate* of a complex number z is denoted by z^* and is obtained by changing the sign of the imaginary part, i.e. $z^* = x - iy$ or equivalently $z^* = re^{-i\theta}$.

The concept of complex numbers can be generalized to hypercomplex numbers, with the next level being a 4-vector, called a *quaternion*, i.e. $q = q_0 + iq_1 + jq_2 + kq_3$, with q_0, q_1, q_2, q_3 being real numbers. A quaternion has a real part, q_0 , and the three imaginary components q_1, q_2, q_3 . The latter can be considered as a vector in a three-dimensional space, where each of the unit vectors have the property $i^2 = j^2 = k^2 = -1$. As one moves up in dimensions in this generalization, common mathematical laws gradually get lost. Quaternions, for example, do not obey the commutative law ($q_a q_b \neq q_b q_a$), while octonions (8-vectors) in addition do not obey the associative law ($(q_a q_b) q_c \neq q_a (q_b q_c)$). Quaternions are encountered for example in relativistic (4-component) quantum mechanics, and they also form a more natural basis for parameterizing the rotation of a three-dimensional structure, rather than the traditional three Euler angles.¹ The latter involve trigonometric functions that are both computationally expensive to evaluate and display singularities. Furthermore, the quaternion formulation treats all the coordinate axes as equivalent, while the Euler parameterization makes the z -axis a special direction.

Vectors can also arise from mathematical operations on functions of coordinates, as for example the gradient being the first derivative of an energy function.

$$\{g_1 \quad g_2 \quad g_3 \quad g_4 \quad \dots\} = \left\{ \frac{\partial E}{\partial x_1} \quad \frac{\partial E}{\partial x_2} \quad \frac{\partial E}{\partial x_3} \quad \frac{\partial E}{\partial x_4} \quad \dots \right\} \quad (16.2)$$

In such cases, it is implicit that the first gradient element is the derivative with respect to the first variable, etc.

The second derivative of the energy is an ordered two-dimensional set of numbers, called a *matrix*.

$$\begin{pmatrix} H_{11} & H_{12} & \cdots \\ H_{21} & H_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} = \begin{pmatrix} \frac{\partial^2 E}{\partial x_1^2} & \frac{\partial^2 E}{\partial x_1 \partial x_2} & \cdots \\ \frac{\partial^2 E}{\partial x_2 \partial x_1} & \frac{\partial^2 E}{\partial x_2^2} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \quad (16.3)$$

The third derivative of the energy is an ordered three-dimensional set of numbers, called a *tensor*, which can be arranged in a cube. Corresponding higher order derivatives can be thought of as ordered sets of numbers in “hypercubes”, called N -order tensors for a hypercube of dimension N .

Ordered sets of numbers may collectively be called tensors, with a matrix being a second-order tensor, and a vector a first-order tensor. Since tensors of order higher than two are relatively rare, the terms vector and matrix are more commonly used. Sometimes it is also convenient to consider a vector as a $1 \times N$ or $N \times 1$ matrix, and a scalar as a 1×1 matrix.

The conversion between a $1 \times N$ and an $N \times 1$ vector, or from an $M \times N$ matrix to an $N \times M$ matrix, is done by *transposition*, indicated by a superscript t . Transposition simply interchanges the i th element with the j th element.

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^t = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$$

$$\begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix}^t = \begin{pmatrix} a & d \\ b & e \\ c & f \end{pmatrix} \quad (16.4)$$

If the matrix elements are complex, the *adjoint* matrix is defined by complex conjugation of the elements followed by transposition, and is denoted with a superscript \dagger . *Hermitian* matrices are very common in quantum chemistry, and are defined as being self-adjoint, i.e. $\mathbf{A} = \mathbf{A}^\dagger$. If all the matrix elements are real, the matrix is called *symmetric*, i.e. $\mathbf{A} = \mathbf{A}^t$.

The addition and subtraction of matrices, which now encompass vectors as well, is directly the addition and subtraction of the elements, analogous to the rules for scalars.

$$\begin{pmatrix} a_1 \\ a_2 \\ \vdots \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ \vdots \end{pmatrix} = \begin{pmatrix} a_1 + b_1 \\ a_2 + b_2 \\ \vdots \end{pmatrix}$$

$$\begin{pmatrix} a_{11} & a_{12} & \cdots \\ a_{21} & a_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} + \begin{pmatrix} a_{11} & a_{12} & \cdots \\ a_{21} & a_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} = \begin{pmatrix} a_{11} + b_{11} & a_{12} + b_{12} & \cdots \\ a_{21} + b_{21} & a_{22} + b_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \quad (16.5)$$

The multiplication of matrices, however, is somewhat different. In standard matrix multiplications, the ij th element in the product is formed by multiplying the elements of

the i th row with the elements of the j th column, and adding all the terms. For the multiplication of two 2×2 matrices the result is given in eq. (16.6).

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{pmatrix} \quad (16.6)$$

Note that this means that matrix multiplication is not necessarily commutative.

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \neq \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \quad (16.7)$$

This is perhaps most easily seen by multiplying two rectangular matrices. The result of multiplying a 2×3 matrix with a 3×2 matrix is a 2×2 matrix.

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{pmatrix} = \begin{pmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{pmatrix} \quad (16.8)$$

While the result of multiplying a 3×2 matrix with 2×3 matrix is a 3×3 matrix.

$$\begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{pmatrix} \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} = \begin{pmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \\ c_{31} & c_{32} & c_{33} \end{pmatrix} \quad (16.9)$$

Even for square matrices, however, the matrix product \mathbf{AB} is not necessarily equal to \mathbf{BA} .

In some cases the matrix elements are multiplied together element by element, which is called an *entrywise* product, and denoted with a “dot” between the two matrices.

$$\begin{pmatrix} a_{11} & a_{12} & \cdots \\ a_{21} & a_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \cdot \begin{pmatrix} b_{11} & b_{12} & \cdots \\ b_{21} & b_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} & a_{12}b_{12} & \cdots \\ a_{21}b_{21} & a_{22}b_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \quad (16.10)$$

For vectors, which can be considered $1 \times N$ or $N \times 1$ matrices, the result of multiplying a $1 \times N$ matrix with an $N \times 1$ matrix is a 1×1 matrix, or a scalar. This is called an *inner* or *dot* product.

$$\mathbf{a}^t \mathbf{b} = (a_1 \quad a_2) \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = (a_1 b_1 + a_2 b_2) \quad (16.11)$$

The length (or *norm*) of a vector follows directly from the interpretation of a vector as a directional line from the origin to a point in space, and is defined as the square root of the dot product of the vector with itself. If the vector components are complex numbers the transposition is replaced by the adjoint instead.

$$|\mathbf{a}| = \sqrt{\mathbf{a}^t \mathbf{a}} = \sqrt{a_1^2 + a_2^2 + a_3^2 + \dots} \quad (16.12)$$

The “physical” interpretation of a dot product of two vectors is related to the angle between them, specifically for two vectors of unit length, the dot product is the cosine of the angle.

$$\mathbf{a}^t \mathbf{b} = |\mathbf{a}| |\mathbf{b}| \cos \alpha \quad (16.13)$$

A dot product of +1 means that the two (unit) vectors are aligned, a value of -1 means that they are aligned, but pointing in opposite directions, while a dot product of 0 means that the two vectors are orthogonal.

The opposite of a dot product is multiplication of an $N \times 1$ matrix with a $1 \times N$ matrix to give an $N \times N$ matrix, and is called an *outer* product.

$$\mathbf{a} \mathbf{b}^t = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \begin{pmatrix} b_1 & b_2 \end{pmatrix} = \begin{pmatrix} a_1 b_1 & a_1 b_2 \\ a_2 b_1 & a_2 b_2 \end{pmatrix} \quad (16.14)$$

The inner and outer products of two vectors produce a scalar and a matrix, respectively. Two 3-vectors may also be multiplied together to generate a new 3-vector, a procedure called a *vector* or *cross* product, with the result given by eq. (16.15).

$$\begin{Bmatrix} x_1 \\ y_1 \\ z_1 \end{Bmatrix} \times \begin{Bmatrix} x_2 \\ y_2 \\ z_2 \end{Bmatrix} = \begin{Bmatrix} y_1 z_2 - z_1 y_2 \\ z_1 x_2 - x_1 z_2 \\ x_1 y_2 - y_1 x_2 \end{Bmatrix} \quad (16.15)$$

The vector product gives a vector perpendicular to both of the original vectors with a length of $|\mathbf{a}| |\mathbf{b}| \sin \alpha$ (compare with eq. (16.13)), and is therefore zero if the two original vectors are aligned. It follows trivially that $\mathbf{a} \times \mathbf{a} = 0$ for any vector \mathbf{a} , and that $\mathbf{a} \times \mathbf{b} = -\mathbf{b} \times \mathbf{a}$.

A matrix *determinant* is denoted $|\mathbf{A}|$ and is given explicitly for the 2×2 and 3×3 cases in eqs (16.16) and (16.17).

$$|\mathbf{A}| = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} = a_{11} a_{22} - a_{12} a_{21} \quad (16.16)$$

$$|\mathbf{A}| = \begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} = \begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} = a_{11} a_{22} a_{33} - a_{11} a_{23} a_{32} - a_{12} a_{21} a_{33} + a_{12} a_{23} a_{31} + a_{13} a_{21} a_{32} - a_{13} a_{22} a_{31} \quad (16.17)$$

The determinant of larger matrices is similarly given as a sum of $N!$ terms, each being the product of N elements. A convenient procedure for the evaluation consists of decomposing the determinant according to a row (or column), with each element being multiplied with a sub-determinant formed by removing the elements of the corresponding rows and columns, and a factor $(-1)^{(i+j)}$.

$$\begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} = a_{11} \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix} - a_{12} \begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix} + a_{13} \begin{vmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{vmatrix} \quad (16.18)$$

This procedure can be applied recursively until only 1×1 determinants remain.

Only square matrices have determinants, and those with determinants equal to $e^{i\theta}$ (Figure 16.2 with $r = 1$) are called *unitary*. A unitary matrix, where all the elements are real (i.e. not complex), is called *orthogonal* and obviously has a determinant equal to +1 or -1. We will in general use the unitary notation, although in most cases the matrices are actually orthogonal.

Determinants have a number of important properties:

- (1) Interchanging two rows or columns in a matrix changes the sign of the determinant. This property is used for parameterizing wave functions in terms of Slater determinants, as the wave function antisymmetry is thereby automatically fulfilled (Section 3.2).
- (2) Adding a row (or a fraction thereof) to another row leaves the determinant unchanged, and similarly for columns. This allows for example representation of a wave function either in terms of canonical or localized molecular orbitals (Section 9.4).
- (3) If two rows or columns are identical except for a multiplicative constant, the determinant is zero. This is easily seen, since one of these rows/columns can be made into a zero-vector by subtraction of the two, and expansion according to this zero row/column by eq. (16.18) will give zero. Such matrices may arise owing to *linear dependencies* of the rows or columns.

Division by matrices is done formally by multiplying with the *inverse* of a matrix, where the inverse is defined such that multiplication of a matrix with its inverse produces a unit matrix.

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{A}\mathbf{A}^{-1} = \mathbf{I}$$

$$\mathbf{I} = \begin{pmatrix} 1 & 0 & 0 & \cdots \\ 0 & 1 & 0 & \cdots \\ 0 & 0 & 1 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (16.19)$$

The elements of a matrix inverse are given by the elements of the matrix itself, and the inverse of the matrix determinant. Specifically, the ij th element in the \mathbf{A}^{-1} matrix is given as the inverse of $|\mathbf{A}|$ times the determinant of the sub-matrix corresponding to removing the j th row and i th column, and a factor $(-1)^{(i+j)}$. Note the transposition, the ij th element in the inverse matrix is formed from the sub-matrix corresponding to the ji th element in the original matrix. For a 3×3 matrix, this is exemplified by the b_{ij} elements in eq. (16.20).

$$\mathbf{A} = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \mathbf{A}^{-1} = |\mathbf{A}|^{-1} \begin{pmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{pmatrix} \quad (16.20)$$

$$b_{11} = \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix} \quad b_{12} = -\begin{vmatrix} a_{12} & a_{13} \\ a_{32} & a_{33} \end{vmatrix} \quad b_{13} = \begin{vmatrix} a_{12} & a_{13} \\ a_{22} & a_{23} \end{vmatrix}$$

It thus follows that only square matrices with determinants different from zero have an inverse matrix. Rectangular matrices can be defined to have a *generalized* inverse

matrix. Such generalized inverse matrices, and more complicated matrix algebra, such as calculating functions of matrices, are considered in the next section. It should be noted that for orthogonal matrices the inverse is simply the transposed matrix, $\mathbf{A}^{-1} = \mathbf{A}^t$, for unitary matrices the transposition must be accompanied by a complex conjugation also, i.e. $\mathbf{A}^{-1} = \mathbf{A}^\dagger$. The names and properties of some special matrices are shown in Table 16.1.

Table 16.1 Some special matrices and their names

Name	Properties
Complex conjugate	\mathbf{A}^* , complex conjugate all elements
Transposed	\mathbf{A}^t , interchange elements ij and ji
Adjoint	\mathbf{A}^\dagger , interchange and complex conjugate elements ij and ji
Symmetric	$\mathbf{A}^t = \mathbf{A}$
Antisymmetric	$\mathbf{A}^t = -\mathbf{A}$
Hermitian	$\mathbf{A}^\dagger = \mathbf{A}$
Anti-Hermitian	$\mathbf{A}^\dagger = -\mathbf{A}$
Inverse	$\mathbf{A}^{-1}\mathbf{A} = \mathbf{A}\mathbf{A}^{-1} = \mathbf{I}$
Orthogonal	$ \mathbf{A} = \pm 1$, real elements, $\mathbf{A}^{-1} = \mathbf{A}^t$
Unitary	$ \mathbf{A} = e^{i\theta}$, complex elements, $\mathbf{A}^{-1} = \mathbf{A}^\dagger$

Matrices arise for example in solving a system of linear equations, where the formal solution can be obtained by multiplication with \mathbf{A}^{-1} on both sides.

$$\begin{aligned}
 a_{11}x_1 + a_{12}x_2 + \dots + a_{1n}x_n &= b_1 \\
 a_{21}x_1 + a_{22}x_2 + \dots + a_{2n}x_n &= b_2 \\
 \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots &= \vdots \\
 a_{n1}x_1 + a_{n2}x_2 + \dots + a_{nn}x_n &= b_n
 \end{aligned} \tag{16.21}$$

$$\mathbf{Ax} = \mathbf{b}$$

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$$

It thus follows that a solution only exists if \mathbf{A}^{-1} exists, i.e. if $|\mathbf{A}|$ is non-zero. In actual calculations, it is rare that matrix determinants are exactly zero. If $|\mathbf{A}|$ is very small, the solution vector \mathbf{x} becomes sensitive to small details in the original \mathbf{A} matrix. Such systems are called *ill-conditioned*, and should be treated by *single value decomposition*, as described in the next section.

For the special case of the right-hand side (\mathbf{b} -vector) in eq. (16.21) being zero, only the trivial $\mathbf{x} = \mathbf{0}$ solution exists *if* \mathbf{A}^{-1} exists. A non-trivial solution is therefore only possible if \mathbf{A}^{-1} does *not* exist, which is equivalent to the condition that $|\mathbf{A}|$ is zero. Linear dependence in the \mathbf{A} matrix is thus a *condition* for a non-trivial solution, and the resulting \mathbf{x} -vector is obtained as a parametric solution of one or more variables. These parameters can be fixed for example by requiring that the \mathbf{x} -vector(s) are normalized and mutually orthogonal.

16.2 Change of Coordinate System

In many cases, it is possible to simplify a problem by choosing a particular coordinate system. It is therefore important to be able to describe how vectors and matrices change when switching from one coordinate system to another.

Some coordinate transformations are *non-linear*, such as converting from a Cartesian to a spherical polar system. Here the r, θ, ϕ coordinates are related to the x, y, z coordinates by square root and trigonometric functions, as shown in Figure 16.1. Other coordinate transformations are *linear*, with the new coordinates given as linear combinations of the old ones. A linear transformation can be described as a *rotation* of the coordinate system.




Figure 16.3 Rotation of a coordinate system

For the 2×2 case, the new coordinates x' and y' are related to the original x and y coordinates by means of a 2×2 matrix containing cosines and sines of the rotational angle α .

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \quad (16.22)$$

The rotation matrix is a unitary (orthogonal) matrix \mathbf{U} , since the determinant is equal to 1 ($\cos^2 \alpha + \sin^2 \alpha = 1$). The significance of a unitary matrix is that it describes a rotation of the coordinate system without changing the length of the coordinate axes. A unitary matrix with a determinant of -1 describes a rotation of the coordinate system, followed by inverting the directions of the coordinate axis, i.e. an improper rotation in the language of point group symmetry.

The connection between the primed and unprimed coordinate systems is given by the unitary matrix in eq. (16.22), and can be written as in eq. (16.23).

$$\mathbf{x}' = \mathbf{U}\mathbf{x} \quad (16.23)$$

The inverse operation \mathbf{U}^{-1} corresponds to rotation with the angle $-\alpha$, and back-transforms the primed coordinates to the unprimed ones.

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} x' \\ y' \end{pmatrix} \quad (16.24)$$

$$\mathbf{x} = \mathbf{U}^{-1}\mathbf{x}'$$

It is easily verified that the matrix product $\mathbf{U}^{-1}\mathbf{U}$ gives a 2×2 unit matrix.

Consider now a (multi-dimensional) linear function \mathbf{f} defined by the action of a matrix \mathbf{A} on a vector \mathbf{x} .

$$\mathbf{f} = \mathbf{A}\mathbf{x} \quad (16.25)$$

In the rotated coordinate system the corresponding connection is given by eq. (16.26).

$$\mathbf{f}' = \mathbf{A}'\mathbf{x}' \quad (16.26)$$

By using the transformations (16.23) between the two coordinate systems, and the fact that a unit matrix of the form $\mathbf{U}^{-1}\mathbf{U}$ can be freely inserted, we get eq. (16.27).

$$\begin{aligned} \mathbf{f} &= \mathbf{A}\mathbf{x} \\ \mathbf{f} &= \mathbf{A}(\mathbf{U}^{-1}\mathbf{U})\mathbf{x} \\ \mathbf{U}\mathbf{f} &= \mathbf{U}\mathbf{A}(\mathbf{U}^{-1}\mathbf{U})\mathbf{x} \\ (\mathbf{U}\mathbf{f}) &= (\mathbf{U}\mathbf{A}\mathbf{U}^{-1})(\mathbf{U}\mathbf{x}) \\ \mathbf{f}' &= (\mathbf{U}\mathbf{A}\mathbf{U}^{-1})\mathbf{x}' \end{aligned} \quad (16.27)$$

Changing the coordinate system thus changes a matrix by pre- and post-multiplication of a unitary matrix and its inverse, a procedure called a *similarity transformation*. Since the \mathbf{U} matrix describes a rotation of the coordinate system in an arbitrary direction, one person's \mathbf{U} may be another person's \mathbf{U}^{-1} . There is thus no significance whether the transformation is written as $\mathbf{U}^{-1}\mathbf{A}\mathbf{U}$ or $\mathbf{U}\mathbf{A}\mathbf{U}^{-1}$, and for an orthogonal transformation matrix ($\mathbf{U}^{-1} = \mathbf{U}^t$), the transformation may also be written as $\mathbf{U}^t\mathbf{A}\mathbf{U}$ or $\mathbf{U}\mathbf{A}\mathbf{U}^t$.

For the case of a symmetric ($a_{12} = a_{21}$) 2×2 matrix, the similarity transformed matrix elements are given in eq. (16.28).

$$\begin{aligned} \mathbf{A}' &= \mathbf{U}\mathbf{A}\mathbf{U}^{-1} \\ \begin{pmatrix} a'_{11} & a'_{12} \\ a'_{12} & a'_{22} \end{pmatrix} &= \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} a_{11} & a_{12} \\ a_{12} & a_{22} \end{pmatrix} \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} \\ a'_{11} &= a_{11} \cos^2 \alpha + a_{22} \sin^2 \alpha + 2a_{12} \cos \alpha \sin \alpha \\ a'_{22} &= a_{22} \cos^2 \alpha + a_{11} \sin^2 \alpha - 2a_{12} \cos \alpha \sin \alpha \\ a'_{12} &= a_{12}(\cos^2 \alpha - \sin^2 \alpha) + (a_{22} - a_{11})\cos \alpha \sin \alpha \end{aligned} \quad (16.28)$$

The off-diagonal element a'_{12} can be made to vanish by choosing a specific rotational angle, as shown in eq. (16.29).

$$\begin{aligned} a'_{12} &= a_{12}(\cos^2 \alpha - \sin^2 \alpha) + (a_{22} - a_{11})\cos \alpha \sin \alpha = 0 \\ a_{12}(\cos^2 \alpha - \sin^2 \alpha) &= (a_{11} - a_{22})\cos \alpha \sin \alpha \\ a_{12}(\cos 2\alpha) &= (a_{22} - a_{11})\left(\frac{1}{2} \sin 2\alpha\right) \\ \tan(2\alpha) &= \frac{2a_{12}}{(a_{11} - a_{22})} \end{aligned} \quad (16.29)$$

In the new coordinate system, the \mathbf{A}' matrix is simplified, as it only contains diagonal elements.

$$\mathbf{A}' = \begin{pmatrix} a'_{11} & a'_{12} \\ a'_{12} & a'_{22} \end{pmatrix} = \begin{pmatrix} \varepsilon_1 & 0 \\ 0 & \varepsilon_2 \end{pmatrix} = \boldsymbol{\varepsilon} \quad (16.30)$$

An $N \times N$ Hermitian (or real symmetric) matrix can always be brought to a diagonal form by a multi-dimensional rotation of the coordinate system, and there are efficient standard computational procedures for diagonalizing matrices. The simplest method consists of an iterative series of 2×2 rotations as in eq. (16.28) which reduces the off-diagonal elements to zero. The rotational matrix \mathbf{U} contains elements corresponding to products of cosines and sines of rotational angles. The elements of the \mathbf{A} matrix in diagonal form (ϵ) are called *eigenvalues*, and the columns of the unitary rotation matrix are called *eigenvectors*. In matrix notation the diagonalization can be written as in eq. (16.31).

$$\epsilon = \mathbf{U} \mathbf{A} \mathbf{U}^{-1} \quad (16.31)$$

A Hermitian matrix will always have real eigenvalues and orthogonal eigenvectors. Matrix diagonalizations play an important role in many areas of computational chemistry, and scientific computations in general, since they correspond to selecting a coordinate system where the variables are (approximately) independent of each other. Furthermore, the magnitude of the eigenvalues indicates the variation along that particular direction. For applications with many variables, it may be possible to describe a significant fraction of the whole variation by taking only a few selected eigenvector directions into account, and this forms the basis for *principal component analysis*, as discussed in Section 17.4.3.

It can be shown that a matrix determinant is independent of a change in the coordinate system, and in the diagonal representation the determinant is simply the product of the eigenvalues. A non-zero determinant is thus equivalent to all the eigenvalues being different from zero. Furthermore, the *trace* of a matrix, defined as the sum of the diagonal elements, is also invariant to a change of the coordinate system, as can be verified for a 2×2 case from eq. (16.28). In the diagonal representation the trace is given by the sum of the eigenvalues.

An alternative way of introducing matrix eigenvalues and eigenvectors is to require non-zero \mathbf{x} -solutions to eq. (16.32).

$$\begin{aligned} \mathbf{A} \mathbf{x} &= \epsilon \mathbf{x} \\ (\mathbf{A} - \epsilon \mathbf{I}) \mathbf{x} &= \mathbf{0} \end{aligned} \quad (16.32)$$

This is a set of linear equations in the form of eq. (16.21) with the right-hand side being zero, and a non-trivial solution therefore only exists when the determinant is zero.

$$|\mathbf{A} - \epsilon \mathbf{I}| = 0 \quad (16.33)$$

Expansion of the determinant (16.33) according to eq. (16.18) produces an N th-order polynomial in ϵ , which can be solved to give N roots (eigenvalues). If some of these are identical, they are called *degenerate* eigenvalues. For each eigenvalue eq. (16.32) can be solved to produce the corresponding eigenvector. In the non-degenerate case (all ϵ_i being different) the one free parameter can be fixed by normalization. For degenerate eigenvectors, the normalization condition must be augmented with a mutual orthogonality condition (Section 16.4) in order to fix all the free parameters.

In the coordinate system (\mathbf{x}') where the \mathbf{A}' matrix is diagonal, it is easy to see that eq. (16.33) holds since the $(\mathbf{A}' - \epsilon \mathbf{I})$ matrix has at least one column consisting of only

zeroes. In this diagonal representation, it is furthermore clear that the eigenvectors are simply unit vectors along the primed coordinate axes, and the eigenvectors in the unprimed coordinate system are therefore given by the elements of the \mathbf{U}^t transformation matrix.

$$\begin{aligned}\mathbf{A}'\mathbf{x}' &= \boldsymbol{\varepsilon}\mathbf{x}' \\ \begin{pmatrix} \varepsilon_1 & 0 \\ 0 & \varepsilon_2 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} &= \varepsilon_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix}; \quad \begin{pmatrix} \varepsilon_1 & 0 \\ 0 & \varepsilon_2 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \varepsilon_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ \mathbf{x} &= \mathbf{U}^t \mathbf{x}'\end{aligned}\tag{16.34}$$

While the polynomial method can be used for solving small eigenvalue problems by hand, all computational implementations rely on iterative similarity transform methods for bringing the matrix to a diagonal form. The simplest of these is the *Jacobi* method, where a sequence of 2×2 rotations analogous to eqs (16.28)–(16.30) can be used to bring all the off-diagonal elements below a suitable threshold value.

In the diagonal form, the matrix $\boldsymbol{\varepsilon}$ contains only elements along the diagonal. The diagonal elements can be treated like regular numbers, allowing calculation of functions of matrices. Calculating for example $\mathbf{A}^{1/2}$ proceeds by first transforming it to a diagonal form, taking the square root of the diagonal elements, and back-transforming to the original coordinate system. This procedure in general allows calculation of functions of matrices, such as $e^{\mathbf{A}}$, $\ln(\mathbf{A})$ or $\cos(\mathbf{A})$.

$$\begin{aligned}\mathbf{A} = \begin{pmatrix} a_{11} & a_{12} & \cdots \\ a_{21} & a_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} &\xrightarrow{\mathbf{U}} \boldsymbol{\varepsilon} = \begin{pmatrix} \varepsilon_1 & 0 & \cdots \\ 0 & \varepsilon_2 & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \xrightarrow{f} \\ f(\boldsymbol{\varepsilon}) = \begin{pmatrix} f(\varepsilon_1) & 0 & \cdots \\ 0 & f(\varepsilon_2) & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} &\xrightarrow{\mathbf{U}^{-1}} f(\mathbf{A})\end{aligned}$$

Figure 16.4 Construction of functions of matrices

This also provides an alternative way of calculating the inverse of a matrix, by simply taking the inverse of the eigenvalues in the diagonal representation and back-transforming the matrix to the original representation.

In some cases, the \mathbf{A} matrix may have eigenvalues that are zero or nearly so. The number of non-zero eigenvalues is called the *rank* of the matrix \mathbf{A} , and corresponds to the number of independent rows/columns in the matrix. In actual applications, it is rare that an eigenvalue is exactly zero, but a very small value will clearly give numerical problems for constructing matrices such as \mathbf{A}^{-1} or $\ln(\mathbf{A})$. The ratio between the largest and smallest eigenvalue is called the *condition number*, and large values ($>10^6$) indicate that the \mathbf{A} matrix is close to having linear dependencies. *Singular value decomposition* constructs \mathbf{A}^{-1} by inverting only those eigenvalues larger than a suitable threshold and setting the rest to zero, before back-transformation to the original coordinate system.

For an $N \times M$ ($N > M$) rectangular matrix \mathbf{A} , a *generalized inverse* can be defined by the matrix $(\mathbf{A}^t \mathbf{A})^{-1} \mathbf{A}^t$. Such generalized inverse matrices correspond to obtaining

the best solution in a least squares sense for an over-determined system of linear equations, as for example arises in statistical applications. Consider for example a system of equations analogous to eq. (16.21), but with more b solution elements than x variables ($n > m$).

$$\begin{aligned}
 a_{11}x_1 + a_{12}x_2 + \dots + a_{1m}x_m &= b_1 \\
 a_{21}x_1 + a_{22}x_2 + \dots + a_{2m}x_m &= b_2 \\
 \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots &= \vdots \\
 a_{n1}x_1 + a_{n2}x_2 + \dots + a_{nm}x_m &= b_n
 \end{aligned} \tag{16.35}$$

$$\begin{aligned}
 \mathbf{Ax} &= \mathbf{b} \\
 \mathbf{A}^t \mathbf{Ax} &= \mathbf{A}^t \mathbf{b} \\
 \mathbf{x} &= (\mathbf{A}^t \mathbf{A})^{-1} \mathbf{A}^t \mathbf{b}
 \end{aligned}$$

Multiplication from the left by \mathbf{A}^t and by the inverse of $\mathbf{A}^t \mathbf{A}$ leads to the formal solution, i.e. $(\mathbf{A}^t \mathbf{A})^{-1} \mathbf{A}^t$ acts as the inverse to the rectangular \mathbf{A} matrix.

16.2.1 Examples of changing the coordinate system

From the “separability” theorem (Section 1.6.3) it follows that if an operator (e.g. the Hamiltonian) depending on N coordinates can be written as a sum of operators that only depend on one coordinate, the corresponding N coordinate wave function can be written as a product of one-coordinate functions, and the total energy as a sum of energies.

$$\begin{aligned}
 \mathbf{H}(x_1, x_2, x_3, \dots) \Psi(x_1, x_2, x_3, \dots) &= E_{\text{tot}} \Psi(x_1, x_2, x_3, \dots) \\
 \mathbf{H}(x_1, x_2, x_3, \dots) &= \sum_i \mathbf{h}_i(x_i) \\
 \mathbf{h}_i(x_i) \phi_i(x_i) &= \varepsilon_i \phi_i(x_i) \\
 E_{\text{tot}} &= \sum_i \varepsilon_i \\
 \Psi(x_1, x_2, x_3, \dots) &= \prod_i \phi_i(x_i)
 \end{aligned} \tag{16.36}$$

Instead of solving one equation with N variables, the problem is transformed into solving N equations with only one variable. When the differential operator is transformed into a matrix representation, the separation is equivalent to finding a coordinate system where the representation is diagonal.

Consider a matrix \mathbf{A} expressed in a coordinate system $\{x_1, x_2, x_3, \dots, x_N\}$. The coordinate axes are the x_i vectors, and these may be simple Cartesian axes, or one-variable functions, or many-variable functions. The matrix \mathbf{A} is typically defined by an operator working on the coordinates. Some examples are:

- (1) The force constant matrix in Cartesian coordinates (Section 13.5.3)
- (2) The Fock matrix in basis functions (atomic orbitals, Section 3.5)
- (3) The CI matrix in Slater determinants (Section 4.2).

Finding the coordinates where these matrices are diagonal corresponds to finding:

- (1) The vibrational normal coordinates
- (2) The molecular orbitals
- (3) The state coefficients, i.e. CI wave function(s).

The coordinate axes are usually orthonormal, but this is not a requirement, since they can be orthogonalized by the methods in Section 16.4.

16.2.2 Vibrational normal coordinates

The potential energy is approximated by a second-order Taylor expansion around the stationary geometry \mathbf{x}_0 .

$$V(\mathbf{x}) \approx V(\mathbf{x}_0) + \left(\frac{dV}{d\mathbf{x}} \right)^t (\mathbf{x} - \mathbf{x}_0) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_0)^t \left(\frac{d^2V}{d\mathbf{x}^2} \right) (\mathbf{x} - \mathbf{x}_0) \quad (16.37)$$

The energy for the expansion point, $V(\mathbf{x}_0)$, may be chosen as zero, and the first derivative is zero since \mathbf{x}_0 is a stationary point.

$$V(\Delta\mathbf{x}) = \frac{1}{2} \Delta\mathbf{x}^t \mathbf{F} \Delta\mathbf{x} \quad (16.38)$$

Here \mathbf{F} is a $3N_{\text{atom}} \times 3N_{\text{atom}}$ (force constant) matrix containing the second derivatives of the energy with respect to the coordinates.

The nuclear Schrödinger equation for an N_{atom} system is given by eq. (16.39).

$$\left[- \sum_{i=1}^{3N_{\text{atom}}} \left(\frac{1}{2m_i} \frac{\partial^2}{\partial x_i^2} \right) + \frac{1}{2} \Delta\mathbf{x}^t \mathbf{F} \Delta\mathbf{x} \right] \Psi_{\text{nuc}} = E_{\text{nuc}} \Psi_{\text{nuc}} \quad (16.39)$$

Eq. (16.39) is first transformed to mass-dependent coordinates by a \mathbf{G} matrix containing the inverse square root of atomic masses (note that *atomic*, not *nuclear*, masses are used, which is in line with the Born–Oppenheimer approximation that the electrons follow the nuclei).

$$\begin{aligned} y_i &= \sqrt{m_i} \Delta x_i \\ \frac{\partial^2}{\partial y_i^2} &= \frac{1}{m_i} \frac{\partial^2}{\partial x_i^2} \\ G_{ij} &= \frac{1}{\sqrt{m_i m_j}} \end{aligned} \quad (16.40)$$

$$\left[- \sum_{i=1}^{3N_{\text{atom}}} \left(\frac{1}{2} \frac{\partial^2}{\partial y_i^2} \right) + \frac{1}{2} \mathbf{y}^t (\mathbf{F} \cdot \mathbf{G}) \mathbf{y} \right] \Psi_{\text{nuc}} = E_{\text{nuc}} \Psi_{\text{nuc}}$$

A unitary transformation is then introduced that diagonalizes the $\mathbf{F} \cdot \mathbf{G}$ (entrywise product, eq. (16.10)) matrix, yielding eigenvalues ε_i and eigenvectors \mathbf{q}_i . The kinetic energy operator is still diagonal in these coordinates.

$$\begin{aligned}
 q &= Uy \\
 \left[- \sum_{i=1}^{3N_{\text{atom}}} \left(\frac{1}{2} \frac{\partial^2}{\partial q_i^2} \right) + \frac{1}{2} \mathbf{q}^t (\mathbf{U}(\mathbf{F} \cdot \mathbf{G}) \mathbf{U}^t) \mathbf{q} \right] \Psi_{\text{nuc}} &= E_{\text{nuc}} \Psi_{\text{nuc}} \\
 \left[- \sum_{i=1}^{3N_{\text{atom}}} \left(\frac{1}{2} \frac{\partial^2}{\partial q_i^2} + \frac{1}{2} \varepsilon_i q_i^2 \right) \right] \Psi_{\text{nuc}} &= E_{\text{nuc}} \Psi_{\text{nuc}} \\
 \sum_{i=1}^{3N_{\text{atom}}} [\mathbf{h}_i(q_i)] &= \Psi_{\text{nuc}} = E_{\text{nuc}} \Psi_{\text{nuc}}
 \end{aligned} \tag{16.41}$$

In the \mathbf{q} -coordinate system, the *vibrational normal coordinates*, the $3N_{\text{atom}}$ -dimensional Schrödinger equation can be separated into $3N_{\text{atom}}$ one-dimensional Schrödinger equations, which are just in the form of a standard harmonic oscillator, with the solutions being Hermite polynomials in the \mathbf{q} -coordinates. The eigenvectors of the $\mathbf{F} \cdot \mathbf{G}$ matrix are the (mass-weighted) vibrational normal coordinates, and the eigenvalues ε_i are related to the vibrational frequencies as shown in eq. (16.42) (analogous to eq. (13.31)).

$$v_i = \frac{1}{2\pi} \sqrt{\varepsilon_i} \tag{16.42}$$

When this procedure is carried out in Cartesian coordinates, there should be six (five for a linear molecule) eigenvalues of the $\mathbf{F} \cdot \mathbf{G}$ matrix being exactly zero, corresponding to the translational and rotational modes. In real calculations, however, these values are not exactly zero. The three translational modes usually have “frequencies” very close to zero, typically less than 0.01 cm^{-1} . The deviation from zero is due to the fact that numerical operations are only carried out with a finite precision, and the accumulations of errors will typically give inaccuracies in v of this magnitude. The residual “frequencies” for the rotational modes, however, may often be as large as 10 – 50 cm^{-1} . This is due to the fact that the geometry cannot be optimized to a gradient of exactly zero, again due to numerical considerations. Typically, the geometry optimization is considered converged if the root mean square (RMS) gradient is less than $\sim 10^{-4}$ – 10^{-5} au , corresponding to the energy being converged to $\sim 10^{-5}$ – 10^{-6} au . The residual gradient shows up as vibrational frequencies for the rotations of the above magnitude.

If there are real frequencies of the same magnitude as the “rotational frequencies”, mixing may occur and result in inaccurate values for the “true” vibrations. For this reason, the translational and rotational degrees of freedom are normally removed by projection (Section 16.4) from the force constant matrix before diagonalization.

If the stationary point is a minimum on the energy surface, the eigenvalues of the \mathbf{F} and $\mathbf{F} \cdot \mathbf{G}$ matrices are all positive. If, however, the stationary point is a transition state (TS), one (and only one) of the eigenvalues is negative. This corresponds to the energy being a maximum in one direction and a minimum in all other directions. The “frequency” for the “vibration” along the eigenvector with a negative eigenvalue will formally be imaginary, as it is the square root of a negative number (eq. (16.42)). The corresponding eigenvector is the direction leading downhill from the TS towards the reactant and product. *At the TS, the eigenvector for the imaginary frequency is the*

reaction coordinate. The whole reaction path may be calculated by sliding downhill to each side from the TS. This can be performed by taking a small step along the TS eigenvector, calculating the gradient and taking a small step in the negative gradient direction. The negative of the gradient always points downhill, and by taking a sufficiently large number of such steps an energy minimum is eventually reached. This is equivalent to a steepest descent minimization, but more efficient methods are available (see Section 12.8 for details). The reaction path in mass-weighted coordinates is called the *Intrinsic Reaction Coordinate* (IRC).

The vibrational Hamiltonian is completely separable within the harmonic approximation, with the vibrational energy being a sum of individual energy terms and the nuclear wave function being a product of harmonic oscillator functions (Hermite polynomial in the normal coordinates). When anharmonic terms are included in the potential, the Hamiltonian is no longer separable, and the resulting nuclear Schrödinger equation can be solved by techniques completely analogous to those used for solving the electronic problem. The vibrational SCF method is analogous to the electronic Hartree–Fock method, with the nuclear harmonic oscillator functions playing the same role as the orbitals in electronic structure theory. Corrections beyond the mean-field approximation can be added by configuration interaction, perturbation theory or coupled cluster methods.²

It should be noted that the force constant matrix can be calculated at any geometry, but the transformation to normal coordinates is only valid at a stationary point, i.e. where the first derivative is zero. At a non-stationary geometry, a set $3N_{\text{atom}} - 7$ *generalized frequencies* may be defined by removing the gradient direction from the force constant matrix (for example by projection techniques, Section 16.4) before transformation to normal coordinates.

16.2.3 Energy of a Slater determinant

The variational problem is to minimize the energy of a single Slater determinant by choosing suitable values for the molecular orbital (MO) coefficients, under the constraint that the MOs remain orthonormal. With ϕ being an MO written as a linear combination of the basis functions (atomic orbitals) χ , this leads to a set of secular equations, \mathbf{F} being the Fock matrix, \mathbf{S} the overlap matrix and \mathbf{C} containing the coefficients (Section 3.5).

$$\begin{aligned}\phi_i &= \sum_{\alpha=1}^{M_{\text{basis}}} c_{\alpha i} \chi_{\alpha} \\ F_{\alpha\beta} &= \langle \chi_{\alpha} | \mathbf{F} | \chi_{\beta} \rangle \\ S_{\alpha\beta} &= \langle \chi_{\alpha} | \chi_{\beta} \rangle \\ \mathbf{FC} &= \mathbf{SC}\epsilon\end{aligned}\tag{16.43}$$

The basis functions (coordinate system) in this case are non-orthogonal, with the overlap elements contained in the \mathbf{S} matrix. By multiplying from the left by $\mathbf{S}^{-1/2}$ and inserting a unit matrix written in the form $\mathbf{S}^{-1/2}\mathbf{S}^{1/2}$, eq. (16.43) may be reformulated as eq. (16.44).

$$\left\{ \mathbf{S}^{-1/2} \mathbf{F} \mathbf{S}^{-1/2} \right\} \left\{ \mathbf{S}^{1/2} \mathbf{C} \right\} = \left\{ \mathbf{S}^{-1/2} \mathbf{S}^{1/2} \right\} \left\{ \mathbf{S}^{1/2} \mathbf{C} \right\} \boldsymbol{\epsilon}$$

$$\mathbf{F}' \mathbf{C}' = \mathbf{C}' \boldsymbol{\epsilon}$$
(16.44)

The latter equation is now in a standard form for determining the eigenvalues of the \mathbf{F}' matrix. The eigenvectors contained in \mathbf{C}' can then be back-transformed to the original coordinate system ($\mathbf{C} = \mathbf{S}^{-1/2} \mathbf{C}'$). This is an example of a symmetrical orthogonalization (Section 16.4) of the initial coordinate system, the basis functions χ . Solving eq. (16.43) corresponds to rotating the original space of basis functions into one of molecular orbitals where the Fock matrix is diagonal.

16.2.4 Energy of a CI wave function

The variational problem may again be formulated as a secular equation, where the coordinate axes are many-electron functions (Slater determinants) Φ_i , which are orthogonal (Section 4.2).

$$\Psi_{\text{CI}} = \sum_i a_i \Phi_i$$

$$H_{ij} = \langle \Phi_i | \mathbf{H} | \Phi_j \rangle$$

$$\mathbf{H} \mathbf{a} = \mathbf{E} \mathbf{a}$$
(16.45)

The \mathbf{a} matrix contains the coefficients of the CI wave function. This problem may again be considered as selecting a basis where the Hamiltonian operator is diagonal (eq. (4.6) and Figure 4.5). In the initial coordinate system, the Hamiltonian matrix will have many off-diagonal elements, but it can be diagonalized by a suitable unitary transformation. The diagonal elements are energies of many-electron CI wave functions, being approximations to the ground and excited states. The corresponding eigenvectors contain the expansion coefficients a_i .

16.2.5 Computational considerations

Finally, a few practical considerations. The time required for diagonalizing a matrix grows as the *cube* of the size of the matrix, and the amount of computer memory necessary for storing the matrix grows as the *square* of the size. Diagonalizing matrices up to $\sim 100 \times 100$ takes insignificant amounts of time, unless there are extraordinarily many such matrices. Matrices up to $\sim 1000 \times 1000$ pose no particular problems, although some consideration should be made as to whether the time required for diagonalization is significant relative to other operations. Matrices larger than this require consideration. Just storing all the elements in a 10000×10000 matrix takes ~ 1 GB of memory (or disk space) on a computer. Determining all eigenvalues and eigenvectors of such a matrix takes a long time. For large-scale problems in quantum chemistry, however, one is usually not interested in all the eigenvalues and eigenvectors. In solving the CI matrix equation, for example, typically only the lowest eigenvalue and eigenvector is of interest, since this is the ground state energy and wave function. Large-scale diagonalizations are therefore normally solved by special iterative schemes, which extract a few selected roots and eigenvectors.

16.3 Coordinates, Functions, Functionals, Operators and Superoperators

A *function* is a recipe for producing a scalar from another set of scalars, for example calculating the energy E from a set of (nuclear) coordinates \mathbf{x} .

$$E(\mathbf{x}) \therefore \mathbf{x} \rightarrow E \quad (16.46)$$

A *functional* is a recipe for producing a scalar from a function, for example calculating the exchange–correlation energy E_{xc} from an electron density ρ depending on a set of (electronic) coordinates \mathbf{x} .

$$E_{\text{xc}}[\rho(\mathbf{x})] \therefore \mathbf{x} \rightarrow \rho \rightarrow E_{\text{xc}} \quad (16.47)$$

An *operator* is a recipe for producing a function from another function, for example the kinetic energy operator acting on a wave function. The operator in this case consists of differentiating the function twice with respect to the coordinates, adding the results and dividing by twice the particle mass.

$$\begin{aligned} \mathbf{T}\Psi(\mathbf{x}) &= \mathbf{T}\Psi(x, y, z) \\ \mathbf{T} &= \frac{1}{2m} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \right) \end{aligned} \quad (16.48)$$

A *superoperator* is a recipe for producing an operator from another operator. This level of abstraction is rarely used, but is for example employed in some formulations of propagator theory (Section 10.9).

$$\hat{\mathbf{O}}(\mathbf{h}(f(\mathbf{x}))) \quad (16.49)$$

In the abstract function (or operator) space, often called a *Hilbert* space, it is possible to consider the functions (or operators) as vectors. The *bra–ket* notation is defined as in eq. (16.50).

$$\begin{aligned} \text{bra: } \langle \mathbf{f} | &= \mathbf{f}^*(\mathbf{x}) \\ \text{ket: } | \mathbf{f} \rangle &= \mathbf{f}(\mathbf{x}) \end{aligned} \quad (16.50)$$

The equivalent of a vector dot product is defined as the integral of the product of the two functions.

$$\langle \mathbf{f} | \mathbf{g} \rangle = \int \mathbf{f}^*(\mathbf{x}) \mathbf{g}(\mathbf{x}) d\mathbf{x} \quad (16.51)$$

The combination of a bra and a ket is called a *bracket*, and the bracket notation is often also used for the dot product of regular coordinate vectors (eq. (16.11)). By analogy with coordinate vectors, the bracket of two functions measures the “angle” or overlap between the functions, with a value of zero indicating that the two functions are orthogonal.

The norm of a function is defined as the square root of the bracket of the function with itself.

$$|\mathbf{f}(\mathbf{x})| = \sqrt{\langle \mathbf{f} | \mathbf{f} \rangle} = \sqrt{\int \mathbf{f}^*(\mathbf{x}) \mathbf{f}(\mathbf{x}) d\mathbf{x}} \quad (16.52)$$

The bracket notation can also be used in connection with functions and operators, as in the example in eq. (16.53).

$$\langle \mathbf{f} | \mathbf{O} | \mathbf{g} \rangle = \int \mathbf{f}^*(\mathbf{x}) \mathbf{O} \mathbf{g}(\mathbf{x}) d\mathbf{x} \quad (16.53)$$

Operator algebra shares the characteristic with matrix algebra; indeed, matrices can be considered as representations of operators in a given set of functions (coordinate system).

Most operators in computational chemistry are *linear*.

$$\begin{aligned} \mathbf{O}(\mathbf{f} + \mathbf{g}) &= \mathbf{O}\mathbf{f} + \mathbf{O}\mathbf{g} \\ \mathbf{O}(c\mathbf{f}) &= c\mathbf{O}\mathbf{f} \end{aligned} \quad (16.54)$$

Operators are *associative*, but not necessarily *commutative*.

$$\begin{aligned} \mathbf{O}(\mathbf{P}\mathbf{Q}) &= (\mathbf{O}\mathbf{P})\mathbf{Q} \\ \mathbf{O}\mathbf{P} &\neq \mathbf{P}\mathbf{O} \end{aligned} \quad (16.55)$$

The commutator of two operators is denoted with square brackets and defined as in eq. (16.56).

$$[\mathbf{O}, \mathbf{P}] = \mathbf{O}\mathbf{P} - \mathbf{P}\mathbf{O} \quad (16.56)$$

Two operators are said to commute when eq. (16.56) is zero.

Operator eigenvalues and eigenfunctions are defined analogously to eq. (16.32).

$$\mathbf{O}\mathbf{f} = \epsilon \mathbf{f} \quad (16.57)$$

16.3.1 Differential operators

Differential operators, which describe the first- and higher order variations of functions, represent an important class of operators in computational chemistry. For a simple one-dimensional function, the first derivative is given by the normal rules for differentiation.

$$f(x) = x^2 \Rightarrow \frac{df}{dx} = 2x \quad (16.58)$$

For a multi-dimensional *scalar function* (i.e. each point in space is associated with a number), the first derivative is a vector containing all the partial derivatives. The corresponding operator is called the *gradient* and denoted with ∇ .

$$\begin{aligned} f(x, y, z) &= x + y^2 + z^3 + xy + xz \\ \nabla f(x, y, z) &= \left\{ \frac{\partial f}{\partial x} \quad \frac{\partial f}{\partial y} \quad \frac{\partial f}{\partial z} \right\} = \{1 + y + z \quad 2y + x \quad 3z^2 + x\} \end{aligned} \quad (16.59)$$

The gradient vector points in the direction where the function increases most.

Two choices are possible for defining the first derivative of a *vector function* (i.e. each point in space is associated with a vector). The *divergence* is denoted with ∇ and produces a *scalar*.

$$\nabla \cdot f(x, y, z) = \frac{\partial f}{\partial x} + \frac{\partial f}{\partial y} + \frac{\partial f}{\partial z} \quad (16.60)$$

The divergence measures how much the vector field “dilutes” or “contracts” at a given point. Alternatively, the first derivative may be defined as the *curl*, denoted with $\nabla \times$, which produces a *vector*.

$$\nabla \times f(x, y, z) = \left\{ \left(\frac{\partial f}{\partial y} - \frac{\partial f}{\partial z} \right) \left(\frac{\partial f}{\partial z} - \frac{\partial f}{\partial x} \right) \left(\frac{\partial f}{\partial x} - \frac{\partial f}{\partial y} \right) \right\} \quad (16.61)$$

The curl describes how fast the vector field rotates, i.e. how rapidly and in which direction the field changes.

Given the above three definitions of first derivatives, there are nine possible combinations for defining second derivatives. Four of these are invalid since the gradient only works on scalar fields and the divergence and curl only work on vector fields. Two of the remaining five combinations can be shown to be zero, leaving only three interesting combinations. Figure 16.5 indicates the action of the three first derivatives (conversion from/to scalar and vector) and their combinations.

		<i>First</i>		
		∇ <small>$S \rightarrow V$</small>	$\nabla \cdot$ <small>$V \rightarrow S$</small>	$\nabla \times$ <small>$V \rightarrow V$</small>
		∇ <small>$S \rightarrow V$</small>	I	$\nabla \nabla \cdot$ I
<i>Second</i>	$\nabla \cdot$ <small>$V \rightarrow S$</small>	∇^2	I	0
	$\nabla \times$ <small>$V \rightarrow V$</small>	0	I	$\nabla \times \nabla \times$

Figure 16.5 Possibilities for second derivatives of multi-dimensional functions, with I indicating an invalid combination

The divergence of the gradient is commonly denoted the *Laplacian*, and is for example involved in the (non-relativistic) quantum mechanical kinetic energy operator. It operates on a scalar function and produces a scalar function.

$$\nabla \cdot \nabla f(x, y, z) = \nabla^2 f(x, y, z) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} + \frac{\partial^2 f}{\partial z^2} \quad (16.62)$$

The Laplacian measures the local depletion or concentration of the function. The two other combinations produce a vector from a vector function and are used less commonly.

16.4 Normalization, Orthogonalization and Projection

The vectors (functions) of a coordinate system may in some cases be given naturally by the problem, and these are not always normalized or orthogonal. For computational purposes, however, it is often advantageous to work in an orthonormal coordinate system. We first note that normalization is trivially obtained by simply scaling each vector by the inverse of its length.

$$\begin{aligned}
\langle \mathbf{x} | \mathbf{x} \rangle &= N^2 \\
\mathbf{x}' &= N^{-1} \mathbf{x} \\
\langle \mathbf{x}' | \mathbf{x}' \rangle &= 1
\end{aligned} \tag{16.63}$$

The orthogonalization of a set of non-orthogonal vectors can be done in many ways, but the two most commonly used are *Gram–Schmidt* and *symmetrical* orthogonalization. The Gram–Schmidt procedure corresponds to sequentially removing the component of all previous vectors, and re-normalizing the remaining component. It should be noted that the final set of orthogonal vectors will depend on the selection of the first vector and the order in which the remaining vectors are orthogonalized, although the total space spanned will of course be the same.

$$\begin{aligned}
\mathbf{x}'_1 &= N^{-1} \mathbf{x}_1 \\
\mathbf{x}'_2 &= N^{-1} (\mathbf{x}_2 - \langle \mathbf{x}_2 | \mathbf{x}'_1 \rangle \mathbf{x}'_1) \\
\mathbf{x}'_3 &= N^{-1} (\mathbf{x}_3 - \langle \mathbf{x}_3 | \mathbf{x}'_1 \rangle \mathbf{x}'_1 - \langle \mathbf{x}_3 | \mathbf{x}'_2 \rangle \mathbf{x}'_2) \\
\mathbf{x}'_k &= N^{-1} \left(\mathbf{x}_k - \sum_{i=1}^{k-1} \langle \mathbf{x}_k | \mathbf{x}'_i \rangle \mathbf{x}'_i \right)
\end{aligned} \tag{16.64}$$

A *symmetrical* orthogonalization corresponds to a transformation that has the property $\mathbf{X}^\dagger \mathbf{S} \mathbf{X} = \mathbf{I}$, where \mathbf{X} denotes all the \mathbf{x}_i coordinate vectors and \mathbf{S} contains the overlap elements. One such transformation is given by the inverse square root of the overlap matrix ($\mathbf{X} = \mathbf{S}^{-1/2}$), a procedure used in solving the self-consistent field equations in Hartree–Fock and Kohn–Sham theories, and for performing the Löwdin population analysis (Section 9.1).

$$\begin{aligned}
S_{ij} &= \langle \mathbf{x}_i | \mathbf{x}_j \rangle \\
\mathbf{x}' &= \mathbf{S}^{-1/2} \mathbf{x}
\end{aligned} \tag{16.65}$$

Alternatively, a *canonical orthogonalization* can be performed by using the unitary matrix obtained by diagonalizing the overlap matrix and weighting by the inverse square root of the eigenvalues.

$$\mathbf{x}' = \left(\mathbf{U} \mathbf{\epsilon}^{-1/2} \right) \mathbf{x} \tag{16.66}$$

The advantage of a canonical orthogonalization is that it allows for handling (near-) linear dependencies in the basis by truncating the transformation matrix by removal of columns with eigenvalues smaller than a suitable cutoff value.

In multi-dimensional coordinate systems it may often be advantageous to work in a subset of the full coordinate system. The component of a function \mathbf{f} along a specific (unit) coordinate vector \mathbf{x}_k is given by the projection.

$$f_k = \langle \mathbf{f} | \mathbf{x}_k \rangle \tag{16.67}$$

For a matrix representation of an operator, \mathbf{A} , the projection onto the \mathbf{x}_k subspace is given by pre- and post-multiplying with a \mathbf{Q}_k matrix defined as the outer vector product of \mathbf{x}_k , or the function equivalent in a ket-bra notation.

$$\begin{aligned}\mathbf{Q}_k &= \mathbf{x}_k \mathbf{x}_k^t = |\mathbf{x}_k\rangle\langle\mathbf{x}_k| \\ \mathbf{A}_k &= \mathbf{Q}_k^t \mathbf{A} \mathbf{Q}_k\end{aligned}\quad (16.68)$$

The reverse process, removing the \mathbf{x}_k subspace, is done by projecting the \mathbf{x}_k subspace direction out by the complementary matrix \mathbf{P}_k , with \mathbf{I} being a unit matrix.

$$\begin{aligned}\mathbf{P}_k &= \mathbf{I} - \mathbf{Q}_k \\ \mathbf{f}_{-k} &= \mathbf{P}_k \mathbf{f} = \mathbf{f} - \langle\mathbf{f}|\mathbf{x}_k\rangle\mathbf{x}_k \\ \mathbf{A}_{-k} &= \mathbf{P}_k^t \mathbf{A} \mathbf{P}_k\end{aligned}\quad (16.69)$$

Projection onto smaller subspaces can be done similarly by adding more vectors to the projection matrix. For the case of removing the translational and rotational degrees of freedom from the vibrational normal coordinates, for example, the (normalized) vector in eq. (16.70) describes a translation in the x -direction.

$$\mathbf{t}_x^t = \frac{1}{\sqrt{N_{\text{atom}}}}(1, 0, 0, 1, 0, 0, 1, 0, 0, \dots) \quad (16.70)$$

The superscript t indicates that \mathbf{t}_x^t is a row vector. The \mathbf{T}_x matrix in eq. (16.71) removes the direction corresponding to translation in the x -direction.

$$\mathbf{T}_x = \mathbf{I} - \mathbf{t}_x \mathbf{t}_x^t \quad (16.71)$$

Extending this to include vectors for all three translational and rotational modes gives a projection matrix for removing the six (five) translational and rotational degrees of freedom.

$$\mathbf{P} = \mathbf{I} - \mathbf{t}_x \mathbf{t}_x^t - \mathbf{t}_y \mathbf{t}_y^t - \mathbf{t}_z \mathbf{t}_z^t - \mathbf{r}_a \mathbf{r}_a^t - \mathbf{r}_b \mathbf{r}_b^t - \mathbf{r}_c \mathbf{r}_c^t \quad (16.72)$$

The \mathbf{r} -vectors are derived from the atomic coordinates and principal axes of inertia determined by diagonalization of the matrix of inertia (eq. (13.27)).³ By forming the matrix product $\mathbf{P}^t \mathbf{F} \mathbf{P}$, the translation and rotational directions are removed from the force constant matrix, and consequently the six (five) trivial vibrations become exactly zero (within the numerical accuracy of the machine).

In some cases, it can be useful to apply an internal double projection onto an auxiliary set of functions \mathbf{k}_i , often called insertion of a *resolution of the identity*. This may for example allow separation of a composite operator.

$$\begin{aligned}\text{Value} &= \langle\mathbf{f}|\mathbf{O}_1 \mathbf{O}_2|\mathbf{g}\rangle \\ \mathbf{I} &= \sum_{ij} |\mathbf{k}_i\rangle M_{ij} \langle\mathbf{k}_j| \quad M_{ij} = \langle\mathbf{k}_i|\mathbf{k}_j\rangle^{-1} \\ \text{Value} &= \sum_{ij} \langle\mathbf{f}|\mathbf{O}_1|\mathbf{k}_i\rangle M_{ij} \langle\mathbf{k}_j|\mathbf{O}_2|\mathbf{g}\rangle\end{aligned}\quad (16.73)$$

When the auxiliary set of functions is complete, the procedure is an exact identity, but the use of a finite number of functions in practice makes this an approximation, which of course can be controlled by the size of the auxiliary basis set.

The advantage is that insertion of a resolution of the identity can allow the computational problem to be separated into two less complicated problems. An example is the occurrence of three-electron integrals in the R12 methods for including electron correlation (Section 4.11), where insertion of a resolution of the identity allows the three-electron integrals to be written as a product of two-electron integrals, which are significantly easier to handle. The technique is also used in methods such as Hartree–Fock, density functional theory and second-order perturbation theory, where the four-index two-electron integrals can be approximated by three- and two-index integrals instead, leading to significant computational savings.

16.5 Differential Equations

Many of the fundamental equations in physics (and science in general) are formulated as differential equations. Typically, the desired mathematical function is known to obey some relationship in terms of its first and/or second derivatives. The task is to solve this differential equation to find the function itself. A complete treatment of the solution of differential equations is beyond the scope of this book, and only a simplified introduction is given here. Furthermore, we will only discuss solutions of differential equations with one variable. In most cases the physical problem gives rise to a differential equation involving many variables, but prior to solution these can often be (approximately) decoupled by separation of the variables, as discussed in Section 1.6.

16.5.1 Simple first-order differential equations

The simplest case is where the first derivative of the unknown function f is equal to the value of the variable x itself times a constant c .

$$\frac{df}{dx} = cx \quad (16.74)$$

The equation can be solved formally by moving dx to the right-hand side and integrating.

$$\begin{aligned} \int df &= c \int x dx \\ f &= c\left(\frac{1}{2}x^2 + a\right) \end{aligned} \quad (16.75)$$

The integral of df is f itself, while the integral of $x dx$ is $\frac{1}{2}x^2$, except that any constant a can be added. This is completely general: any first-order differential equation will give one additional integration constant that will have to be determined by some other means, for example from knowing the functional value at some point. That the found solution indeed is a solution to the differential equation can be verified by differentiation.

16.5.2 Less simple first-order differential equations

Differential functions where the right-hand side depends only on the variable are relatively simple. A slightly more difficult problem arises when the right-hand side depends on the function itself.

$$\frac{df}{dx} = cf \quad (16.76)$$

The task is now to find a function that upon differentiation gives the same function, except for a multiplicative constant. Formally it can be solved as in eq. (16.77), by separating the variables and integrating.

$$\begin{aligned} \int f^{-1} df &= c \int dx \\ \ln(f) &= cx + a \\ f &= e^{cx+a} = e^a e^{cx} \\ f &= Ae^{cx} \end{aligned} \quad (16.77)$$

The solution is an exponential function where the integration constant can be written as a multiplicative factor A . One may again verify that the solution indeed satisfies the original differential equation by differentiation.

16.5.3 Simple second-order differential equations

A second-order differential equation involves the second derivative of the function.

$$\frac{d^2 f}{dx^2} = cf \quad (16.78)$$

Since the second derivative may be written as two consecutive differentiations, it can formally be solved by applying the above techniques twice. The first integration gives the same solution as in eq. (16.75), with an integration constant a_1 .

$$\begin{aligned} \frac{d^2 f}{dx^2} &= \frac{d}{dx} \left(\frac{df}{dx} \right) = cx \\ \left(\frac{df}{dx} \right) &= c \int x dx \\ \frac{df}{dx} &= \frac{1}{2} cx^2 + a_1 \end{aligned} \quad (16.79)$$

The second integration is now analogous to a simple first-order differential equation.

$$\begin{aligned} \frac{df}{dx} &= \frac{1}{2} cx^2 + a_1 \\ \int df &= \frac{1}{2} c \int x^2 dx + a_1 \int dx \\ f &= \frac{1}{6} cx^3 + a_1 x + a_2 \end{aligned} \quad (16.80)$$

Solving the second-order differential equation produces two integration constants, which must be assigned based on knowledge of the function at two points.

16.5.4 *Less simple second-order differential equations*

Analogously to first-order differential equations, second-order differential equations may have the function itself on the right-hand side, as for example in eq. (16.81).

$$\frac{d^2 f}{dx^2} = cf \quad (16.81)$$

Another example is when the right-hand side involves both the function and its first derivative.

$$\frac{d^2 f}{dx^2} = c_1 f + c_2 \frac{df}{dx} \quad (16.82)$$

The right-hand side may also involve both the unknown function f and another known function of the variable, such as x^2 .

$$\frac{d^2 f}{dx^2} = c_1 f + c_2 x^2 \quad (16.83)$$

Equations of this type are representative of the Schrödinger equation, although in this case it is often written in a slightly different form with the kinetic energy operator plus the potential energy on the left-hand side, and with the c_1 constant written as an energy ϵ .

$$\frac{d^2 f}{dx^2} + cx^2 = \epsilon f \quad (16.84)$$

Eq. (16.84) for example arises for a harmonic oscillator, where the potential energy depends on the square of the variable.

The task in these cases is to find a function that upon differentiation twice gives some combination of the same function, its derivative and variable. Such differential equations cannot be solved by the above “separation of variables” technique. A detailed discussion of how to solve second-order differential equations is beyond the scope of this book, but we will consider two special cases that often arise in computational chemistry.

16.5.5 *Second-order differential equations depending on the function itself*

A second-order differential equation with the function itself on the right-hand side times a positive constant is for example involved in solving the radial part of the Schrödinger equation for the hydrogen atom.

$$\frac{d^2 f}{dx^2} = c^2 f \quad (16.85)$$

For reasons that will become clear shortly, we have written the constant as c^2 , rather than c . By reference to the corresponding first-order equation (eq. (16.77)), we may guess that a possible solution is an exponential function.

$$f = Ae^{cx} \quad (16.86)$$

That this is indeed a solution can be verified by explicit differentiation twice. Recognizing that the corresponding exponential function with a negative argument also is a solution, we can write a more general solution as a linear combination of the two.

$$f = A_1e^{cx} + A_2e^{-cx} \quad (16.87)$$

This contains two constants A_1 and A_2 , as required for a solution to a second-order differential equation, and it is indeed the complete solution. The two integration constants A_1 and A_2 must be assigned based on physical arguments. For the radial part of the hydrogen atom, for example, the A_1 constant is zero since the wave function must be finite for all values of x , and A_2 becomes a normalization constant.

A slightly different situation arises when the differential equation contains a negative constant on the right-hand side, such as that involved in solving the angular part of the Schrödinger equation for the hydrogen atom.

$$\frac{d^2f}{dx^2} = -c^2f \quad (16.88)$$

The solutions are analogous to those above, except for the presence of a factor i in the exponentials.

$$f = A_1e^{icx} + A_2e^{-icx} \quad (16.89)$$

However, since complex exponentials can be combined to give sine and cosine functions, the complete solution can also be written as a linear combination of real functions.

$$f = B_1 \sin(cx) + B_2 \cos(cx) \quad (16.90)$$

The constants A_1/A_2 or B_1/B_2 must again be assigned based on physical arguments.

16.6 Approximating Functions

Although the fundamental mathematical equations describing a physical phenomenon are often very compact, as for example the Schrodinger equation written in operator form $\mathbf{H}\Psi = E\Psi$, their application to all but the simplest model systems usually leads to equations that cannot be solved in analytical form. Even if the equations could be solved, one may only be interested in the solution for a certain limited range of variables. In many cases, it is therefore of interest to obtain an approximate solution, and preferably in a form where the accuracy of the solution can be improved in a systematic fashion. We will here consider three approaches for obtaining such approximate solutions:

- (1) *Taylor expansion*: The real function is approximated by a polynomial that is constructed such that it becomes more and more accurate the closer the variable is to the expansion point. Away from the expansion point, the accuracy can be improved by including more terms in the polynomial.
- (2) *Basis set expansion*: The unknown function is written as a (linear) combination of known functions. The accuracy is determined by the number and mathematical form of the expansion functions. In contrast to a Taylor expansion, which has an error that increases as the variable is removed from the expansion point, a basis set expansion tends to distribute the error over the whole variable range. The error can be reduced by adding more functions in the expansion.
- (3) *Grid representation*: This is similar to expansion in a basis set, except that the known functions are points rather than continuous functions. The accuracy is determined by the number of grid points, and their location.

16.6.1 Taylor expansion

The idea in a *Taylor expansion* is to approximate the unknown function by a polynomial centred at an expansion point x_0 , typically at or near the “centre” of the variable of interest. The coefficients of an N th-order polynomial are determined by requiring that the first N derivatives match those of the unknown function at the expansion point. For a one-dimensional case this can be written as in eq. (16.91).

$$f(x) = f(x_0) + \left. \frac{\partial f}{\partial x} \right|_{x_0} (x - x_0) + \frac{1}{2!} \left. \frac{\partial^2 f}{\partial x^2} \right|_{x_0} (x - x_0)^2 + \frac{1}{3!} \left. \frac{\partial^3 f}{\partial x^3} \right|_{x_0} (x - x_0)^3 + \dots \quad (16.91)$$

For a many-dimensional function, the corresponding second-order expansion can be written as in eq. (16.92).

$$f(\mathbf{x}) = f(\mathbf{x}_0) + \mathbf{g}^t(\mathbf{x} - \mathbf{x}_0) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_0)^t \mathbf{H}(\mathbf{x} - \mathbf{x}_0) + \dots \quad (16.92)$$

Here \mathbf{g}^t is a transposed vector (gradient) containing all the partial first derivatives, and \mathbf{H} is the (Hessian) matrix containing the partial second derivatives. In many cases, the expansion point x_0 is a stationary point for the real function, making the first derivative disappear, and the zeroth-order term can be removed by a shift of the origin.

$$f(x) \approx \frac{1}{2} k_2 (x - x_0)^2 + \frac{1}{6} k_3 (x - x_0)^3 + \frac{1}{24} k_4 (x - x_0)^4 + \dots \quad (16.93)$$

A Taylor expansion is an approximation to the real function by a polynomial terminated at order N . For a given (fixed) N , the Taylor expansion becomes a better approximation as the variable x approaches x_0 . For a fixed point x at a given distance from x_0 , the approximation can be improved by including more terms in the polynomial. Except for the case where the real function is a polynomial, however, the Taylor expansion will always be an approximation. Furthermore, as one moves away from the expansion point, the rate of convergence slows down, i.e. more and more terms are

required to reach a given accuracy. At some point the expansion may become divergent, i.e. even inclusion of all terms up to infinite order does not lead to a well-defined value, and this point is called the *radius of convergence*. It is determined by the distance from the expansion point to the nearest point (which may be in the complex plane) where the function has a singularity. A Taylor expansion of the function $\ln(1+x)$ around $x=0$, for example, has a convergence radius of 1, as the logarithm function is not defined for $x=-1$. Attempting to approximate $\ln(1+x)$ by a Taylor expansion for x -values near -1 or 1 will thus require inclusion of a very large number of terms, and will not converge if $x \geq 1$.

A specific example of a Taylor expansion is the molecular energy as a function of the nuclear coordinates. The real energy function is quite complex, but for describing a stable molecule at sufficiently low temperatures, only the functional form near the equilibrium geometry is required. Terminating the expansion at second order corresponds to modelling the nuclear motion by harmonic vibrations, while higher order terms introduce anharmonic corrections.

For illustration, we will consider the Morse potential in reduced units.

$$y_{\text{Morse}} = (1 - e^{-(x-1)})^2 \quad (16.94)$$

Figure 16.6 shows the second-, third- and fourth-order Taylor approximations to the Morse potential.

Approximating the real function by a second-order polynomial forms the basis for the Newton–Raphson optimization techniques described in Section 12.2.




Figure 16.6 Taylor approximations to the Morse potential

16.6.2 Basis set expansion

An alternative way of modelling an unknown function is to write it as a linear combination of a set of known functions, often called a *basis set*. The basis functions may or may not be orthogonal.

$$f(x) = \sum_{i=1}^M c_i \chi_i(x) \quad (16.95)$$

This corresponds to describing the function f in an M -dimensional space of the basis functions χ . For a fixed basis set size M , only the components of f that lie within this space can be described, and f is therefore approximated. As the size of the basis set M is increased, the approximation becomes better since more and more components of f can be described. If the basis set has the property of being *complete*, the function f can be described to any desired accuracy, provided that a sufficient number of functions are included. The expansion coefficients c_i are often determined either by variational or perturbational methods. For the expansion of the molecular orbitals in a Hartree–Fock wave function, for example, the coefficients are determined by requiring the total energy to be a minimum.

The basis set expansion can be illustrated by using polynomials as basis functions for reproducing the Morse potential in eq. (16.94), i.e. the approximating function is given by eq. (16.96).

$$f(x) = \sum_{i=0}^M a_i (x-1)^i \quad (16.96)$$

The fitting coefficients a_i can be determined by requiring that the integrated difference in a certain range $[a,b]$ is a minimum.

$$\text{ErrF} = \int_a^b (y_{\text{Morse}} - f(x))^2 dx \quad (16.97)$$

Taking the range to be either $[0.5, 2.0]$ or $[0.2, 2.5]$ produces the fits for a second-, third- and fourth-order polynomial shown in Figure 16.7.

Note that the polynomial no longer has the same minimum as the Morse potential but provides a smaller *average* error over the fitting range than the corresponding Taylor polynomial. The Taylor expansion provides an exact fit at the expansion point, which rapidly deteriorates as the variable moves away from the reference point, while the basis set expansion provides a rather uniform accuracy over the fitting range, at the price of sacrificing local accuracy.

16.7 Fourier and Laplace Transformations

Transforming functions between different coordinate systems can often simplify the description. In some cases, it may also be advantageous to switch between different *representations* of a function. A function in real space, for example, can be transformed into a reciprocal space, where the coordinate axes have units of inverse length.




Figure 16.7 Approximations to the Morse potential by expansion in a polynomial basis set; vertical lines indicate the range used for the fitting

Similarly, a function defined in time may be transformed into a representation of inverse time, or frequency. Such interconversions can be done by *Fourier transformations*. The Fourier transform g of a function f is defined as in eq. (16.98).

$$g(k) = \int_{-\infty}^{\infty} f(r) e^{-ikr} dr \quad (16.98)$$

The inverse transformation is given as in eq. (16.99).

$$f(r) = \frac{1}{2\pi} \int_{-\infty}^{\infty} g(k) e^{ikr} dk \quad (16.99)$$

The factor of 2π can also be included as the square root in both the forward and reverse transformation or included in the complex exponent.

While the integral form of the Fourier transform is useful in analytical work, the computational implementation is often done by a discrete representation of the function(s) on a grid, in which case the integrals are replaced by a sum over grid points.

$$g(k_n) = \sum_{n=0}^{N-1} f(r_n) e^{-ikr_n/N} \quad (16.100)$$

In a straightforward implementation of the discrete Fourier transform, the computational time increases as the square of the number of grid points. If, however, the number of grid points is an integer power of two, i.e. $N_{\text{grid}} = 2^m$, the Fourier transform can be done recursively, and is called *Fast Fourier Transform* (FFT). The FFT involves (only) a computational effort proportional to $N_{\text{grid}} \ln N_{\text{grid}}$, which is a substantial reduction relative to the general case of N_{grid}^2 for large values of N_{grid} .

Fourier transforms are often used in connection with periodic functions, for example for evaluating the kinetic energy operator in a density functional calculation where the orbitals are expanded in a plane wave basis.

The *Laplace transform* is defined as in eq. (16.101), where the integral can again be approximated as a finite sum in practical applications.

$$g(k) = \int_0^{\infty} f(r)e^{-kr} dr \quad (16.101)$$

The inverse orbital energy differences in the MP2 method, for example, can be re-written as a sum over the auxiliary variable r , and a sufficient accuracy can often be obtained by including only a few points in the sum.⁴

16.8 Surfaces

A one-dimensional function $f(x)$ can be visualized by plotting the functional value against the variable x . A two-dimensional function $f(x,y)$ can similarly be visualized by plotting the functional value against the variables x and y . However, since plotting devices (paper or an electronic screen) are inherently two-dimensional, the functional value must be plotted in a pseudo-three-dimensional fashion, with the three-dimensional object being imagined by the viewer's brain. Functions with more than two variables cannot readily be visualized.

Functions in computational chemistry typically depend on many variables, often several hundreds, thousands or millions. For analysis purposes, it is possible to visualize the behaviour of such functions in a reduced variable space, i.e. keeping some (most) of the variables constant. Figure 16.8 shows the value of the acrolein LUMO (lowest unoccupied molecular orbital) in a two-dimensional cut 1 Å above the molecular plane.⁵ The magnitude and sign of the orbital is plotted along the third perpendicular dimension.




Figure 16.8 Representation of the acrolein LUMO orbital

An alternative way of visualizing multi-variable functions is to condense or contract some of the variables. An electronic wave function, for example, is a multi-variable function, depending on $3N$ electron coordinates. For an independent-particle model, such as Hartree–Fock or density functional theory, the total (determinantal) wave function is built from N orbitals, each depending on three coordinates.

$$\Phi_{\text{HF/DFT}} = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(1) & \phi_2(1) & \cdots & \phi_N(1) \\ \phi_1(2) & \phi_2(2) & \cdots & \phi_N(2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(N) & \phi_2(N) & \cdots & \phi_N(N) \end{vmatrix} \quad (16.102)$$

The electron density can be obtained by integrating the coordinates for $N - 1$ electrons, giving a function depending on only three coordinates.

$$\rho(x, y, z) = \int \Phi^2(x_1, y_1, z_1, x_2, y_2, \dots, z_N) dx_2 dy_2 \dots dz_N \quad (16.103)$$

Functions of three variables can be visualized by generating a surface in the three-dimensional space corresponding to a constant value of the function, e.g. $\rho(x, y, z) = \text{constant}$. Such surfaces can be plotted, again using the viewer's brain for generating the illusion of a three-dimensional object. The full three-dimensional figure can be visualized by plotting surfaces for different values. Figure 16.9 shows the total electron density of cyclohexane, plotted for decreasing density values. The scales of the seven plots are the same, i.e. the sizes of the surfaces are directly comparable.




Figure 16.9 Surfaces corresponding to the indicated value for the electron density of cyclohexane

The first box corresponds to $\rho = 0.50$, and only the core electrons for the carbon atoms are visible. For $\rho = 0.32$ the hydrogens also appear; and bonds can be recognized for $\rho = 0.20$. Further reduction in the electron density level used for generating the surface obscures the bonding information, and for $\rho = 0.001$ there is little information about the underlying molecular structure left. A surface generated by a constant value of the electron density defines the size and shape of a molecule, but the exact size and shape clearly depend on the value chosen. It can be noted that an isocontour value of 0.001 has often been taken to represent a van der Waals type surface.

More information can be added to surface plots by colour-coding. Orbitals, for example, have a sign associated with the overall shape, which can be visualized by adding two different colors or greyscale shading to the surface corresponding to a constant (numerical) value of the orbital. Figure 16.10 shows the acrolein LUMO in a grey-coded surface representation, which can be compared with the two-dimensional plot shown in Figure 16.8.




Figure 16.10 Grey-coded surface representation of the acrolein LUMO

Other properties have a continuous range of values, not just a sign. An example is the electrostatic potential, which by itself is a function of three coordinates. The combined information of the molecular shape and the value of the electrostatic potential can be visualized by colour-coding the value of the electrostatic potential onto a surface corresponding to a constant value of the electron density. Figure 16.11 shows the electrostatic potential for the acrolein molecule, although the greyscale shading does not provide the level of detail available in a colour-coding.




Figure 16.11 Electrostatic potential superimposed on a surface corresponding to a fixed value of the electron density for acrolein

References

1. G. R. Kneller, *Mol. Sim.*, **7** (1991), 113.
2. O. Christiansen, *J. Chem. Phys.*, **120** (2003), 2149.
3. W. H. Miller, N. C. Handy, J. E. Adams, *J. Chem. Phys.*, **72** (1980), 99.
4. P. Y. Ayala, G. E. Scuseria, *J. Chem. Phys.*, **110** (1999), 3660.
5. G. Schaftenaar, J. H. Noordik, "Molden: a pre- and post-processing program for molecular and electronic structures", *J. Comput.-Aided Mol. Design*, **14** (2000), 123.

17 Statistics and QSAR

17.1 Introduction

An essential component of calculations is to calibrate new methods, and to use the results of calculations to predict or rationalize the outcome of experiments. Both of these types of investigation compare two types of data and the interest is in characterizing how well one set of data can represent or predict the other. Unfortunately, one or both sets of data usually contain “noise”, and it may be difficult to decide whether a poor correlation is due to noisy data or to a fundamental lack of connection. Statistics is a tool for quantifying such relationships. We will start with some philosophical considerations and move into elementary statistical measures, before embarking on more advanced tools.

The connection between reality and the outcome of a calculation can be illustrated as shown in Figure 17.1.




Figure 17.1 Relationship between *Model* and *Reality*

A specific example for “*Reality*” could be the (experimental) atomization energy of a molecule, defined as the energy required to separate a molecule into atoms, which is equivalent to the total binding energy. The atomization energy is closely related to the heat of formation, differing only by the zero point reference state and neglect of vibrational effects. For the atomization energy the zero point for the energy scale is the isolated atoms, while for the heat of formation it is the elements in their most stable form (e.g. H₂ and N₂). Since the dissociation energies for the reference molecules can also be measured, the atomization energy is an experimental observable quantity.

It is important to realize that each element in Figure 17.1 contains errors, and these can be either *systematic* or *random*. A systematic error is one due either to an

inherent bias or to a user-introduced error. A random error is, as the name implies, a non-biased deviation from the “true” result. A systematic error can be removed or reduced, once the source of the error is identified. A random, also sometimes called a *statistical* error, can be reduced by averaging the results of many measurements. Note that random errors can be truly random, for example due to thermal fluctuations or a cosmic ray affecting a detector, but may also be due to many small unrecognized systematic errors adding up to an apparent random noise.

Experimental measurements may contain both systematic and random errors. The latter can be quantified by repeating the experiment a number of times and taking the deviation between these results as a measure for the uncertainty of the (average) result. Systematic errors, however, are difficult to identify. One possibility for detecting these is to measure the same quantity by different methods, or using the same method in different laboratories. The literature is littered with independent investigations reporting conflicting results for a given quantity, each with error bars smaller than the deviation between the results. Such cases clearly indicate that at least one of the experiments contains unrecognized systematic errors.

Theory almost always contains “errors”, but these are called “approximations” in the community. The Hartree–Fock method, for example, systematically underestimates atomization energies since it neglects electron correlation, and the correlation energy is larger for molecules than for atoms. For other properties, the Hartree–Fock method has the same fundamental flaw, neglect of electron correlation, but this may not necessarily lead to systematic errors. For energy barriers for rotation around single bonds, which are differences between two energies for the same molecule with (slightly) different geometries, the contribution from the correlation energy is small, and Hartree–Fock calculations do not systematically over- or underestimate rotational barriers.

The use of a basis set also introduces a systematic error but the direction depends on the specific basis set and the molecule at hand. For a system composed of first row elements (such as C, N, O), the isolated atoms can be completely described with *s*- and *p*-functions at the Hartree–Fock level, but molecules require the addition of higher angular momentum (polarization) functions. Using a basis set containing only *s*- and *p*-functions will systematically underestimate the atomization energy, while a basis set containing few *s*- and *p*-functions but many polarization functions may overestimate the atomization energy. In principle one should choose a balanced basis set, defined as one where the error for the molecule is almost the same as for the atoms, but since the number of basis functions of each kind necessarily is quantized (one cannot have a fractional number of basis functions), this is not rigorously possible, and will depend on the computational level in any case. A very large (complete) basis set will fulfil the balance criteria but is usually impossible in practice. An example of a (systematic) user error is the use of one basis set for the molecule and another for the atoms, as is sometimes done by inexperienced users of electronic structure methods.

The computational implementation of a Hartree–Fock calculation involves choosing a specific algorithm for calculating the integrals and solving the HF equations. In addition, various cutoff parameters are usually implemented for deciding whether to neglect certain integral contributions, and a tolerance is set for deciding when the iterative HF equations are considered to be converged. Since computers perform arithmetic with a finite precision, given by the number of bits chosen to represent a number, this introduces truncation errors, which are of a random nature (roughly as many neg-

ative as positive deviations for a large sample). These random errors can be reduced by increasing the number of bits per number, by using smaller cutoffs and convergence criteria, but can never be completely eliminated. Usually these errors can be controlled, and reduced to a level where they are insignificant compared with the other approximations, such as neglect of electron correlation and the use of a finite basis set.

17.2 Elementary Statistical Measures

Statistics is a tool for characterizing a large amount of data by a few key quantities and it may therefore also be considered as information compression. Consider a data set containing N data points with values x_i ($i = 1, 2, 3, \dots, N$). One important quantity is the *mean* or *average* value, denoted with either a bar or an angle bracket.

$$\bar{x} = \langle x \rangle = \frac{1}{N} \sum_{i=1}^N x_i \quad (17.1)$$

The average is the “middle” point, or the “centre of gravity”, of the data set but it does not tell how wide the data point distribution is. The data sets $\{3.0, 4.0, 5.0, 6.0, 7.0\}$ and $\{4.8, 4.9, 5.0, 5.1, 5.2\}$ have the same average value of 5.0.

In computational chemistry, the mean may depend on an external parameter, such as time. In a molecular dynamics simulation, for example, the average energy (NVT ensemble) or temperature (NVE ensemble) will depend on the simulation time. Indeed, a plot of the average energy or temperature against time can be used as a measure of whether the system is sufficiently (thermal) equilibrated to provide a realistic model.

The *width* or the *spread* of the data set can be characterized by the second moment, the *variance*.

$$\sigma^2 = \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2 \quad (17.2)$$

The “normalization” factor is $N - 1$ when the average is calculated from eq. (17.1); if the exact average is known from another source, the normalization is just N . For large samples the difference between the two is minute and the normalization is often taken as N . The square root of the variance (i.e. σ) is called the *standard deviation*. The above two data sets have standard deviations of 1.6 and 0.2, clearly showing that the first set contains elements further from the mean than the second.

If the distribution of the data is given by a Gaussian function ($\exp(-ax^2)$), then σ determines how large a percentage of the data is within a given range of the mean. Specifically, 68% of the data is within one σ of the mean, 95% is within 2σ and 99.7% is within 3σ . For experimental quantities, the measured result is often given by the notation $\langle x \rangle \pm \sigma$. The σ is loosely called the “error bar”, reflecting the common procedure of drawing a line stretching from $\langle x \rangle - \sigma$ to $\langle x \rangle + \sigma$ in a plot diagram. Note, however, that for a Gaussian data distribution there is a 32% chance that the actual value is outside this interval. Furthermore, for actual data, the distribution is rarely exactly Gaussian. Note also that the standard deviation depends inversely on the square root of the number of data points, i.e. for truly random errors, the standard deviation can be reduced by increasing the number of points.

The third and fourth moments are called the *skew* and *kurtosis*.

$$\begin{aligned}\text{skew} &= \frac{1}{N} \sum_{i=1}^N \left(\frac{x_i - \bar{x}}{\sigma} \right)^3 \\ \text{kurtosis} &= \frac{1}{N} \sum_{i=1}^N \left(\frac{x_i - \bar{x}}{\sigma} \right)^4 - 3\end{aligned}\tag{17.3}$$

These quantities are dimensionless, in contrast to the first and second moments (mean and variance). The skew, kurtosis and corresponding higher moments are rarely used.

The mean and variance are closely connected with the qualifiers *accurate* and *precise*. An accurate measure is one where the mean is close to the real value. A precise measure is one that has a small variance. The goal is an accurate and precise measurement (many data points close to the “real” value). An accurate but imprecise measurement (good mean, large variance) indicates large random and small systematic errors, while a precise but inaccurate measurement (poor mean, small variance) indicates small random but large systematic errors.

Some data sets are (almost) symmetric, such as the digits in the phone book of a large city, containing almost the same number of elements below and above the mean value. Others may be asymmetric, for example containing many points slightly below the mean, but relatively few with much larger values than the mean (for example the income profile for the US population or the Boltzmann energy distribution). Higher order moments such as the skew can be used to characterize such cases. Two alternative quantities can also be used, the *median* and the *mode*. The *median* is the *value in the middle* of the data points, i.e. 50% of the data are below the median and 50% are above. The *mode* is the *most probable element*, i.e. the one that occurs with the highest frequency. In some cases, there may be more than one maximum in the probability distribution, for example a *bimodal* distribution for a probability function with two maxima. For a symmetric distribution, the median and mean are identical, and thus a large difference between these two quantities indicates an asymmetric distribution.

One should be aware that essential information can easily be lost in the data compression of a statistical analysis. European women, for example, have on average 1.5 children, but none have 1.5 children (but 0, 1, 2, 3, . . . children). Such “paradoxes” are at the root of characterizing statistics as “a precise and concise method of communicating half-truths in an inaccurate way”.

17.3 Correlation Between Two Sets of Data

In science, one is often interested in whether one type of data is connected with another type, i.e. whether the data points from one set can be used to predict the other. We will denote such two data sets \mathbf{x} and \mathbf{y} , and ask whether there is a function $f(\mathbf{x})$ that can model the \mathbf{y} data. When the function $f(\mathbf{x})$ is defined or known *a priori*, the question is how well the function can reproduce the \mathbf{y} data.

Two quantities are commonly used for qualifying the “goodness of fit”, the *Root Mean Square* (RMS) deviation and the *Mean Absolute Deviation* (MAD), which for a set of N data points are defined in eq. (17.4).

$$\begin{aligned}\text{RMS} &= \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - f(x_i))^2} \\ \text{MAD} &= \frac{1}{N} \sum_{i=1}^N |y_i - f(x_i)|\end{aligned}\quad (17.4)$$

The MAD represents a uniform weighting of the errors for each data point, while the RMS quantity has a tendency of being dominated by the (few) points with largest deviations. Note that the function $f(\mathbf{x})$ can be very complicated, as for example calculating the atomization energy by the Hartree–Fock method from a set of nuclear coordinates (Figure 17.1).

When the functional form $f(\mathbf{x})$ is unknown, correlation analysis may be used to seek an approximate function connecting the two sets of data. The simplest case corresponds to a linear correlation.

$$y_i = f(x_i) = ax_i + b \quad (17.5)$$

We want to determine the a (slope) and b (intersection) parameters to give the best possible fit, i.e. in a plot of y against x , we seek the best straight line.




Figure 17.2 An approximate linear correlation between x and y

The familiar least squares linear fit arises by defining the “best” line as the one that has the smallest deviation between the actual y_i -points and those derived from eq. (17.5), and taking the error to be the deviation squared. The equations defining a and b can be derived by minimizing (setting the first derivations to zero) an error function.

$$\begin{aligned}\text{ErrF} &= \sum_{i=1}^N w_i (y_i - ax_i - b)^2 \\ \left. \begin{aligned} \frac{\partial \text{ErrF}}{\partial a} &= 0 \\ \frac{\partial \text{ErrF}}{\partial b} &= 0 \end{aligned} \right\} \Rightarrow a, b\end{aligned}\quad (17.6)$$

We note in passing that the minimum number of data points is two, since there are two fitting parameters a and b , i.e. the correlation between any two points can be modelled by a straight line. The data points can be weighted by w_i factors, for example related to the uncertainty with which the y_i data points are measured. The actual equation for a and b can be written in several different ways. One convenient form is by introducing some auxiliary sum-functions.

$$\begin{aligned} S &= \sum_{i=1}^N w_i & S_x &= \sum_{i=1}^N w_i x_i & S_y &= \sum_{i=1}^N w_i y_i \\ S_{xx} &= \sum_{i=1}^N w_i x_i^2 & S_{yy} &= \sum_{i=1}^N w_i y_i^2 & S_{xy} &= \sum_{i=1}^N w_i x_i y_i \end{aligned} \quad (17.7)$$

$$K = (SS_{xx} - S_x^2)^{-1}$$

In terms of these quantities, the optimum a and b parameters are given by eq. (17.8).

$$\begin{aligned} a &= K(SS_{xy} - S_x S_y) \\ b &= K(S_{xx} S_y - S_x S_{xy}) \end{aligned} \quad (17.8)$$

The associated variances are given in eq. (17.9).

$$\begin{aligned} \sigma_a^2 &= KS \\ \sigma_b^2 &= KS_{xx} \end{aligned} \quad (17.9)$$

The “goodness of fit” for such xy -plots is commonly measured by the *correlation coefficient*, r , which is defined in eq. (17.10).

$$r = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}} \quad (17.10)$$

By construction, the correlation coefficient is constrained to the interval $[-1,1]$, where $r = 1$ indicates that all points lie exactly on a line with a positive slope ($a > 0$), $r = -1$ indicates that all points lie exactly on a line with a negative slope ($a < 0$), while $r = 0$ indicates two sets of uncorrelated data. Note that the “correlation coefficient” is often given as r^2 instead, which of course is constrained to the interval $[0,1]$.

The error function in eq. (17.6) is defined by the *vertical* distance, i.e. assuming that the error is located mainly in the y data set. If both data sets have similar errors, the *perpendicular* distance from the data points to the line can be used instead in the error function. This, however, leads to somewhat complicated non-linear equations for the fitting parameters,¹ and is rarely used.

Non-linear correlations (e.g. a quadratic correlation, $f(x) = ax^2 + bx + c$) can be treated completely analogously to the linear case above, by defining an error function and setting the first derivatives with respect to the fitting parameters to zero. Non-linear correlations, however, are used much less than linear ones, for four reasons:

- (1) Many non-linear connections can be linearized by a suitable variable transformation. An exponential dependence, for example, can be made linear by taking the logarithm ($y = ae^{bx}$ becomes $\ln(y) = \ln(a) + bx$).
- (2) Increasing the number of fitting parameters will always produce a better fit, since the fitting function becomes more flexible (a quadratic fitting function has three parameters, while a linear has only two). The data set, however, usually contains noise (random errors), and a more flexible fitting function may simply try to fit the noise rather than improving the fit of the “true” data. For polynomial fitting functions of increasing degree, oscillations of the fitting function between the data points are often seen, which is a clear indication of “overfitting”.
- (3) Any function connecting two sets of data can be Taylor expanded and, to a first approximation, the connection will be linear. All correlations will therefore be linear within a sufficiently “small” interval.
- (4) For functions where the fitting parameters enter in a linear fashion, the equations defining the parameters can be solved analytically. For a function with non-linear parameters, however, the resulting equations must be solved by iterative techniques, with the risk of divergence or convergence to a non-optimal solution.

Points 2–4 suggest that a non-linear fitting function should only be attempted if there are sound theoretical reasons for expecting a non-linear correlation between the data. One such example is the often observed parabolic dependence of the biological activity on the lipophilicity for a series of compounds. Compounds with a low lipophilicity will have difficulty entering the cells and therefore often have a low activity. Compounds with a high lipophilicity, however, will have a tendency of accumulating in the fat cells and therefore also have a low activity. A quadratic dependence with a negative second-order term is therefore expected based on physical arguments.

17.4 Correlation between Many Sets of Data

17.4.1 Multiple-descriptor data sets and quality analysis

In the previous section there were only two sets of data, the \mathbf{y} data, we wanted to model and the variable \mathbf{x} , each being a vector of dimension $N \times 1$. In many cases, however, one has several sets of \mathbf{x} variables ($\mathbf{x1}, \mathbf{x2}, \mathbf{x3}, \dots, \mathbf{xM}$), each of which can potentially describe some of the variation in the \mathbf{y} data set.² There may also be several different sets of \mathbf{y} data that we want to model with the same \mathbf{x} descriptors but for simplicity we will only consider a single set of \mathbf{y} data. The \mathbf{x} variables can be arranged into an $N \times M$ matrix.

$$\mathbf{X} = \begin{bmatrix} x1_1 & x2_1 & x3_1 & \cdots & xM_1 \\ x1_2 & x2_2 & x3_2 & \cdots & xM_2 \\ x1_3 & x2_3 & x3_3 & \cdots & xM_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x1_N & x2_N & x3_N & \cdots & xM_N \end{bmatrix} \quad (17.11)$$

The \mathbf{x} descriptors are often derived from many different sources and may have different units, means and variances. Prior to any correlation analysis, each \mathbf{x} vector is usually

centred to have a mean value of zero (i.e. subtracting the mean value from each vector element), as this eliminates any constant terms in the correlation analysis and focuses on describing the variation in the \mathbf{y} data.

Each \mathbf{x} vector may also be *scaled* with a suitable factor to take into account for example different units for the variables. This, however, is non-trivial and requires careful consideration. A common procedure, which avoids a user decision, is to normalize each \mathbf{x} vector to have a variance of 1, a procedure called *autoscaling*. Autoscaling equalizes the variance of each descriptor and can thus amplify random noise in the sample data and reduce the importance of a variable having a large response and a good correlation with the \mathbf{y} data.

Analogous to the correlation coefficient in eq. (17.10), we want a measure of the quality of fit produced by a given correlation model. Two commonly used quantities are the *Predicted RESidual Sum of Squares* (*PRESS*) and the correlation coefficient R^2 defined by the normalized *PRESS* value and the variance of the \mathbf{y} data (σ_y^2).

$$\begin{aligned} PRESS &= \sum_{i=1}^N (y_i^{\text{actual}} - y_i^{\text{predicted}})^2 \\ \sigma_y^2 &= \frac{1}{N} \sum_{i=1}^N (y_i^{\text{actual}} - \bar{y})^2 \\ R^2 &= 1 - \frac{PRESS}{N\sigma_y^2} \end{aligned} \quad (17.12)$$

R^2 thus measures how well the model reproduces the variation in \mathbf{y} , compared with a model that just predicts the average \mathbf{y} value for all variables.

A straightforward inclusion of more and more \mathbf{x} variables having some relationship with \mathbf{y} in a correlation analysis will necessarily monotonically increase the amount of \mathbf{y} variation described and thus produce an R^2 converging towards 1. Inclusion of variables that primarily serve to describe the noise in the \mathbf{y} data, however, will lead to a model with *less* predictive value for the real variation. This is clearly something that should be avoided but in many cases it is not obvious when the additional components included primarily serve to model the noise in the \mathbf{y} data. To make an unbiased judgement, it is of interest to introduce a quantity that does not measure how well the variables can *fit* the \mathbf{y} data, but one that measures how well the variables can *predict* the \mathbf{y} data. Since we are ultimately interested in predicting \mathbf{y} data from the independent variables, such *cross-validations* are more useful quantities.

One possible measure is to make a correlation analysis using only $N - 1$ data points and ask how well this model predicts the point left out. This can be performed for each of the data points, i.e. a total of N correlation analyses is required. Such “*leave-one-out*” or “*jackknife*” models give an independent measure of the predictive power of a correlation model using a given number of variables. Either the *PRESS* calculated from summing over all the N correlation models or the *predictive* correlation coefficient Q^2 , defined analogously to R^2 in eq. (17.12), can be used to measure the predictive capabilities of the model. Q^2 has in analogy with R^2 a maximum value of 1 but can achieve negative values for models with poor predictive capabilities. A small *PRESS* or a large Q^2 value thus indicates a model with good predictive powers, and models with Q^2 of 0.5–0.6 are usually considered acceptable.

Even for medium-sized data sets the leave-one-out cross-correlation model tends to overestimate the predictive capabilities, as the predicted fraction of the sample is only $1/N$, which rapidly approaches zero for a large sample. Alternative models can be generated by randomly leaving out for example k data points, rather than just one, or by forming subgroups of the data set, and either leaving one point out in each group, or leaving all points in one group out.

17.4.2 Multiple linear regression

For multiple-descriptor data sets, one could use the methods in Section 17.3 to derive a correlation between \mathbf{y} and \mathbf{x}_1 , between \mathbf{y} and \mathbf{x}_2 , between \mathbf{y} and \mathbf{x}_3 , etc., to find the \mathbf{x}_k data set giving the best correlation with \mathbf{y} . It is very likely, however, that one of the other \mathbf{x} variables can describe part of the \mathbf{y} variation, which is not described by \mathbf{x}_k , and a third \mathbf{x} variable describing some of the remaining variation, etc. Since the \mathbf{x} vectors may be internally correlated, however, the second most important \mathbf{x} vector found in a one-to-one correlation is not necessarily the most important once the \mathbf{x}_k vector has been included.

In order to use all the information in the \mathbf{x} variables, a *Multiple Linear Regression* (MLR) of the type indicated in eq. (17.13) can be attempted.

$$\mathbf{y} = \sum_{k=1}^M a_k \mathbf{x}_k \quad (17.13)$$

Note that each data set (\mathbf{y} and \mathbf{x}_k) is a vector containing N data points, and the constant corresponding to b in eq. (17.5) is eliminated if the data are centred with a mean value of zero. Since the expansion coefficients are multiplied directly onto the \mathbf{x} variables, MLR is independent of a possible scaling of the \mathbf{x} data (a scaling just affects the magnitude of the a coefficients but does not change the correlation).

The number of fitting parameters is M , and N must therefore be larger than or equal to M , and in practice one should not attempt fitting unless $N > 5M$, as overfitting is otherwise a strong possibility. The fitting coefficients contained in the \mathbf{a} vector can be obtained from the generalized inverse (Section 16.2) of the \mathbf{X} matrix.

$$\begin{aligned} \mathbf{y} &= \mathbf{X}\mathbf{a} \\ \mathbf{a} &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y} \end{aligned} \quad (17.14)$$

This procedure works fine as long as there are relatively few \mathbf{x} variables that are not internally correlated. In reality, however, it is very likely that some of the \mathbf{x} vectors describe almost the same variation, and in such cases there is a large risk of overfitting the data. This can also be seen from the solution vector in eq. (17.14), the $(\mathbf{X}'\mathbf{X})^{-1}$ matrix has dimension $M \times M$, and will be poorly conditioned (Section 16.2) if the \mathbf{x} vectors are (almost) linearly dependent. Note that the presence of (experimental) noise in the \mathbf{x} data can often mask the linear dependence, and MLR methods are therefore sensitive to noisy \mathbf{x} data.

MLR works best if $N \gg M$ and when the \mathbf{x} data are not internally correlated. If either of these criteria is not fulfilled, one can try to form MLR models by selecting subsets of the descriptors. Searching all possible combinations of descriptors from a total of M data sets rapidly leads to a large number of possibilities, which may be

impossible to search in a systematic fashion. Global optimization methods such as genetic algorithms or simulated annealing (Sections 12.6.3 and 12.6.4) can be used to hunt for the best combination of number and types of descriptors. Such optimizations clearly should focus on maximizing the Q^2 value and not the R^2 value. One may also consider weighting Q^2 with a factor depending (inversely) on the number of components, as a slight increase in Q^2 by including one more components may not be beneficial. Alternatively, and somewhat more systematically, one of the methods in the next section can be used.

17.4.3 Principal component and partial least squares analysis

Multiple linear regression cannot easily handle situations where $M > N$ or when the \mathbf{x} variables are (almost) linearly correlated. These problems can be removed by introducing a modified set of descriptors, so-called *latent variables*. The idea is to extract linear combinations of the \mathbf{x} variables that are orthogonal and ranked according to their variation, and only use a limited set of these variables for performing the correlation with the \mathbf{y} variables.

The eigenvalues of the $\mathbf{X}^t\mathbf{X}$ matrix (eq. (17.11)) contain information on the correlation between the \mathbf{x} variables: an eigenvalue of zero indicates that one column can be written as a linear combination of the other columns, and a small non-zero value indicates that one column contains almost redundant information. The eigenvector corresponding to the largest eigenvalue contains the linear combination of \mathbf{x} descriptors having the largest variation in the \mathbf{x} data, the eigenvector corresponding to the second largest eigenvalue has the second largest variations, etc. Furthermore, since the eigenvectors are orthogonal, different eigenvectors describe different parts of the variation.

Using the concepts in Chapter 16, the original \mathbf{x} vectors can be thought of as being the (non-orthogonal and unnormalized) basis vectors of an M -dimensional coordinate system. The eigenvectors of the $\mathbf{X}^t\mathbf{X}$ matrix describe the same fundamental coordinate system but with (orthonormal) basis vectors that have been rotated with respect to the original \mathbf{x} vectors. We note that if there are eigenvalues close to zero then the effective dimension of the coordinate system is less than M . Figure 17.3 shows a two-dimensional example where the points for the two non-orthogonal \mathbf{x} descriptors have an internal correlation and display a similar variation along the two directions. The new \mathbf{z} variables are orthogonal and most of the variation is now located in the \mathbf{z}_1 variable, while \mathbf{z}_2 describes a much smaller variation.




Figure 17.3 Illustrating the relationship between the original (\mathbf{x}) and latent (\mathbf{z}) variables

If we denote the rotated basis vectors ($\mathbf{X}'\mathbf{X}$ eigenvectors) with \mathbf{z} , we can write the connection (neglecting normalization) as in eq. (17.15).

$$\begin{aligned}\mathbf{z}_j &= \sum_{k=1}^M \mathbf{x}_k u_{jk} \\ \mathbf{Z} &= \mathbf{XU}\end{aligned}\tag{17.15}$$

Since the $\mathbf{X}'\mathbf{X}$ matrix depends on the relative magnitudes of the individual \mathbf{x} vectors, the \mathbf{z} vectors depend on a possible scaling of the original \mathbf{x} descriptors.

The idea in a *Principal Component Analysis* (PCA) is to use the \mathbf{z} variables as the descriptive variables. If all M eigenvectors are used the result is identical to using the original \mathbf{x} variables (i.e. MLR). However, the premise of the PCA method is to include only a few (J) \mathbf{z} variables, and to select them according to their eigenvalues. A series of multiple linear regressions are done using more and more eigenvectors, first \mathbf{z}_1 , then \mathbf{z}_1 and \mathbf{z}_2 , then including also \mathbf{z}_3 , etc. At each stage the predictive capabilities of the model are calculated, for example quantified by Q^2 . If the original \mathbf{x} data have a reasonable correlation with the \mathbf{y} data, then a plot of Q^2 against the number of variables included will typically display an initial steep increase, but then level off or even start to decrease slightly as the number of latent variables is increased. The point where Q^2 levels off indicates that the optimal number of components has been reached, i.e. at this point the predictive power of the model cannot be increased further by including more components.

The main problem with the PCA method is that some of the \mathbf{x} variables may not be particularly good at describing the \mathbf{y} variables, i.e. the first few PCA vectors describing the largest variation among the \mathbf{x} variables may correlate poorly with the variation in the \mathbf{y} data. In such cases, a global optimization search can be made for a model based on a relatively small number of components selected from all the PCA vectors with eigenvalues above a suitable cutoff.

The *Partial Least Squares* (PLS, also sometimes called *Projection to Latent Structures*) method attempts to improve on the selection of the latent variables by weighting the \mathbf{X} matrix with the \mathbf{y} vector prior to diagonalization, i.e. diagonalizing the $\mathbf{X}'\mathbf{y}\mathbf{y}'\mathbf{X}$ (equivalent to $(\mathbf{y}'\mathbf{X})'(\mathbf{y}'\mathbf{X})$) matrix instead of $\mathbf{X}'\mathbf{X}$. This ensures that the eigenvectors with the largest eigenvalues will be biased towards describing the variation in \mathbf{y} . The only difference between PCA and PLS is thus in how the latent \mathbf{z} variables are generated, either by diagonalization of the $\mathbf{X}'\mathbf{X}$ matrix, or from the corresponding \mathbf{y} -weighted matrix. The PLS latent variables will naturally be ordered according to their ability to describe the \mathbf{y} variation, alleviating the necessity for performing a combinatorial search for which latent vectors to use in the regression. For optimal cases, a plot of Q^2 against the number of PLS components will rapidly reach a maximum and provide a compact model with good predictive capabilities.

A disadvantage of the PLS method is the inherent bias towards selecting latent variables describing noise in the \mathbf{y} data, i.e. \mathbf{x} variables that only have a small internal variation but that correlate with the noise in the \mathbf{y} data are selected as important. For this reason, \mathbf{x} variables with small internal variance over the \mathbf{y} data points are often removed from the descriptor data set prior to performing the PLS analysis. This pre-selection procedure, however, requires user involvement and it is not always easy to decide which variables to remove. Unfortunately, the predictive capabilities of a PLS

model are often sensitive to elimination of one or more \mathbf{x} variables. A global optimization scheme may again be employed in such cases, i.e. performing a search for which \mathbf{x} components to *remove* from the PLS analysis in order to provide a model with a high Q^2 value.

17.4.4 Illustrative example

Consider the \mathbf{X} and \mathbf{y} matrices in eq. (17.16), where variables have been centred to give a mean of zero.

$$\mathbf{X} = \begin{pmatrix} 1.0 & 1.0 \\ -1.0 & -1.0 \\ 0.1 & -0.1 \\ -0.1 & 0.1 \end{pmatrix} \quad \mathbf{y} = \begin{pmatrix} 0.1 \\ 0.1 \\ 0.9 \\ -1.1 \end{pmatrix} \quad (17.16)$$

Clearly the first two rows show a large variation in x with no change in y , i.e. these variables are not related to the response. The last two rows display correlation and anti-correlation, respectively, with the y data in equal amounts. Solving with MLR equation (17.14) gives the solution vector \mathbf{a} in eq. (17.17), which shows that both \mathbf{x} columns are equally important in describing the \mathbf{y} variation. The difference between the actual and predicted \mathbf{y} indicates that there is a residual \mathbf{y} variation that cannot be modelled by the \mathbf{x} variables.

$$\mathbf{a} = \begin{pmatrix} 5 \\ -5 \end{pmatrix} \quad \mathbf{y}_{\text{pred}} = \begin{pmatrix} 0.0 \\ 0.0 \\ 1.0 \\ -1.0 \end{pmatrix} \quad \mathbf{y}_{\text{res}} = \begin{pmatrix} 0.1 \\ 0.1 \\ -0.1 \\ -0.1 \end{pmatrix} \quad (17.17)$$

Diagonalizing the $\mathbf{X}^t\mathbf{X}$ matrix for the PCR analysis gives eigenvalues of 4.00 and 0.04, with corresponding (unnormalized) eigenvectors (1,1) and (1,-1). In the transformed coordinate system, the principal components \mathbf{z} (eq. (17.15)) are given in eq. (17.18).

$$\mathbf{Z} = \begin{pmatrix} 1.0 & 0.0 \\ -1.0 & 0.0 \\ 0.0 & 0.1 \\ 0.0 & -0.1 \end{pmatrix} \quad (17.18)$$

Clearly the eigenvector corresponding to the largest eigenvalue (the first principal component, i.e. the first column in \mathbf{Z}) has a zero overlap with the y data, while the second eigenvector accounts for all the important variation in the original \mathbf{x} variables. The PLS components arising from diagonalization of the $\mathbf{X}^t\mathbf{y}\mathbf{y}^t\mathbf{X}$, on the other hand, have eigenvalues of 0.0 and 0.08 with corresponding (unnormalized) eigenvectors (1,1) and (1,-1). The transformed \mathbf{X} matrix is identical to eq. (17.18), except that it is now the *second* column that corresponds to the *largest* eigenvalue. The direction that has the largest eigenvalue in the PCR case has an eigenvalue of zero in the PLS case, showing that it contains no information of the y variation. In both the PCR and PLS cases, the y variation that can be described is contained in only one component, but

the PLS procedure identifies the most important component as the one with the largest eigenvalue. The predicted \mathbf{y} is identical to the MLR results (eq. (17.17)) for both PCR and PLS; the only difference is that only one component is required, rather than two.

17.5 Quantitative Structure–Activity Relationships (QSAR)

One important application of PCA and PLS techniques is in the lead development and optimization of new drugs in the pharmaceutical industry. From the basic chemical and physical theories, it is clear that there must be a connection between the structure of a molecule and its biological activity. It is also clear, however, that the connection is very complex, and it is very unlikely that the biological activity can ever be predicted completely *a priori*. A drug taken orally, for example, must first survive the digestive enzymes and acidic conditions in the stomach, cross over into the blood stream, possibly also cross over into the brain, diffuse to the target enzyme without binding to other enzymes, bind the target enzyme at a specific site and with a large binding constant, and finally have a suitable half-life in the organism before being degraded into non-toxic components. Each of these quantities depends on different parts of the molecular structure and the combined effect is therefore very difficult to predict. Each quantity, however, may possibly be correlated with (different) molecular properties, but adequate data for each effect is rarely available.

In the initial stages of developing a new drug, the focus is usually on the binding to the target enzyme and having a suitable lipophilicity, the latter ensuring a reasonable transfer rate for crossing between the blood and cells. A lead compound is somehow determined, more often than not by serendipity. From this lead, a small initial pool of compounds is synthesized and their biological activity, of which the binding constant to the target enzyme is an important quantity, is measured by a suitable biological assay. At this point, statistical methods are often used to correlate the molecular structure and properties to the observed activity, a *Quantitative Structure–Activity Relationship* (QSAR), thereby providing a tool for “guessing” which new modified compounds should be tried next.

In a traditional QSAR study, a variety of molecular descriptors are included. Typically, these include a measure of the lipophilicity (often taken as the partition coefficient between water and 1-octanol), electronic and steric substituent parameters (Hammett³ and Taft constants), and $pK_{a/b}$ values for acidic and/or basic groups. These are rather obvious molecular descriptors, but many other less obvious descriptors have also been used, such as the molecular weight, IR frequencies, dipole moments, NMR chemical shifts, etc. The philosophy is to include, rather indiscriminately, as many descriptors as can easily be generated and then let the statistical method sort out which of these are actually important. With many of the descriptors having only remote connection with the measured activity, classical or MLR correlations are clearly not suitable methods. PCA or PLS methods are better at detecting poor descriptors and dealing with the fact that the measured biological activities often have rather large uncertainties.

Classical QSAR methods focus on correlating experimental activities with experimental descriptors. This allows an identification of important structural features, such as having a pK_a value close to 5 or having an electron-withdrawing group in a *para*-position of a phenyl ring, thereby limiting the field of possible new compounds to

prepare and test. More recently, the focus has been on “virtual” (*in silico*) screening, i.e. correlating experimental activities with theoretical descriptors. If a good QSAR model can be constructed from such data, this allows prediction of the biological activities of molecules that exist only in the computer. The activity of many thousands of (possibly computer-generated) structures can thus be predicted, and only those that are predicted to be reasonably active need to be synthesized. The theoretical descriptors can be similar to those in traditional QSAR methods, for example replacing the experimental water–octanol partition coefficient with a corresponding theoretical estimate, etc. The so-called 3D-QSAR models, however, are more representative of these new QSAR methods, and the *Comparative Molecular Field Analysis* (COMFA) method is probably one of the most popular of these.⁴

In the COMFA method, the molecular descriptors are taken as steric and electric fields calculated at a large number of points surrounding each molecule. The molecule is placed in a “box” and a few thousand points are selected between the surface of the molecule and a few van der Waals radii outwards. At each of these points, the steric repulsion or attraction from a probe atom (typically a carbon atom) is calculated by a force field van der Waals term. The electric attraction or repulsion is similarly calculated by placing a point charge with magnitude +1 in each point. The complete set of molecular descriptors thus consists of a few thousand data points, representing the steric and electric interaction of the molecule with other (possible) atoms in the near vicinity. Clearly these data are highly redundant; the value in a certain point will be very close to the average of the neighbouring points. Deriving QSAR models with such large sets of data descriptors is only possible using PCA and PLS methods. Such 3D-QSAR methods are primarily used when the structure or identity of the receptor enzyme is unknown. If the enzyme and binding site is known from an X-ray structure, the testing of possible drug candidates can be done by docking methods, as described in Section 12.7.

The main problem with 3D-QSAR methods such as COMFA is the alignment of the molecules in the test set.⁵ First, it must be assumed that each of the molecules binds to the enzyme at the same site. Second, it is not always clear that all the molecules bind to the active site in the same overall orientation. Third, if the molecule has several conformations available, one has to guess or estimate which conformation actually binds to the enzyme. Furthermore, even if the overall orientation of all the molecules is assumed to be the same, the specific alignment is not unambiguous. Even for molecules having the same major structural features, one can choose either to align on the best RMS fit of all atoms, on only the non-hydrogen atoms, or on only the atoms in a substructure of the molecule (e.g. a phenyl ring). Figure 17.4 illustrates the ambiguity for aligning the compound on the right with that on the left, should the imine or nitro group be used for alignment?




Figure 17.4 Illustrating the alignment problem in COMFA methods

When the molecular structure of the compounds in the training set has few or no common elements, the alignment may instead be done based on for example the electric moments (dipole, quadrupole, etc.) or on the electrostatic potential on a suitable molecular surface.

Since the alignment of the molecules influences the values calculated at the steric and electric grid points, this is a feature that influences the statistical correlation. If a successful QSAR model can be obtained from such data, however, the model will provide information on which of the grid points are important in a steric and electric sense. Analysis of such data provides a virtual mapping of the receptor, i.e. identifying regions of the active site where the drug candidates should have large/small groups, and where they should have positively/negatively charged groups.

References

1. D. York, *Can. J. Phys.*, **44** (1966), 1079.
2. G. R. Famini, L. Y. Wilson, *Rev. Comp. Chem.*, **18** (2002), 211.
3. C. Hansch, A. Leo, R. W. Taft, *Chem. Rev.*, **91** (1991), 165.
4. R. D. Cramer, D. E. Patterson, J. D. Bunce, *J. Am. Chem. Soc.*, **110** (1989), 5959.
5. F. Melani, P. Gratteri, M. Adamo, C. Bonaccini, *J. Med. Chem.*, **46** (2004), 1359.

18 Concluding Remarks

The real world is very complex and a complete description is therefore also very complicated. Only by imposing a series of often quite stringent limitations and approximations can a problem be reduced in complexity such that it may be analyzed in some detail, as for example by calculations. A chemical reaction in the laboratory may involve perhaps 10^{20} molecules surrounded by 10^{24} solvent molecules, in contact with a glass surface and interacting with gases (N_2 , O_2 , CO_2 , H_2O , etc.) in the atmosphere. The whole system will be exposed to a flux of photons of different frequency (light) and a magnetic field (from the earth), and possibly also a temperature gradient from external heating. The dynamics of all the particles (nuclei and electrons) is determined by relativistic quantum mechanics, and the interaction between particles is governed by quantum electrodynamics. In principle the gravitational and strong (nuclear) forces should also be considered. For chemical reactions in biological systems, the number of different chemical components will be large, involving various ions and assemblies of molecules behaving intermediately between solution and solid state (e.g. lipids in cell walls).

Except for a couple of rather extreme areas (such as the combination of general relativity and quantum mechanics, or the unification of the strong and gravitational forces with the electroweak interaction), we believe that all the fundamental physics is known. The “only” problem is that the real world contains so many (different) components interacting by complicated potentials that a detailed description is impossible.

As this book hopefully has given some insight into, the key is to know what to neglect or approximate when trying to obtain answers to specific questions in pre-defined systems. For chemical problems, only the electrostatic force needs to be considered; the gravitational interaction is a factor of 10^{39} weaker and can be completely neglected. Similarly, the strong nuclear force is so short-ranged that it has no effect on chemical phenomena (although the brief claims regarding “cold” fusion for a period seemed to contradict this). The weak interaction, which is responsible for radioactive decay by the $n \rightarrow p + e$ process, is also much smaller than the electrostatic, although there have been various estimates of whether it might give rise to a symmetry breaking (i.e. preference for one enantiomer over its mirror image) which possibly could be detected

experimentally. Similarly, the earth's magnetic field is so tiny that only under very special circumstances can it have any detectable influence on the outcome of a chemical reaction.

For electronic structure calculations within the wave function approach, the starting point is usually an independent-particle model, which for electrons is the Hartree–Fock model. The results from this model can be improved by adding electron correlation corrections and increasing the basis set. The resulting two-dimensional diagram shown in Figure 4.3 indicates that the “exact” result can be obtained by systematically increasing the level of sophistication along both axes until convergence is reached. Usually the desired level of accuracy is such that the convergence cannot be reached owing to limitations in computational resources, and the results thus suffer from approximations in the one-particle (basis set) and many-particle (configurations) space. Even if the residual errors can be reduced below the target accuracy, the “exact” solution is still subject to a number of limitations that must also be considered in order to compare with the experimental results. These may for example be:

- (1) Neglect of relativistic effects, by using the Schrödinger instead of the Dirac equation. This is reasonably justified in the upper part of the periodic table but not in the lower half. For some phenomena, such as spin–orbit coupling, there is no classical counterpart and only a relativistic treatment can provide an understanding. The relativistic effects may be incorporated by a one-component (mass–velocity and Darwin terms), two-component (spin–orbit) or full four-component methods (Figure 8.2).
- (2) The effects of the environment, such as solvent effects. These may be modelled for example by a continuum model, by treating the solvent as an ensemble of classical particles (QM/MM methods), or by including them in the quantum description (e.g. Car–Parrinello methods).
- (3) Vibrational corrections. For energies, this would typically be inclusion of zero point energies while for properties this may correspond to a vibrational averaging. The corrections may again be done at several levels of accuracy, for example using a harmonic approximation or also including anharmonic effects.
- (4) Finite temperature effects. This would correspond for example to a molecular structure not being represented as a fixed geometry but rather as an ensemble of structures corresponding to an average over accessible geometries at a given temperature.
- (5) Non-Born–Oppenheimer effects. The assumption of a rigorous separation of nuclear and electronic motions is in most cases a quite good approximation and there is a good understanding of when it will fail. Methods for going beyond the Born–Oppenheimer approximation are still somewhat limited in term of generality and applicability.
- (6) Quantum effects for the nuclei. One may argue that the vibrational effects are the most important of these, but in some cases other effects such as tunnelling may also be important.
- (7) Quantum mechanics being replaced (wholly or partly) by classical mechanics. For electrons such an approximation would lead to disastrous results, but for nuclei (atoms) the quantum effects are sufficiently small that they in most cases can be neglected.

- (8) Approximating the intermolecular interactions to only include two-body effects, for example electrostatic forces are only calculated between pairs of fixed atomic charges in force field techniques. Or the discrete interactions between molecules may be treated only in an average fashion, by using Langevin dynamics instead of molecular dynamics.
- (9) Calculating ensemble or time averages over a relatively small number of points (perhaps a few million) and a limited number of particles (perhaps a few hundred), instead of something that approach the macroscopic sample of perhaps 10^{20} molecules/configurations.
- (10) Finite temperature being reduced to zero kelvin, i.e. the use of static structures to represent molecules, rather than treating them as an ensemble of molecules in a distribution of states (translational, rotational and vibrational) corresponding to a (macroscopic) temperature.
- (11) Making approximations in the Hamiltonian describing the system, for example semi-empirical electronic structure methods.
- (12) Approximating external fields (electric or magnetic) by only considering their linear components. For normal conditions, this will be a quite good approximation, but this is not the case in for example intense laser fields.
- (13) Treating the nuclei as point particles. In reality a nucleus has a finite size ($\sim 10^{-15}$ m) and since the electrons can penetrate the nucleus the potential felt inside the nucleus will differ from that of a point particle, and consequently lead to changes in the electronic energy.
- (14) QED corrections. The interaction between charged particles is normally described by the Coulomb interaction, but when the quantization of the field is considered, there are additional higher order correction terms.

Most of these approximations are mainly of a computational nature, as there are well-defined methods available for going beyond the approximations, but they are computationally too demanding. The key is therefore to be able to evaluate what level of theory (i.e. which approximations are appropriate) is required for obtaining results that are sufficiently accurate to provide useful information about the question at hand. Hopefully this book has given a few clues as how to select a suitable method for a given problem.

Appendix A

Notation

Bold quantities are operators, vectors, matrices or tensors. Plain symbols are scalars.

α	Polarizability
$\alpha\beta$	Spin functions
$\alpha\beta$	Dirac 4×4 spin matrices
$\alpha\beta\gamma\delta$	Summation indices for basis functions
$\alpha, \beta, \gamma, \delta, \zeta$	Basis function exponents
α_A, β_{AB}	Hückel parameters for atom A, and between atoms A and B
a	Born radius for solvation cavities
$abcd$	Summation indices for virtual MOs
$a_n, a_i, b_i, c_i, \dots$	Expansion coefficients
a	Acceleration
A	Helmholtz free energy
A	Antisymmetrizing operator
A	Vector potential
A	Hyperfine coupling constant
β	First hyperpolarizability
β_μ	Resonance parameter in semi-empirical theory
B	Magnetic field (magnetic induction)
$\chi, \mu, \nu, \lambda, \sigma$	Basis functions (atomic orbitals), <i>ab initio</i> or semi-empirical methods
χ	Electronegativity
χ^B	Out-of-plane angle for atom B
χ	Magnetic susceptibility
X	Gauge including basis function
c	Speed of light
$c_{\alpha i}$	MO expansion coefficients
δ	An infinitesimal variation or quantity

δ_{ij}	Kronecker delta ($\delta_{ij} = 1$ for $i = j$, $\delta_{ij} = 0$ for $i \neq j$)
$\delta(\mathbf{r})$	Dirac delta function ($\delta(\mathbf{r}) = 0$ for $\mathbf{r} \neq 0$)
Δ	A finite difference or quantity
d	Distance
d	Deexcitation operator
D	Dissociation energy
D	Density matrix
$D_{\alpha\beta}$	Density matrix element in AO basis
ε	Matrix eigenvalue
ε	van der Waals parameter
ε	Dielectric constant
ε	Energy, for one electron or as an individual term in a sum
ε	Energy matrix
e	Excitation operator
E	Energy, many particles or terms
E_e	Electronic energy
$E[\rho]$	Energy functional
EA	Electron affinity
ϕ	Molecular orbital
ϕ	Electrostatic potential
Φ	Slater determinant or similar approximate wave function
f_i	Gradient component along a Hessian eigenvector
F	Electric field
F	Force constant matrix
F	Fock operator or Fock matrix
$F_{ij}, F_{\alpha\beta}$	Fock matrix element in MO and AO basis
γ	Second hyperpolarizability
γ_k	Reduced density matrix of order k
g_e	Electronic g-factor
g_A	Nuclear g-factor
g	Two-electron repulsion operator
g	Gradient (first derivative)
G	Gibbs free energy
G_{xy}	Coulomb type matrix elements in semi-empirical theory ($x, y = s, p, d$)
G	Matrix containing square root of inverse atomic masses
η	An infinitesimal scalar
η	Absolute hardness
h	Planck's constant
\hbar	$h/2\pi$
h	Core or other effective one-electron operator
$h_{ij}, h_{\alpha\beta}$	Matrix element of a one-electron operator in MO and AO basis
$h_{\mu\nu}$	Matrix element of a one-electron operator in semi-empirical theory
h, h₁, h₂, . . .	Excitation and deexcitation operators
H	Enthalpy
H	Hessian (second derivative matrix of a function)

H, H_e, H_n	Hamiltonian operator or Hamiltonian matrix (general, electronic, nuclear)
H_{ij}	Matrix element of a Hamiltonian operator between Slater determinants
H_{xy}	Exchange type matrix elements in semi-empirical theory ($x, y = s, p, d$)
$ijkl$	Summation indices for occupied MOs
I	Moment of inertia
I	Unit matrix
I	Nuclear magnetic moment or spin
IP, I_μ	Ionization potential
J	Spin–spin coupling matrix
J	Coulomb operator
J_{ij}	Coulomb integral
$J[\rho]$	Coulomb functional
κ	Lagrange multiplier
κ	Transmission coefficient
κ	Compressibility constant
k	Boltzmann's constant
k	Rate constant
$k, k^{\text{AB...}}$	Force constant (for atoms A, B, . . .)
K	Anharmonic constants (third derivative)
K	Exchange operator
K_{ij}	Exchange integral
$K[\rho]$	Exchange functional
λ	Lagrange multiplier
λ	General perturbation strength
λ	Hessian shift parameter
l, L	Angular momentum quantum number
L	Lagrange function
l, L	Angular momentum operator
μ	Mulliken electronegativity, chemical potential
μ	Reduced mass
μ	Dipole moment
μ_0	Vacuum permeability
μ_B	Bohr magneton
μ_N	Nuclear magneton
m	Mass, general or electron mass
M_A	Nuclear mass
ν	Vibrational frequency
n_i	Orbital occupation number
N	Number of particles
N_A	Avogadro's number
O, P, Q, R, S	General operators
π	Generalized momentum operator
Π	Product of diagonal elements in a Slater determinant
p	Momentum operator or vector

P	Pressure
P_i	Legendre polynomial
$\mathbf{P}, \mathbf{P}_{ij}$	Permutation operators (permuting indices i and j)
$\mathbf{P}_1, \mathbf{P}_2$	Perturbation operators (one- and two-electron)
q	Charge on a particle (integer)
q	Partition function (one particle)
Q	Atomic charge (can be fractional), fitted or from population analysis
Q	Partition function (many particles)
\mathbf{q}	Normal or generalized coordinate
Q	Predictive correlation coefficient
\mathbf{Q}	Quadrupole moment
ρ	Electron density
ρ	Bond order
\mathbf{r}	Position vector(s), general or electronic
r_{ij}	Distance between electrons i and j
R	Trust radius
R	Gas constant
r, R	Correlation coefficient
\mathbf{R}	Position vector, nuclear
R_{ij}, R_{AB}, R^{AB}	Distance between atoms or nuclei, i and j or A and B
r, θ, ϕ	Polar coordinates
σ	Order of rotational subgroup
σ	Charge density
σ^2	Variance
$\sigma_{x,y,z}$	Pauli 2×2 spin matrices
\mathbf{s}	Electron spin operator
S	Entropy
S	Switching function
\mathbf{S}^2	Spin squared operator
$S_{\alpha\beta}$	Overlap matrix element in AO basis
$\theta(t)$	Heaviside step function ($\theta(t) = 0$ for $t < 0$, $\theta(t) = 1$ for $t > 0$)
θ^{ABC}	Angle between atoms A, B and C
$\Theta_{s,i}^N$	Spin coupling function
Δt	Small (finite) time step
t	Time
\mathbf{t}	Translational vector
τ	Heat or pressure bath coupling parameter
τ	Phase factor
τ	Imaginary time variable
τ	Orbital kinetic energy density
t_i, t_{ij}	Cluster amplitudes
T	Temperature
$\mathbf{T}, \mathbf{T}_1, \mathbf{T}_2, \dots$	Cluster operator (general, single, double, . . . excitations)
$\mathbf{T}, \mathbf{T}_e, \mathbf{T}_n$	Kinetic energy operator (general, electronic, nuclear)
$T[\rho]$	Kinetic energy functional, exact
$T_s[\rho]$	Kinetic energy functional, calculated from a Slater determinant

U	Internal energy
\mathbf{U}	Unitary matrix
U_i	Matrix element of a semi-empirical one-electron operator, usually parameterized
\mathbf{v}	Velocity
V	Volume
V, V^{AB}	Potential energy (between atoms A and B)
V_{ij}	Coulomb potential between particles i and j
$V[\rho]$	Potential energy functional
$\mathbf{V}, \mathbf{V}_{\text{ee}}, \mathbf{V}_{\text{ne}}, \mathbf{V}_{\text{nn}}$	Potential (Coulomb) energy operator (general, electron–electron, nuclear–electron, nuclear–nuclear)
ω	Frequency associated with an electric or magnetic field
ω	Harmonic vibrational frequency
ω^{ABCD}	Torsional angle between atoms A, B, C and D
Ω	Two-electron operator
W	Energy of an approximate wave function
W_i	Perturbation energy correction at order i
W_k	Wigner intracule of order k
\mathbf{W}	Energy-weighted density matrix
$W_{\alpha\beta}$	Energy-weighted density matrix element in AO basis
$W_{\sigma\beta}^*$	Weighting factor in pseudospectral methods
ξ	Magnetizability
x_i, y_i, z_i	Cartesian coordinates for particle i
Δx_i	Component in a vector
$\Psi, \Psi_{\text{e}}, \Psi_{\text{n}}$	Exact or multi-determinant wave function (general, electronic, nuclear)
ζ	Spin polarization
ζ	Friction coefficient
ζ	Molecular surface parameter for calculating solvation energies
Z	Nuclear charge, exact
Z'	Nuclear charge, reduced by the number of core electrons
$\langle n $	Bra, referring to a function characterized by quantum number n .
$ n\rangle$	Ket, referring to a function characterized by quantum number n .
$\langle n \mathbf{O} m\rangle$	Bracket (matrix element) of operator \mathbf{O} between functions n and m
$\langle \mathbf{O} \rangle$	Average value of \mathbf{O}
$ \mathbf{O} $	Norm or determinant of \mathbf{O}
$\langle\langle \mathbf{P}; \mathbf{Q} \rangle\rangle$	Propagator of \mathbf{P} and \mathbf{Q}
$[\mathbf{P}, \mathbf{Q}]$	Commutator of \mathbf{P} and \mathbf{Q} ($[\mathbf{P}, \mathbf{Q}] = \mathbf{PQ} - \mathbf{QP}$)
∇	Gradient operator
∇^2	Laplace operator
\cdot	Entrywise matrix product or tensor contraction
\times	Cross product
$\nabla \cdot$	Divergence operator
$\nabla \times$	Curl operator
t	Vector transposition
\dagger	Complex conjugate

Appendix B

B.1 The Variational Principle

The *Variational Principle* states that an approximate wave function has an energy that is above or equal to the exact energy. The equality holds only if the wave function is exact. The proof is as follows.

Assume that we know the exact solutions to the Schrödinger equation.

$$\mathbf{H}\Psi_i = E_i\Psi_i \quad i = 0, 1, 2, \dots, \infty \quad (\text{B.1})$$

There are infinitely many solutions and we assume that they are labelled according to their energies, E_0 being the lowest. Since the \mathbf{H} operator is Hermitian, the solutions form a *complete* basis. We may furthermore choose the solutions to be orthogonal and normalized.

$$\langle \Psi_i | \Psi_j \rangle = \delta_{ij} \quad (\text{B.2})$$

An approximate wave function can be expanded in the exact solutions, since they form a complete set.

$$\Phi = \sum_{i=0}^{\infty} a_i \Psi_i \quad (\text{B.3})$$

The energy of an approximate wave function is calculated as in eq. (B.4).

$$W = \frac{\langle \Phi | \mathbf{H} | \Phi \rangle}{\langle \Phi | \Phi \rangle} \quad (\text{B.4})$$

Inserting the expansion (B.3) we obtain eq. (B.5).

$$W = \frac{\sum_{i=0}^{\infty} \sum_{j=0}^{\infty} a_i a_j \langle \Psi_i | \mathbf{H} | \Psi_j \rangle}{\sum_{i=0}^{\infty} \sum_{j=0}^{\infty} a_i a_j \langle \Psi_i | \Psi_j \rangle} \quad (\text{B.5})$$

Using the fact that $\mathbf{H}\Psi_i = E_i\Psi_i$ and the orthonormality of the Ψ_i (eqs (B.1) and (B.2)), we obtain eq. (B.6).

$$W = \frac{\sum_{i=0}^{\infty} a_i^2 E_i}{\sum_{i=0}^{\infty} a_i^2} \quad (\text{B.6})$$

The variational principle states that $W \geq E_0$ or, equivalently, $W - E_0 \geq 0$.

$$W - E_0 = \frac{\sum_{i=0}^{\infty} a_i^2 E_i}{\sum_{i=0}^{\infty} a_i^2} - E_0 = \frac{\sum_{i=0}^{\infty} a_i^2 (E_i - E_0)}{\sum_{i=0}^{\infty} a_i^2} \geq 0 \quad (\text{B.7})$$

Since a_i^2 is always positive or zero, and $E_i - E_0$ is always positive or zero (E_0 is by definition the lowest energy), this completes the proof. Furthermore, in order for the equal sign to hold, all $a_{i \neq 0} = 0$ since $E_{i \neq 0} - E_0$ is non-zero (neglecting degenerate ground states). This in turn means that $a_0 = 1$, owing to the normalization of Φ , and consequently the wave function is the exact solution.

This proof shows that any approximate wave function will have an energy above or equal to the exact ground state energy. There is a related theorem, known as *MacDonald's Theorem*, which states that the n th root of a set of secular equations (e.g. a CI matrix) is an upper limit to the $(n - 1)$ th excited exact state, within the given symmetry subclass.¹ In other words, the lowest root obtained by diagonalizing a CI matrix is an upper limit to the lowest exact wave functions, the second root is an upper limit to the exact energy of the first excited state, the third root is an upper limit to the exact second excited state, and so on.

B.2 The Hohenberg-Kohn Theorems

In wave mechanics, the electron density is given by the square of the wave function integrated over $N - 1$ electron coordinates and the wave function is determined by solving the Schrödinger equation. For a system of N_{nuclei} nuclei and N_{elec} electrons, the electronic Hamiltonian operator contains the terms given in eq. (B.8).

$$\mathbf{H}_e = - \sum_{i=1}^{N_{\text{elec}}} \frac{1}{2} \nabla_i^2 - \sum_{i=1}^{N_{\text{elec}}} \sum_{A=1}^{N_{\text{nuclei}}} \frac{Z_A}{|\mathbf{R}_A - \mathbf{r}_i|} + \sum_{i=1}^{N_{\text{elec}}} \sum_{j>i}^{N_{\text{elec}}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_{A=1}^{N_{\text{nuclei}}} \sum_{B=1}^{N_{\text{nuclei}}} \frac{Z_A Z_B}{|\mathbf{R}_A - \mathbf{R}_B|} \quad (\text{B.8})$$

Within the Born-Oppenheimer approximation, the last term is a constant. It is seen that the Hamiltonian operator is uniquely determined by the number of electrons and the potential created by the nuclei, \mathbf{V}_{ne} , i.e. the nuclear charges and positions. This means that the ground state wave function (and thereby the electron density) and ground state energy are also given uniquely by these quantities.

Assume now that two different external potentials (which may be from nuclei), \mathbf{V}_{ext} and \mathbf{V}'_{ext} , result in the same electron density, ρ . Two different potentials imply that the two Hamiltonian operators are different, \mathbf{H} and \mathbf{H}' , and the corresponding lowest energy wave functions are different, Ψ and Ψ' . Taking Ψ' as an approximate wave function for \mathbf{H} and using the variational principle yields eq. (B.9).

$$\begin{aligned} \langle \Psi' | \mathbf{H} | \Psi' \rangle &> E_0 \\ \langle \Psi' | \mathbf{H} | \Psi' \rangle + \langle \Psi' | \mathbf{H} - \mathbf{H}' | \Psi' \rangle &> E_0 \\ E'_0 + \langle \Psi' | \mathbf{V}_{\text{ext}} - \mathbf{V}'_{\text{ext}} | \Psi' \rangle &> E_0 \\ E'_0 + \int \rho(\mathbf{r})(\mathbf{V}_{\text{ext}} - \mathbf{V}'_{\text{ext}}) d\mathbf{r} &> E_0 \end{aligned} \quad (\text{B.9})$$

Similarly, taking Ψ as an approximate wave function for \mathbf{H}' yields eq. (B.10).

$$E_0 + \int \rho(\mathbf{r})(\mathbf{V}_{\text{ext}} - \mathbf{V}'_{\text{ext}})d\mathbf{r} > E'_0 \quad (\text{B.10})$$

Addition of these two inequalities gives $E'_0 + E_0 > E'_0 + E_0$, showing that the assumption was wrong. In other words, for the ground state there is a one-to-one correspondence between the electron density and the nuclear potential, and thereby also with the Hamiltonian operator and the energy. In the language of density functional theory, the energy is a unique functional of the electron density, $E[\rho]$.

Using the electron density as a parameter, there is a variational principle analogous to that in wave mechanics. Given an approximate electron density ρ' (assumed to be positive definite everywhere) that integrates to the number of electrons, the energy given by this density is an upper bound to the exact ground state energy, provided that the exact functional is used.

$$\begin{aligned} \int \rho(\mathbf{r})d\mathbf{r} &= N_{\text{elec}} \\ E_0[\rho'] &\geq E_0[\rho] \end{aligned} \quad (\text{B.11})$$

B.3 The Adiabatic Connection Formula

The Hellmann–Feynman theorem (eq. (10.28)) is given in eq. (B.12).

$$\frac{\partial}{\partial \lambda} \langle \Psi_\lambda | \mathbf{H}_\lambda | \Psi_\lambda \rangle = \left\langle \Psi_\lambda \left| \frac{\partial \mathbf{H}_\lambda}{\partial \lambda} \right| \Psi_\lambda \right\rangle \quad (\text{B.12})$$

With the Hamiltonian in eq. (6.5), this gives eq. (B.13).

$$\begin{aligned} \mathbf{H}_\lambda &= \mathbf{T} + \mathbf{V}_{\text{ext}}(\lambda) + \lambda \mathbf{V}_{\text{ee}} \\ \frac{\partial}{\partial \lambda} \langle \Psi_\lambda | \mathbf{H}_\lambda | \Psi_\lambda \rangle &= \left\langle \Psi_\lambda \left| \frac{\partial \mathbf{V}_{\text{ext}}(\lambda)}{\partial \lambda} + \mathbf{V}_{\text{ee}} \right| \Psi_\lambda \right\rangle \end{aligned} \quad (\text{B.13})$$

Integrating over λ between the limits 0 and 1 corresponds to smoothly transforming the non-interacting reference to the real system.

$$\begin{aligned} \int_0^1 \frac{\partial}{\partial \lambda} \langle \Psi_\lambda | \mathbf{H}_\lambda | \Psi_\lambda \rangle d\lambda &= \int_0^1 \left\langle \Psi_\lambda \left| \frac{\partial \mathbf{V}_{\text{ext}}(\lambda)}{\partial \lambda} + \mathbf{V}_{\text{ee}} \right| \Psi_\lambda \right\rangle d\lambda \\ \langle \Psi_1 | \mathbf{H}_1 | \Psi_1 \rangle - \langle \Psi_0 | \mathbf{H}_0 | \Psi_0 \rangle &= E_1 - E_0 = \int_0^1 \left\langle \Psi_\lambda \left| \frac{\partial \mathbf{V}_{\text{ext}}(\lambda)}{\partial \lambda} + \mathbf{V}_{\text{ee}} \right| \Psi_\lambda \right\rangle d\lambda \end{aligned} \quad (\text{B.14})$$

This integration is done under the assumption that the density remains constant, i.e. Ψ_0 and Ψ_1 yield the *same* density. For the term involving the external potential, this allows the integration to be written in terms of the two limits.

$$\begin{aligned} \int_0^1 \left\langle \Psi_\lambda \left| \frac{\partial \mathbf{V}_{\text{ext}}(\lambda)}{\partial \lambda} \right| \Psi_\lambda \right\rangle d\lambda &= \int \rho(\mathbf{r}) \left(\int_0^1 \frac{\partial \mathbf{V}_{\text{ext}}(\mathbf{r}, \lambda)}{\partial \lambda} d\lambda \right) d\mathbf{r} \\ &= \int \rho(\mathbf{r})(\mathbf{V}_{\text{ext}}(1) - \mathbf{V}_{\text{ext}}(0))d\mathbf{r} \\ &= \int \rho(\mathbf{r})\mathbf{V}_{\text{ext}}(1)d\mathbf{r} - \int \rho(\mathbf{r})\mathbf{V}_{\text{ext}}(0)d\mathbf{r} \end{aligned} \quad (\text{B.15})$$

The energy of the non-interacting (E_0) system is given by eq. (B.16) since \mathbf{V}_{ee} makes no contribution.

$$E_0 = \langle \Psi_0 | \mathbf{T} | \Psi_0 \rangle + \int \rho(\mathbf{r}) \mathbf{V}_{\text{ext}}(0) d\mathbf{r} \quad (\text{B.16})$$

Combining eqs (B.14), (B.15) and (B.16) yields eq. (B.17).

$$E_1 = \langle \Psi_0 | \mathbf{T} | \Psi_0 \rangle + \int \rho(\mathbf{r}) \mathbf{V}_{\text{ext}}(1) d\mathbf{r} + \int_0^1 \langle \Psi_\lambda | \mathbf{V}_{ee} | \Psi_\lambda \rangle d\lambda \quad (\text{B.17})$$

Using the fact that $\mathbf{V}_{\text{ext}}(1) = \mathbf{V}_{\text{ne}}$ and the definition of E_1 (eq. (6.9)) we obtain eq. (B.18).

$$J[\rho] + E_{\text{xc}}[\rho] = \int_0^1 \langle \Psi_\lambda | \mathbf{V}_{ee} | \Psi_\lambda \rangle d\lambda \quad (\text{B.18})$$

The exchange–correlation energy can thus be obtained by integrating the electron–electron interaction over the λ variable and subtracting the Coulomb part. The right-hand side of eq. (B.18) can be written in terms of the second-order reduced density matrix eq. (6.14), and the definition of the exchange–correlation hole in eq. (6.21) allows the Coulomb energy to be separated out.

$$\begin{aligned} \langle \Psi_\lambda | \mathbf{V}_{ee} | \Psi_\lambda \rangle &= \frac{1}{2} \int \frac{\rho_2(\lambda, \mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2 \\ &= \frac{1}{2} \int \frac{\rho_1(\mathbf{r}_1) \rho_1(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2 + \frac{1}{2} \int \frac{\rho_1(\mathbf{r}_1) h_{\text{xc}}(\lambda, \mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2 \\ &= J[\rho] + \frac{1}{2} \int \frac{\rho_1(\mathbf{r}_1) h_{\text{xc}}(\lambda, \mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2 \end{aligned} \quad (\text{B.19})$$

Defining $\mathbf{V}_{\text{xc}}^{\text{hole}}$ as in eq. (B.20) gives the *adiabatic connection formula* (eq. (6.50)).

$$\begin{aligned} \mathbf{V}_{\text{xc}}^{\text{hole}}(\lambda, \mathbf{r}_1) &= \frac{1}{2} \int \frac{h_{\text{xc}}(\lambda, \mathbf{r}_1, \mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_2 \\ E_{\text{xc}} &= \int_0^1 \langle \Psi_\lambda | \mathbf{V}_{\text{xc}}^{\text{hole}}(\lambda) | \Psi_\lambda \rangle d\lambda \\ E_{\text{xc}} &= \int \rho(\mathbf{r}) \left(\int_0^1 \mathbf{V}_{\text{xc}}^{\text{hole}}(\lambda, \mathbf{r}) d\lambda \right) d\mathbf{r} \end{aligned} \quad (\text{B.20})$$

Reference

1. J. K. L. MacDonald, *Phys. Rev.*, **43** (1933), 830.

Appendix C

Atomic Units

In electronic structure calculations it is convenient to work in the *atomic unit* (au) system, which is defined by setting $m_e = e = \hbar = 1$. From these values follow related quantities, as shown in Table C.1.

Table C.1 The atomic unit system

Symbol	Quantity	Value in au	Value in SI units
m_e	Electron mass	1	$9.110 \times 10^{-31} \text{ kg}$
e	Electron charge	1	$1.602 \times 10^{-19} \text{ C}$
t	Time	1	$2.419 \times 10^{-17} \text{ s}$
\hbar	$\hbar/2\pi$ (atomic momentum unit)	1	$1.055 \times 10^{-34} \text{ J s}$
h	Planck's constant	2π	$6.626 \times 10^{-34} \text{ J s}$
a_0	Bohr radius (atomic distance unit)	1	$5.292 \times 10^{-11} \text{ m}$
E_H	Hartree (atomic energy unit)	1	$4.360 \times 10^{-18} \text{ J}$
c	Speed of light	137.036	$2.998 \times 10^8 \text{ m/s}$
α	Fine structure constant ($= e^2/\hbar c 4\pi\epsilon_0 = 1/c$)	0.00729735	0.00729735
μ_B	Bohr magneton ($= e\hbar/2m_e$)	$1/2$	$9.274 \times 10^{-24} \text{ J/T}$
μ_N	Nuclear magneton	2.723×10^{-4}	$5.051 \times 10^{-27} \text{ J/T}$
$4\pi\epsilon_0$	Vacuum permittivity	1	$1.113 \times 10^{-10} \text{ C}^2/\text{J m}$
μ_0	Vacuum permeability ($4\pi/c^2$)	6.692×10^{-4}	$1.257 \times 10^{-6} \text{ N s}^2/\text{C}^2$

Appendix D

Z-Matrix Construction

All calculations need as input a molecular geometry. This is commonly given by one of the following three methods:

- (1) Cartesian coordinates
- (2) Internal coordinates
- (3) Via a graphical interface.

Generating Cartesian coordinates by hand is only realistic for small molecules. If, however, the geometry is taken from outside sources, such as an X-ray structure, Cartesian coordinates are often the natural choice. Similarly, a graphical interface produces a set of Cartesian coordinates for the underlying program, which carries out the actual calculation.

Generating internal coordinates such as bond lengths and angles by hand is relatively simple, even for quite large molecules. One widely used method is the *Z-matrix* where each atom is specified in terms of a distance, angle and torsional angle to other atoms. It should be noted that internal coordinates are *not* necessarily related to the actual bonding, they are only a convenient method for specifying the geometry. The internal coordinates are usually converted to Cartesian coordinates before any calculations are carried out. Geometry optimizations, however, are often done in internal coordinates in order to remove the six (five) translational and rotational degrees of freedom.

Construction of a *Z-matrix* begins with a drawing of the molecule and a suitable numbering of the atoms. *Any* numbering will result in a valid *Z-matrix*, although assignment of numerical values to the parameters is greatly facilitated if the bonding and symmetry of the molecule is considered when the numbering is performed (see the examples below). The *Z-matrix* specifies the position of each atom in terms of a distance, an angle and a torsional angle relative to other atoms. The first three atoms, however, are slightly different. The first atom is always positioned at the origin of the coordinate system. The second atom is specified as having a distance to the first atom,

and is placed along one of the Cartesian axis (usually x or z). The third atom is specified by a distance to either atom 1 or 2, and an angle to the other atom. All subsequent atoms need a distance, an angle and a torsional angle to uniquely specify the position. The atoms are normally identified either by the chemical symbol or by their atomic number.

If the molecular geometry is optimized by the program then only rough estimates of the parameters are necessary. In terms of internal coordinates, this is fairly easy. Some typical bond lengths (Å) and angles are given below.

A—H: A=C: 1.10; A=O,N: 1.00; A=S,P: 1.40

A—B: A,B=C,O,N: 1.40–1.50

A=B: A,B=C,O,N: 1.20–1.30

A≡B: A,B=C,N: 1.20

A—B: A=C, B=S,P: 1.80

Angles around sp^3 -hybridized atoms: 110°

Angles around sp^2 -hybridized atoms: 120°

Angles around sp -hybridized atoms: 180°

Torsional angles around sp^3 -hybridized atoms: separated by 120°

Torsional angles around sp^2 -hybridized atoms: separated by 180°

Such estimates allow specification of molecules with up to 50–100 atoms fairly easy. For larger molecules, however, it becomes cumbersome. In such cases the molecule is often built from pre-optimized fragments. This is typically done by means of a graphical interface, i.e. the molecule is pieced together by selecting fragments (such as amino acids) and assigning the bonding between the fragments.

Below are some examples of how to construct Z-matrices. Figure D.1 shows acetaldehyde.




Figure D.1 Atom numbering for acetaldehyde

C1	0	0.00	0	0.0	0	0.0
O2	1	1.20	0	0.0	0	0.0
H3	1	1.10	2	120.0	0	0.0
C4	1	1.50	2	120.0	3	180.0
H5	4	1.10	1	110.0	2	0.0
H6	4	1.10	1	110.0	2	120.0
H7	4	1.10	1	110.0	2	-120.0

The definition of the torsional angles is illustrated in Figure 2.7. To emphasize the symmetry (C_s) of the above conformation, the Z-matrix may also be given in terms of *symbolic* variables, where variables that are equivalent by symmetry have identical names.

C1						
O2	1	R1				
H3	1	R2	2	A1		
C4	1	R3	2	A2	3	D1
H5	4	R4	1	A3	2	D2
H6	4	R5	1	A4	2	D3
H7	4	R5	1	A4	2	-D3

R1 = 1.20

R2 = 1.10

R3 = 1.50

R4 = 1.10

R5 = 1.10

A1 = 120.0

A2 = 120.0

A3 = 110.0

A4 = 110.0

D1 = 180.0

D2 = 0.0

D3 = 120.0

Some important things to notice:

- (1) Each atom must be specified in terms of atoms already defined, i.e. relative to atoms *above*.
- (2) Each specification atom can only be used *once* in each line.
- (3) The specification in terms of distance, angle and torsional angle has *nothing* to do with the bonding in the molecule, e.g. the torsional angle for C₄ in acetaldehyde is given to H₃, but there is no bond between O₂ and H₃. A Z-matrix, however, is usually constructed such that the distances, angles and torsional angles follow the bonding. This makes it much easier to estimate reasonable values for the parameters.
- (4) *Distances* should always be *positive*, and *angles* always in the range 0°–180°. Torsional angles may be taken in the range –180°–180°, or 0°–360°.
- (5) The symbolic variables show explicitly which parameters are constrained to have the same values, i.e. H₆ and H₇ are symmetry equivalent and must therefore have the same distances and angles, and a sign difference in the torsional angle. Although the R4 and R5 (and A3 and A4) parameters have the same values initially, they will be different in the final optimized structure.

The limitation that the angles must be between 0° and 180° introduces a slight complication for linear arrays of atoms, such as a cyano group.

Specification of the nitrogen in term of a distance to C₂ and an angle to C₁ does not allow a unique assignment of a torsional angle since the C₁—C₂—N₆ angle is linear, which makes the torsional angle undefined. There are two methods for solving this problem, either by specifying N₆ relative to C₁ with a long distance:




Figure D.2 Atom numbering for methyl cyanide

C1						
C2	1	R1				
H3	1	R2	2	A1		
H4	1	R2	2	A1	3	D1
H5	1	R2	2	A1	3	-D1
N6	1	R3	3	A2	4	D2

$$R1 = 1.50$$

$$R2 = 1.10$$

$$R3 = 2.70$$

$$A1 = 110.0$$

$$A2 = 110.0$$

$$D1 = 120.0$$

$$D2 = 120.0$$

Note that the variables imply that the molecule has C_{3v} symmetry. Alternatively, a *Dummy Atom* (X) may be introduced.




Figure D.3 Atom numbering for methyl cyanide including a dummy atom

C1						
C2	1	R1				
H3	1	R2	2	A1		
H4	1	R2	2	A1	3	D1
H5	1	R2	2	A1	3	-D1
X6	2	R3	1	A2	3	D2
N7	2	R4	6	A3	1	D3

R1 = 1.50
 R2 = 1.10
 R3 = 1.00
 R4 = 1.20
 A1 = 110.0
 A2 = 90.0
 A3 = 90.0
 D1 = 120.0
 D2 = 0.0
 D3 = 180.0

A dummy atom is just a point in space and has no significance in the actual calculation. The above two Z-matrices give identical Cartesian coordinates. The R3 variable has arbitrarily been given a distance of 1.00, and the D2 torsional angle of 0.0° is also arbitrary – any other values may be substituted without affecting the coordinates of the real atoms. Similarly, the A2 and A3 angles should just add up to 180° ; their individual values are not significant. The function of a dummy atom in this case is to break up the problematic 180° angle into two 90° angles. It should be noted that the introduction of dummy atoms does not increase the number of (non-redundant) parameters, although there are formally three more variables for each dummy atom. The dummy variables may be identified by excluding them from the symbolic variable list, or by explicitly forcing them to be non-optimizable parameters.

When a molecule is symmetric, it is often convenient to start the numbering with atoms lying on a rotation axis or in a symmetry plane. If there are no real atoms on a rotation axis or in a mirror plane, dummy atoms can be useful for defining the symmetry element. Consider for example the cyclopropenyl system, which has D_{3h} symmetry. Without dummy atoms, one of the C—C bond lengths will be given in terms of the two other C—C distances and the C—C—C angle, and it will be complicated to force the three C—C bonds to be identical. By introducing two dummy atoms to define the C_3 axis, this becomes easy.




Figure D.4 Atom numbering for the cyclopropyl system

X1						
X2	1		1.00			
C3	1	R1	2	90.0		
C4	1	R1	2	90.0	3	120.0
C5	1	R1	2	90.0	3	-120.0
H6	1	R2	2	90.0	3	0.0
H7	1	R2	2	90.0	3	120.0
H8	1	R2	2	90.0	3	-120.0

R1 = 0.80

R2 = 1.90

In this case there are only two genuine variables, the others are fixed by symmetry.

Let us finally consider two *Z*-matrices for optimization to transition structures, the Diels–Alder reaction of butadiene and ethylene, and the [1,5]-hydrogen shift in (*Z*)-1,3-pentadiene. To enforce the symmetries of the TS's (C_s in both cases) it is again advantageous to use dummy atoms.




Figure D.5 Atom numbering for the transition structure of the Diels–Alder reaction of butadiene and ethylene

X1						
X2	1	1.00				
C3	1	R1	2	90.0		
C4	1	R1	2	90.0	3	180.0
C5	3	R2	1	A1	2	180.0
C6	4	R2	1	A1	2	180.0
C7	3	R3	1	A2	2	D1
C8	4	R3	1	A2	2	–D1
H9	3	R4	5	A3	6	D2
H10	3	R5	5	A4	6	–D3
H11	4	R4	6	A3	5	–D2
H12	4	R5	6	A4	5	D3
H13	5	R6	3	A5	1	–D4
H14	6	R6	4	A5	1	D4
H15	7	R7	8	A6	4	D5
H16	7	R8	8	A7	4	–D6
H17	8	R7	7	A6	3	–D5
H18	8	R8	7	A7	3	D6

R1 = 1.40

R2 = 1.40

R3 = 2.20

R4 = 1.10

R5 = 1.10

$R6 = 1.10$
 $R7 = 1.10$
 $R8 = 1.10$
 $A1 = 60.0$
 $A2 = 70.0$
 $A3 = 120.0$
 $A4 = 120.0$
 $A5 = 120.0$
 $A6 = 120.0$
 $A7 = 120.0$
 $D1 = 60.0$
 $D2 = 170.0$
 $D3 = 30.0$
 $D4 = 170.0$
 $D5 = 100.0$
 $D6 = 100.0$

The mirror plane is defined by the two dummy atoms and the fixing of the angles and torsional angle of the first two carbons. The torsional angles for atoms C5 and C6 are dummy variables as they only define the orientation of the plane of the first four carbon atoms relative to the dummy atoms, and may consequently be fixed at 180° . Note that the C₅—C₆ and C₇—C₈ bond distances are given implicitly in terms of the R2/A1 and R3/A2 variables. The presence of such “indirect” variables means that some experimentation is necessary for assigning proper values to the “direct” variables. The forming C—C bond is given directly as one of the Z-matrix variables, R3, which facilitates a search for a suitable start geometry for the TS optimization, for example by running a series of constrained optimizations with fixed R3 distances.

The [1,5]-hydrogen shift in (*Z*)-1,3-pentadiene is an example of a “narcissistic” reaction, with the reactant and product being identical. The TS is therefore located exactly at the halfway point, and has a symmetry different from either the reactant or product. By suitable constraints on the geometry the TS may therefore be located by a *minimization* within a symmetry-constrained geometry.




Figure D.6 Atom numbering for the transition structure for the [1,5]-hydrogen shift in (*Z*)-1,3-pentadiene

X1						
X2	1	1.00				
X3	1	1.00	2	90.0		
C4	1	R1	2	90.0	3	90.0
C5	1	R1	2	90.0	3	-90.0
C6	4	R2	1	A1	2	180.0
C7	5	R2	1	A1	2	180.0
C8	1	R3	3	A2	2	180.0
H9	4	R4	6	A3	8	-D1
H10	4	R5	6	A4	8	D2
H11	5	R4	7	A3	8	D1
H12	5	R5	7	A4	8	-D2
H13	6	R6	4	A5	1	D3
H14	7	R6	5	A5	1	-D3
H15	1	R7	3	A6	2	180.0
H16	1	R8	2	A7	3	0.0

R1 = 1.30

R2 = 1.40

R3 = 2.10

R4 = 1.10

R5 = 1.10

R6 = 1.10

R7 = 3.20

R8 = 0.70

A1 = 80.0

A2 = 90.0

A3 = 120.0

A4 = 120.0

A5 = 120.0

A6 = 90.0

A7 = 60.0

D1 = 160.0

D2 = 60.0

D3 = 160.0

The mirror plane is defined by the dummy atoms. The migrating hydrogen H₁₆ is not allowed to move out of the plane of symmetry, and must consequently have the same distance to C₄ and C₅. A minimization will locate the lowest energy structure within the given C_s symmetry, and a subsequent frequency calculation will reveal that the optimized structure is a TS, with the imaginary frequency belonging to the *a*'' representation (breaking the symmetry).

Index

- 3-21G basis sets 203–4, 218
 - dissociation curves 361
 - relative energies of isomers 375
- 3D-QSAR 560
- 5Z *see* pentuple zeta
- 6-21G basis sets 203–4
- 6-31G basis sets 203–4, 216, 218, 375–7
- 6-311G basis sets 203–4, 216–18
- STO-3G basis sets 375–7
- AC *see* asymptotic corrected
- accuracy 550
- accuracy-per-function criteria 192
- ACF *see* adiabatic connection formula
- ACM *see* adiabatic connection model
- ACPF *see* averaged coupled-pair functional
- addition reactions 489–90, 497–506
- adiabatic approximation 84–5
- adiabatic connection formula (ACF) 252, 572–3
- adiabatic connection model (ACM) 252
- Ahlrichs type basis sets 205
- AM1 *see* Austin model 1
- AMBER force field 42, 63–4
- angle bending curves 370, 371
- angular correlation 196
- ANO *see* atomic natural orbitals
- anti* conformations, force field methods 31–4
- approximating functions 538–41
- approximations 548
- APT *see* atomic polar tensor
- asymptotic corrected (AC) functionals 259
- atom types 23, 24
- atomic natural orbitals (ANO) 202, 205–6
- atomic polar tensor (APT) charges 304
- atomic units (Appendix C) 574
- atomization energy 547–8
- atoms in molecules (AIM) method
 - qualitative theories 493
 - wave function analysis 299–303
- augmented Hessian 387
- Austin model 1 (AM1) 121–2, 124, 125–7, 130–1
- autocorrelation functions 472
- autoscaling 554
- average values 549
- averaged coupled-pair functional (ACPF) 176
- avoided crossings 85
- B3 *see* Becke 3 parameter functional
- B3LYP
 - density functional methods 255–7, 264
 - dipole moment convergence 357–8
 - dissociation curves 369–70
 - geometry convergence 353
 - simulation techniques 484
 - vibrational frequency convergence 360–1
- B97 model 249
- basis functions 293–6
- basis set superposition errors (BSSE) 225–7
- basis sets 192–231
 - Ahlrichs type 205
 - approximation 93–8
 - atomic natural orbitals 202, 205–6
 - balance 197
 - classification 194–8
 - composite extrapolation procedures 213–21
 - computational issues 212–13
 - contracted 200–11
 - convergence 208
 - correlation consistent 206–11, 219
 - dipole moment convergence 357, 372
 - dissociation curves 361–2
 - Dunning–Huzinaga 204–5
 - effective core potentials 222–5
 - even-tempered 198–200
 - extrapolation 208–11, 213–21
 - Gaussian type orbitals 192–4
 - geometry convergence 350–3, 371–2
 - isodesmic reactions 221–2

- isogyric reactions 221–2
- mathematical methods 539, 541
- MINI, MIDI and MAXI 205
- minimum 194
- molecular properties 348–9
- optimization 199–200
- parameterization 199, 212
- plane wave basis functions 211–12
- polarization 196–8, 206, 228
- polarization consistent 207–8
- Pople style 202–4
- pseudopotential 222–3
- pseudospectral methods 227–9
- recent developments 212–13
- relative energies of isomers 375–7
- simulation techniques 484
- Slater type orbitals 192–4
- split valence 195, 205
- superposition errors 225–7
- total energy convergence 354–6
- vibrational frequency convergence 359–60, 373
- wave function analysis 295, 297
- well-tempered 198–200
- Becke 3 parameter functional (B3) 252
- Becke–Roussel (BR) functional 251
- Bell correction 461
- Bell–Evans–Polanyi (BEP) principle 506–10
- bending energies, force field methods 27–30, 59–61, 64–5
- BEP *see* Bell–Evans–Polanyi
- BFGS *see* Broyden–Fletcher–Goldfarb–Shanno
- bimodal distribution 550
- Bloch theorem 114
- blue moon sampling 464
- BLYP
 - density functional methods 255–6, 264
 - dipole moment convergence 357–8
 - dissociation curves 369–70
 - geometry convergence 353
 - vibrational frequency convergence 360–1
- BO *see* bond order
- Boltzmann distributions
 - condensed phases 440–1
 - simulation techniques 445–6
 - statistical mechanics 427
 - transition state theory 425–6
- bond critical points 302
- bond dissociation
 - basis sets 221–2, 361–2
 - configuration interaction 145–53
 - curves 361–70
 - density functional theory 259, 369–70
 - force field methods 25, 59
 - Hartree–Fock theory 361–2
 - restricted Hartree–Fock methods 145–53
 - transition state theory 425–6
 - unrestricted Hartree–Fock methods 148–53
 - wave function differences 363–9
- bond order (BO) 296
- Born model 480–1
- Born–Oppenheimer approximation
 - chemistry principles 19, 20
 - electronic structure methods 80, 82–6
 - Hohenberg–Kohn theorems 571
 - molecular properties 333
 - rigid-rotor harmonic-oscillator approximation 429
 - separation of variables 9, 11
 - simulation techniques 457, 458–9, 463
 - transition state theory 422
- bound solutions 13, 16
- Boys localization scheme 305–8
- BR *see* Becke–Roussel
- bra-ket notation 82, 530–1
- branching 2
- Bravais lattices 113
- Breit correction 5
- Breit interaction 285, 289, 291
- Brillouin zones 113–14
- Brillouin’s theorem 140, 164, 171
- Brownian dynamics 455
- Broyden–Fletcher–Goldfarb–Shanno (BFGS) method 388
- Brueckner theory 175–6
- BSSE *see* basis set superposition errors
- Buckingham-type potentials, force field methods 36–8
- cake critical points 302
- canonical orthogonalization 533
- Car–Parrinello (CP) methods 457–9, 476
- cartesian polar systems 514–15
- CASSCF *see* complete active space self-consistent field
- CBS *see* complete basis set
- cc/CC *see* correlation consistent; coupled cluster
- CCD *see* coupled cluster doubles
- CCSD(T) model
 - dipole moment convergence 357, 372
 - dissociation curves 367–9
 - geometry convergence 350–2, 371–2
 - total energy convergence 354–6
 - vibrational frequency convergence 359
- central field model 13
- centre of mass coordinate systems 1, 8–9
- CEPA *see* coupled electron pair approximation
- CF *see* Coulson–Fischer
- CG *see* conjugate gradient
- CGTO *see* contracted Gaussian type orbitals
- CHA *see* chemical Hamiltonian approach
- chain optimization 399–400
- charge controlled reactions 488
- charge iteration Hückel methods 128
- charges *see* electrostatic energies
- CHARMM force field 63–4
- chemical Hamiltonian approach (CHA) 227
- chemist’s notation 96

- CHF *see* coupled Hartree–Fock
Cholesky decomposition method 183
CI *see* configuration interaction
CI-NEB *see* climbing image NEB
cis isomerism 32, 34
Claisen reactions 435–6
classical mechanics 6, 12–14
 molecular properties 333–4
 see also force field methods
classical valence bond (VB) theory 269–70
climbing image NEB (CI-NEB) optimization 401
CNDO *see* complete neglect of differential overlap
combined parameterization 56
comparative molecular field analysis (COMFA) 560
complete active space self-consistent field (CASSCF) method 155–8
 dissociation curves 363–5, 367
 valence bond theory 273–5
complete basis set (CBS) 214–18
complete neglect of differential overlap (CNDO) approximation 117–18
complex
 conjugates 82, 515
 numbers 514–15
composite extrapolation procedures 213–21
condensed phases 439–43
condition number 524
conductor-like screening model (COSMO) 483
configuration interaction (CI) 137–59, 183–5
 beryllium atom 177–8
 coupled cluster theory 172–8
 direct methods 144–5
 dissociation curves 367–9
 mathematical methods 525–6, 529
 matrix dimensions 141–3
 matrix elements 138–41
 molecular properties 322
 multi-configuration self-consistent field 153–8
 multi-reference 158–9
 optimization techniques 382
 perturbation theory 165, 174–8, 183–5
 quadratic 176
 RHF dissociation 145–53
 simulation techniques 456
 size consistency 153
 size extensivity 153
 spin contamination 148–53
 state-selected 159
 truncated 143–4
 UHF dissociation 148–53
configurational state functions (CSF) 139, 141–2, 144–5
conformational sampling 409–15
conical intersection 505
conjugate gradient (CG) methods 384–5
conjugate peak refinement (CPR) 400
conjugated systems 48–50, 58–62
constrained
 optimization 407–9
 sampling methods 463–4
continuum models 476–84
contracted basis sets 200–11
 Ahlricks type basis sets 205
 atomic natural orbital basis sets 205–6
 correlation consistent basis sets 206–11, 219
 degree of contraction 201
 Dunning–Huzinaga basis sets 204–5
 extrapolation 208–11
 general contraction 201–2
 MINI, MIDI and MAXI 205
 polarization consistent basis sets 207–8
 Pople style basis sets 202–4
 segmented contraction 201–2
contracted Gaussian type orbitals (CGTO) 200–6
coordinate driving 394–5
coordinate selection 390–4, 405–6
coordinate transformations 520–9
 CI wave functions 529
 computational considerations 529
 examples 525–6
 rotations 520–2
 similarity transformations 522
 Slater determinants 528–9
 unitary transformations 526
 vibrational normal coordinates 526–8
coordination compounds
 force field methods 58–62
 parameterization 58–62
Cope rearrangements 508–12
core–core repulsion 120
correlation
 coefficients 552
 energy 19
 functions 471–2
 illustrative example 558–9
 many data sets 553–9
 multiple linear regression 555–6, 558–9
 multiple-descriptor data sets 553–5
 partial least squares 557–9
 principle component analysis 556–7, 558–9
 quality analysis 553–5
 two data sets 550–3
correlation consistent (cc) basis sets 206–11, 219
 dipole moment convergence 357, 372
 dissociation curves 361–2
 geometry convergence 350–3, 371–2
 relative energies of isomers 376
 simulation techniques 484
 total energy convergence 354–6
 vibrational frequency convergence 359–60, 373
COSMO *see* conductor-like screening model

- Coulomb
 - correlation 134
 - gauge 284
 - holes 134, 242–3
 - integrals 89–90, 96, 228
 - interactions 1, 9, 238–9, 496–7
 - potential 465–6
- Coulson–Fischer (CF) functions 270
- counterpoise (CP) correction 226
- coupled cluster (CC) theory 137, 169–78, 183–6
 - beryllium atom 177–8
 - configuration interaction 172–8
 - perturbation theory 174–8
 - truncated CC methods 172–4
- coupled cluster doubles (CCD) 172, 176
- coupled electron pair approximation (CEPA) 176
- coupled Hartree–Fock (CHF) theory 328
- coupled perturbed Hartree–Fock (CPHF) theory 325–9, 343
- CP *see* Car–Parrinello; counterpoise correction
- CPHF *see* coupled perturbed Hartree–Fock
- CPR *see* conjugate peak refinement
- cross-correlation functions 472
- cross products 518
- cross terms
 - force field methods 47–8, 51–2, 62
 - molecular properties 319–20
- cross-validations 554–5
- crystalline orbitals 114
- CSF *see* configurational state functions

- damping 101
- Darwin correction 281–2, 286
- Davidson algorithm 145
- Davidson correction 174–5
- Davidson–Fletcher–Powell (DFP) method 388
- DBOC *see* diagonal Born–Oppenheimer correction
- degenerate eigenvalues 523
- degree of contraction 201
- delocalized internal coordinates 394
- density functional theory (DFT) 232–67
 - computational considerations 260–3
 - Coulomb holes 242–3
 - dipole moment convergence 357–8, 372
 - dissociation curves 369–70
 - electronic structure methods 81–2, 111
 - exchange–correlation functionals 243–55, 259
 - exchange–correlation holes 240–3
 - Fermi holes 242–3
 - generalized gradient approximation 248–56
 - generalized random phase approximations 253–4
 - geometry convergence 353–4
 - gradient-corrected methods 248–56
 - Hohenberg–Kohn theorem 232, 239
 - hyper-GGA methods 252–6
 - Jacob’s ladder classification 246–55
 - Kohn–Sham theory 235–6, 239, 257–8, 260–3
 - limitations 258–60
 - local density approximation 245, 246–8, 254, 257
 - mathematical methods 544
 - meta-GGA methods 250–2, 254–6
 - molecular properties 346
 - orbital-free 233–5
 - parameterization 238, 247–8
 - performance and properties 255–8
 - qualitative theories 492–4, 496
 - reduced density matrix methods 236–40
 - simulation techniques 459
 - time-dependent 346
 - vibrational frequency convergence 360–1, 374
- detailed balance condition 449
- determinants 518–19
- DFP *see* Davidson–Fletcher–Powell
- DFT *see* density functional theory
- DHF *see* Dirac–Hartree–Fock
- diagonal Born–Oppenheimer correction (DBOC) 84, 86
- diamagnetic shielding 332–3
- diamagnetism 335
- Diels–Alder reactions
 - qualitative theories 490
 - rigid-rotor harmonic-oscillator approximation 435–6
- Z-matrix construction 580
- different orbitals for different spins (DODS) 99
- differential
 - equations 535–8
 - operators 531–2
- diffuse functions 203
- diffusion methods 414
- diffusion quantum Monte Carlo 188
- DIIS *see* direct inversion in the iterative subspace
- dimer method 405
- dipole moment convergence 356–8
 - Ab initio* methods 356–7
 - density functional theory 357–8, 372
 - problematic systems 372
- dipole–dipole interactions 34–5, 40–2, 66–7
 - see also* van der Waals energies
- Dirac equation 6–8, 278–88
 - electric potentials 280–4
 - four-component calculations 287–8
 - magnetic potentials 282–4
 - many-particle systems 284–7
 - molecular properties 331
- Dirac–Fock equation 288, 289, 291
- Dirac–Fock–Breit 352
- Dirac–Hartree–Fock (DHF) theory 223
- direct
 - configuration interaction methods 144–5
 - electron correlation methods 181–2
 - minimization techniques 103–4
 - self-consistent field theory 108–10

- direct inversion in the iterative subspace (DIIS) 102–3, 104
- dispersion forces 35
- dissociation *see* bond dissociation
- distance geometry methods, global minima 414–15
- distributed multipole analysis (DMA) 44, 298–9
- distribution functions 470–1
- DMA *see* distributed multipole analysis
- docking 415–16
- DODS *see* different orbitals for different spins
- dot products 517–18
- double zeta (DZ) basis sets 194–7
- double zeta plus polarization (DZP) 196–7, 214
 - wave function analysis 295, 297
- Douglas–Kroll transformations 289
- DREIDING force field 63–4
- dummy atoms 578–82
- Dunning–Huzinaga basis sets 204–5
- dynamic methods 406–7
- dynamical effects 425–6
- dynamical equations 3, 4, 5–12
 - nuclear and electronic variables 10–11
 - separation of variables 8–12
 - solving 8
 - space and time variables 10
- DZ *see* double zeta
- DZP *see* double zeta plus polarization

- EC *see* electron correlation
- ECP *see* effective core potentials
- Edmiston–Ruedenberg localization scheme 306–8
- EF *see* eigenvector following
- effective core potentials (ECP) 222–5
- effective fragment method 75
- efficiency-per-function criteria 192
- EHT *see* extended Hückel theory
- eigenvector following (EF) 387
- Einstein dynamical equation 6, 8
- electric fields
 - external 315, 316–17, 329
 - internal 329
- electric potentials 280–4
- electromagnetic interactions 4–5, 17–18
- electron
 - density 299–304
 - propagators 344
 - spin 333
- electron correlation (EC) methods 133–91
 - beryllium atom 177–8
 - configuration interaction 137–59, 183–5
 - convergence 136, 152, 154, 166–8, 180–1
 - coupled cluster theory 137, 169–78, 183–6
 - direct methods 181–2
 - dissociation 145–53
 - excited Slater determinants 135–7, 139, 146, 163–5
 - excited states 186–7
 - interelectronic distance 178–81
 - localized orbital methods 182–3
 - many-body perturbation theory 137, 159–69, 174–8
 - Møller–Plesset perturbation theory 162–9, 174–8
 - multi-configuration self-consistent field 153–8, 187
 - projected Møller–Plesset methods 168–9
 - quantum Monte Carlo methods 187–9
 - resolution of the identity method 180–1, 183
 - size consistency 153
 - size extensivity 153
 - spin contamination 148–53
 - summary of methods 183–6
 - truncated coupled cluster methods 172–4
 - unrestricted Møller–Plesset methods 168–9
- electron–nuclear dynamics (END) method 463
- electronic
 - chemical potential 493
 - degrees of freedom 433
 - embedding 75
 - Hamiltonian operators 83
- electronic structure methods 80–132
 - adiabatic approximation 84–5
 - basis set approximation 93–8
 - Born–Oppenheimer approximations 80–92
 - Hartree–Fock theory 80–2, 87, 91–2, 93–100
 - independent-particle models 80–1
 - Koopmans’ theorem 92–3
 - parameterization 118–25, 130
 - periodic systems 113–15
 - self-consistent field theory 86–7, 92, 96–7, 100–13
 - semi-empirical methods 115–18
 - Slater determinants 81, 87–92
 - variational problem 98–9
- electrophilic reactions 489, 492, 512
- electrostatic energies
 - charges 40–2
 - computational considerations 65–7
 - fluctuating charge model 44–5
 - force field methods 40–7, 57, 65–7
 - multipoles 43–7
 - parameterization 57
 - polarizabilities 43–7
 - see also* dipole–dipole interactions
- electrostatic potential (ESP)
 - mathematical methods 545
 - wave function analysis 296–9, 312
- END *see* electron–nuclear dynamics
- entropy 429, 433–9
- entrywise products 517
- equation of motion (EOM) methods 346
- ergodic hypothesis 441, 447
- errors 547–9
- ESP *see* electrostatic potential
- ETS *see* extended transition state

- Euler algorithm 417–18
EVB *see* extended valence bond
even-tempered basis sets 198–200
Ewald sums 67, 466–7
exact wave functions 321
exchange energy 18
exchange integrals 238–9
 basis sets 228
 electronic structure methods 89–90, 96
exchange–correlation functionals 243–55
 generalized gradient approximation 248–56
 generalized random phase approximations 253–4
 gradient-corrected methods 248–56
 hyper-GGA methods 252–6
 Jacob's ladder classification 246–55
 limitations 259
 local density approximation 246–8, 254, 257
 meta-GGA methods 250–2, 254–6
exchange–correlation holes 240–3
excited electron correlation methods 186–7
excited Slater determinants 135–7, 139, 146, 163–5
extended Hückel theory (EHT) 107, 127–8
extended Lagrange methods 45, 457–9
extended transition state (ETS) approach 496
extended valence bond (EVB) method 73
external electric fields 315, 316–17, 329
external magnetic fields 315, 318, 331–2
extrapolation 101

fast Fourier transforms (FFT) 115, 542
fast multiple sums 67
fast multipole moment (FMM) method 111, 467–8
Fermi
 contact 287
 contact operators 332–3, 334, 336
 correlation 134
 holes 134, 242–3
FF *see* force field methods
FFT *see* fast Fourier transforms
first-order differential equations 535–6
first-order regular approximation (FORA)
 method 282
fixed node approximation 189
Fletcher–Reeves (FR) method 385
fluctuating charge model 44–5
fluctuation potential 163
FMM *see* fast multipole moment
FMO *see* frontier molecular orbital theory
Fock matrices *see* Hartree–Fock theory
Fock operators 91–2, 99, 104–5
Foldy–Wouthuysen transformations 289
FORA *see* first-order regular approximation
force field (FF) methods 22–79
 accuracy/generalizability 53–5, 71–2
 advantages 69–70, 72
 atom types 23, 24
 bending energies 27–30, 59–61, 64–5
 computational considerations 65–7
 conjugated systems 48–50, 58–62
 coordination compounds 58–62
 cross terms 47–8, 51–2, 62
 differences in force fields 62–5
 electrostatic energies 40–7, 57, 65–7
 energy comparisons 50–1
 energy types 23–51
 errors 68
 functional forms 62–3
 functional groups 22–3, 53–4
 generic parameters 57–8
 hybrid force field electronic structure methods 74–7
 hydrogen bonds 39–40
 hyperconjugation 48
 limitations 69–70, 72–3
 out-of-plane bending energies 30
 parameterization 51–62, 68–9
 practical considerations 69
 reactive energy surfaces 73–4
 relative energies of isomers 378
 small rings 48–50
 stretch energies 26–7, 64–5
 structurally different molecules 50–1
 torsional energies 30–4, 42–3, 48, 57, 63
 transition structure modelling 70–4
 universal force fields 62
 validation of force fields 67–9
 van der Waals energies 34–40, 42–3, 52–3, 57, 61, 65–7
forces 4–5
FORS *see* full optimized reaction space
four index transformations 141
Fourier transformations (FT) 541–2
FR *see* Fletcher–Reeves
free energy methods
 simulation techniques 472–5
 thermodynamic integration 473–5
 thermodynamic perturbation 472–3
frontier molecular orbital (FMO) theory 487–92
frozen-core approximation 136
FT *see* Fourier transformations
Fukui function 492–4
full optimized reaction space (FORS) 155–8
functional forms 62–3
functional groups 22–3, 53–4
functionals 530
functions 530, 531
fundamental forces 4–5

GA *see* genetic algorithms
GAPT *see* generalized atomic polar tensor
gauge dependence 338–9
gauge including/invariant atomic orbitals (GIAO) 338–9
gauge origin 282, 330

- Gaunt interaction 285, 289, 291
 Gaussian type orbitals (GTO) 192–4, 200–6, 214–15
 GB/SA *see* generalized Born/surface area
 GDIIS *see* geometry direct inversion in the iterative subspace
 general contraction 201–2
 general functions
 conjugate gradient methods 384–5
 coordinate selection 390–4, 405–6
 GDIIS extrapolations 389–90
 Hessian computation 385–9
 Newton–Raphson methods 385–94
 optimization techniques 381, 383–407
 saddle points 381, 394–407
 Simplex method 383
 steepest descent method 383–4
 step control 386–7
 general relativity 279
 generalized atomic polar tensor (GAPT)
 charges 304, 311
 generalized Born/surface area (GB/SA) model 480
 generalized gradient approximation (GGA)
 methods 248–56
 generalized inverse matrices 519–20, 524–5
 generalized random phase approximations (GRPA) 253–4
 generalized valence bond (GVB) theory 275
 generic parameters 57–8
 genetic algorithms (GA) 413–14
 geometry convergence 350–4
 Ab initio methods 350–3
 density functional theory 353–4
 problematic systems 371–2
 geometry direct inversion in the iterative subspace (GDIIS) extrapolations 389–90
 geometry perturbations 315, 319, 339–43
 GGA *see* generalized gradient approximation
 ghost orbitals 227
 GIAO *see* gauge including/invariant atomic orbitals
 Gibbs free energy 472
 global minima
 diffusion methods 414
 distance geometry methods 414–15
 genetic algorithms 413–14
 molecular dynamics 412–13
 Monte Carlo methods 411–12
 optimization techniques 380–1, 409–15
 simulated annealing 413
 stochastic methods 411–12
 Gonzalez–Schlegel optimization 418
 gradient norm minimization 402–3
 gradient-corrected methods 248–56
 Gram–Schmidt orthogonalization 533
 grand unified theory 5
 gravitational interactions 4–5, 9
 Greens functions *see* propagator methods
 grid representation 539
 GROMOS force field 63–4
 GRPA *see* generalized random phase approximations
 GTO *see* Gaussian type orbitals
guache conformations 31–4
 GVB *see* generalized valence bond

 half-and-half (H + H) method 252
 half-electron method 100
 Hamilton formulation 453
 Hamiltonian operators
 density functional theory 240
 dynamical equation 7
 electron correlation methods 138–40, 159–66, 170–4, 179, 187–8
 electronic structure methods 82–4, 87, 88, 91, 104–5
 force field methods 74
 mathematical methods 525, 528–9, 538
 molecular properties 315–16, 332, 345–7
 quantum mechanics 15
 relativistic methods 283, 284, 286–7
 separation of variables 10–12
 simulation techniques 459–60, 482
 statistical mechanics 428
 superoperators 345–7
 Hammett-type effects 71
 Hammond postulate 506–10
 Hamprecht–Cohen–Tozer–Handy (HCTH)
 model 249, 251
 geometry convergence 353
 vibrational frequency convergence 360–1
 hard and soft acid and base (HSAB) principle 493–4
 harmonic expansions *see* Taylor expansions
 Hartree–Fock (HF) theory
 basis set approximation 93–8
 basis sets 213, 223, 227–9
 classical mechanics 13–14
 coupled 328
 coupled perturbed 325–9, 343
 density functional theory 233, 236–40, 248, 255–7, 259, 262–3
 dissociation curves 361–2
 electron correlation methods 133–4, 137, 138–40, 189
 electronic structure methods 80–2, 87, 91–2, 93–100
 force field methods 62, 70
 geometry convergence 353–4
 Hartree–Fock limit 97–8
 mathematical methods 541, 544
 molecular properties 322, 325–9, 339–41, 346
 numerical 93–8
 optimization techniques 380, 382
 qualitative theories 496
 quantum mechanics 18–19
 relativistic methods 288, 289

- separation of variables 9
- Slater determinants 87, 91–2
- statistical methods 551
- time-dependent 346
- total energy convergence 355
- vibrational frequency convergence 358–9
- wave function analysis 304–5
- see also* restricted Hartree–Fock; self-consistent field theory; unrestricted Hartree–Fock
- HCTH *see* Hamprecht–Cohen–Tozer–Handy
- Heisenberg uncertainty principle 20
- Heitler–London (HL) functions 269–70
- helium atoms 17–19
- Hellmann–Feynman theorem 322–3, 339–40, 572
- Helmholtz free energy
 - condensed phases 441–2
 - simulation techniques 472
 - statistical mechanics 428
- Hermitian matrices 83, 91, 516, 523
- Hermitian operators 159
- Hessian computation 385–9
- Hestenes–Stiefel (HS) method 385
- HF *see* Hartree–Fock
- higher order gradient methods 250–2, 254–6
- higher random phase approximation (HRPA) 347
- highest occupied molecular orbitals (HOMO) 488–95
- Hilbert space 530
- Hill-type potentials 36–8
- Hirshfeld atomic charges 303–4, 311
- HL *see* Heitler–London
- Hohenberg–Kohn theorem 232, 239, 571–2
- HOMO *see* highest occupied molecular orbitals
- HRPA *see* higher random phase approximation
- HS *see* Hestenes–Stiefel
- HSAB *see* hard and soft acid and base principle
- Hückel theory 107, 127–9
- hybrid force field electronic structure methods 74–7
- hybrid GGA methods 252–6
- hydrogen bonds 39–40
- hydrogen shifts 504, 580–2
- hydrogen-like atoms 14–17
- Hylleras type wave functions 179
- hyper-GGA methods 252–6
- hyperconjugation 48
- hypercubes 516
- hyperpolarizability 317
- idempotent density matrices 103–4
- IGLO *see* individual gauge for localized orbitals
- ill-conditioned systems 520
- imaginary numbers 515
- independent-particle models 80–1
- individual gauge for localized orbitals (IGLO) 338
- INDO *see* intermediate neglect of differential overlap
- infrared (IR) absorption 319
- initial guess orbitals 107
- interactions
 - description 3–4
 - fundamental forces 4–5
- interelectronic distance 178–81
- intermediate neglect of differential overlap (INDO) approximation 107, 117, 118
- internal electric fields 329
- internal magnetic moment *see* nuclear magnetic moment
- intrinsic activation energy 507–8
- intrinsic reaction coordinates (IRC)
 - mathematical methods 528
 - optimization techniques 395, 416–18
 - simulation techniques 461, 463
- introductory material *see* theoretical chemistry
- intruder states 166
- inverse matrices 519–20, 524–5
- IR *see* infrared
- IRC *see* intrinsic reaction coordinates
- isodesmic reactions 221–2
- isogyric reactions 221–2
- isomers 374–8
- jackknife models 554–5
- Jacobi method 524
- Jacob’s ladder classification 246–55
- Janak theorem 257
- Jastrow factors 189
- k-nlmG basis sets 203
- Keal–Tozer (KT) functionals 250
- kinetic balance condition 288
- Kirkwood model 480–1, 483
- Kirkwood–Westheimer model 481
- Kohn–Sham (KS) theory 235–6, 239, 257–8, 260–3
- Koopmans’ theorem 92–3, 99, 493–4
- KS *see* Kohn–Sham
- KT *see* Keal–Tozer
- Lagrange techniques
 - constrained sampling methods 464
 - electronic structure methods 90–1, 98, 102
 - extended 457–9
 - force field methods 45
 - molecular properties 324–5, 328
 - optimization techniques 408–9, 418
 - simulation techniques 452–3
- Langevin methods 455, 476
- LAO *see* London atomic orbitals
- Laplace transforms 543
- Laplacians 532
- large curvature ground state (LCG) approximation 462–3
- latent variables 556

- LCAO *see* linear combination of atomic orbitals
 LCCD *see* linear coupled cluster doubles
 LCG *see* large curvature ground state
 LDA *see* local density approximation
 leap-frog algorithms 452
 least squares linear fit 551
 leave-one-out models 554–5
 Lee–Yang–Parr (LYP) model 249–50
 Legendre parameterization 199, 212
 Lennard-Jones (LJ) potential 35–8, 40, 62
 leptons 4–5
 level shifting
 Newton–Raphson methods 387
 self-consistent field theory 101–2
 level/basis notation 137
 LIE *see* linear interaction energy
 Lieb–Oxford condition 245
 line-then-plane (LTP) optimization 398
 linear combination of atomic orbitals (LCAO) 94
 linear correlation 551
 linear coupled cluster doubles (LCCD) 176
 linear interaction energy (LIE) method 475
 linear synchronous transit (LST) 395
 linearised Poisson–Boltzmann equation (LPBE) 479
 LJ *see* Lennard-Jones potential
 LMOs *see* localized molecular orbitals
 local density approximation (LDA) 245, 246–8, 254, 257
 local minima 380–1, 383–90
 local spin density approximation (LSDA)
 basis sets 225
 density functional methods 246–8, 251, 255–6, 258, 263–4
 dipole moment convergence 357–8
 geometry convergence 353
 vibrational frequency convergence 360–1
 localized molecular orbitals (LMOs) 304–8
 localized orbital methods 182–3
 localized orbital/local origin (LORG) 338
 locally updated planes (LUP) optimization 400
 London atomic orbitals (LAO) 338–9
 London forces 35
 long-range solvation 475–6
 looping 2
 loosely bound electrons 258
 Lorentz transformations 277
 LORG *see* localized orbital/local origin
 Löwdin partitioning 294–6, 310, 311
 lowest unoccupied molecular orbitals (LUMO) 488–95, 543–6
 LPBE *see* linearised Poisson–Boltzmann equation
 LSDA *see* local spin density approximation
 LST *see* linear synchronous transit
 LTP *see* line-then-plane
 LUMO *see* lowest unoccupied molecular orbitals
 LUP *see* locally updated planes
 LYP *see* Lee–Yang–Parr
 MacDonald’s theorem 571
 McWeeny procedure 104
 MAD *see* mean absolute deviation
 magnetic fields
 diamagnetic contribution 335
 external 315, 318, 331–2
 gauge dependence 338–9
 molecular properties 315, 318–19, 329–39
 nuclear magnetic moment 315, 318–19, 332
 paramagnetic contribution 335
 magnetic potentials 282–4
 magnetizability 318, 335
 many-body perturbation theory (MBPT) 137, 159–69, 183–6
 beryllium atom 177–8
 configuration interaction 174–8, 183–5
 coupled cluster theory 174–8
 Møller–Plesset perturbation theory 162–9, 174–8
 projected Møller–Plesset methods 168–9
 unrestricted Møller–Plesset methods 168–9
 many-body problem 2
 Marcus equation 71, 506–10
 mass-polarization term 83–5
 mass-velocity correction 281–2
 mathematical methods 514–46
 approximating functions 538–41
 basis set expansion 539, 541
 computational considerations 529
 coordinate transformations 520–9
 differential equations 535–8
 differential operators 531–2
 Fourier transformations 541–2
 functionals 530
 functions 530, 531
 Laplace transforms 543
 matrices 516–20, 523–4
 normalization 532–5
 numbers 514–15
 operations 2
 operators 530–2
 orthogonalization 533–5
 projection 534–5
 Slater determinants 528–9
 surfaces 543–6
 vectors 514–15, 517, 532
 matrices 516–20
 determinants 518–19
 eigenvalues/eigenvectors 523–4, 526–8
 inverses 519–20, 524–5
 multiplications 516–18
 rank 524
 transpositions 516
 Z-matrix construction 575–82
 matrix elements 82

- MAXI basis sets 205
MBPT *see* many-body perturbation theory
MC *see* Monte Carlo
MCM *see* multi-configurations molecular mechanics
MCRPA *see* multi-configuration random phase approximation
MCSCF *see* multi-configuration self-consistent field
MD *see* molecular dynamics
mean absolute deviation (MAD) 550–1
mean values 549
mean-field approximations *see* Hartree–Fock theory
mechanical embedding 74–5
median 550
MEP *see* minimum energy path; molecular electrostatic potential
MEPSAC *see* minimum energy path semi-classical adiabatic ground state
Merck molecular force field (MMFF) 35–6
meta-GGA methods 250–2, 254–6
metal coordination compounds *see* coordination compounds
methyl shifts 505
Metropolis algorithms 188
microcanonical transition state theory 424–5
MIDI basis sets 205
migrations 504
MINDO *see* modified intermediate neglect of differential overlap
MINI basis sets 205
minimum basis set 194
minimum energy path (MEP) 417, 461
minimum energy path semi-classical adiabatic ground state (MEPSAC) 462
minimum energy structures 70–3
mixed derivatives 319–20
MLR *see* multiple linear regression
MM (molecular mechanics) *see* force field methods
MMFF *see* Merck molecular force field
MNDO *see* modified neglect of diatomic overlap
mode 550
modified intermediate neglect of differential overlap (MINDO) approximation 119
modified NDDO approximations 119–20
modified neglect of diatomic overlap (MNDO) 121–7, 130
MOJ *see* More O’Farrell–Jencks
molecular docking 415–16
molecular dynamics (MD) 445–8, 451–4
condensed phases 440–3
constrained sampling methods 464
extracting information from simulations 468–9
global minima 412–13
molecular electrostatic potential (MEP) 42, 296
molecular mechanics (MM) *see* force field methods
molecular orbital theory *see* frontier molecular orbital theory; qualitative molecular orbital theory
molecular properties 315–49
basis sets 348–9
classical terms 333–4
derivative techniques 321–4
electron spin 333
examples 316–20
external electric fields 315, 316–17, 329
external magnetic fields 315, 318, 331–2
gauge dependence 338–9
internal electric fields 329
Lagrangian techniques 324–5, 328
magnetic field perturbations 315, 318–19, 329–39
mixed derivatives 319–20
nuclear geometry perturbations 315, 319, 339–43
nuclear magnetic moment 315, 318–19, 332
perturbation methods 321
propagator methods 343–8
relativistic methods 324
response methods 343–8
Møller–Plesset perturbation theory 158, 162–9, 183–6
beryllium atom 177–8
configuration interaction 174–8, 183–5
coupled cluster theory 174–8
dipole moment convergence 357, 372
geometry convergence 350–3
projected methods 168–9
total energy convergence 354–6
unrestricted methods 168–9
vibrational frequency convergence 359–60, 373
Monte Carlo (MC) methods 445–50
condensed phases 440–3
constrained sampling methods 464
density functional theory 247
extracting information from simulations 468–9
global minima 411–12
non-natural ensembles 450
see also quantum Monte Carlo
More O’Farrell–Jencks (MOJ) diagrams 510–12
Morokuma energy decomposition 496–7
Morse potentials
force field methods 25–7, 36–8, 47, 59
mathematical methods 540, 541
MRCI *see* multi-reference configuration interaction
Mulliken
electronegativity 493
notation 96
population analysis 294–6, 310–12

- multi-configuration random phase approximation (MCRPA) 347
- multi-configuration self-consistent field (MCSCF)
 - electron correlation methods 153–8, 187
 - molecular properties 322
 - qualitative theories 505
 - valence bond theory 273–5
- multi-configurations molecular mechanics (MCMM) method 73
- multi-determinant wave functions
 - electron correlation methods 134–5
 - electronic structure methods 81
- multi-dimensional energy surfaces 381
- multi-reference configuration interaction (MRCI) 158–9, 456
- multi-reference wave functions 157
- multiple linear regression (MLR) 555–6, 558–9
- multiple-descriptor data sets 553–5
- multipoles 43–7

- N*-order tensors 516
- N*-representability 238, 240
- natural atomic orbital (NAO) analysis 309–12
- natural bond orbital (NBO) analysis 309–12
- natural geminals 308–9
- natural internal coordinates 393–4
- natural orbitals (NO) 308–9
- NBO *see* natural bond orbital
- NDDO *see* neglect of diatomic differential overlap
- NEB *see* nudged elastic band
- neglect of diatomic differential overlap (NDDO) approximation 116–17, 118, 119–20, 130
- neighbour lists 65
- Newton formulation 453
- Newton–Raphson (NR) methods
 - electron correlation methods 154
 - electronic structure methods 103–4
 - mathematical methods 540
 - minima 385–94
 - optimization techniques 385–94, 403–5
 - saddle points 403–5
- Newtonian mechanics 5–8, 12, 22
- NMR *see* nuclear magnetic resonance
- NO *see* natural orbitals
- noise 548
- non-adiabatic coupling elements 84
- non-bonded energies *see* electrostatic energies; van der Waals energies
- non-degenerate eigenvalues 523
- non-linear correlations 552–3
- non-natural ensembles 450, 454–5
- non-specific solvation 475–6
- norm-conserving pseudopotentials 224
- norm-extended Hessian 387
- normalization 532–5
- Nosé–Hoover methods 455

- notation (Appendix A) 565–9
- NR *see* Newton–Raphson
- nuclear geometry
 - perturbations 315, 319, 339–43
 - transition state theory 423–4
- nuclear magnetic moment 315, 318–19, 332
- nuclear magnetic resonance (NMR) 320, 337
- nucleophilic reactions 489, 492
- nudged elastic band (NEB) optimization 400–1
- numbers 514–15
- numerical Hartree–Fock methods 93

- occupation numbers 308–9
- OEP *see* optimized effective potential
- one-centre one-electron integrals 119
- one-electron integrals 97, 116–18
- ONIAM *see* our own *n*-layered integrated molecular orbital molecular mechanics
- Onsager model 480–1, 482, 483
- operators 530–2
- optimization techniques 380–420
 - conformational sampling 409–15
 - conjugate gradient methods 384–5
 - constrained optimization 407–9
 - coordinate selection 390–4, 405–6
 - GDIIS extrapolations 389–90
 - general functions 381, 383–407
 - global minima 380–1, 409–15
 - Hessian computation 385–9
 - intrinsic reaction coordinate methods 395, 416–18
 - local minima 380–1, 383–90
 - molecular docking 415–16
 - Newton–Raphson methods 385–94
 - quadratic functions 380–2
 - saddle points 381, 394–407
 - Simplex method 383
 - steepest descent method 383–4
 - step control 386–7
- optimized effective potential (OEP) methods 253–4
- optimized exchange (OPTX) model 249–50
- orbital
 - controlled reactions 488
 - correlation diagrams 497–8, 501–3
- orbital-free density functional theory 233–5
- orbital-Zeeman term 331–2
- ortho* conformations 32
- orthogonalization 533–5
- our own *n*-layered integrated molecular orbital molecular mechanics (ONIAM) method 76
- out-of-plane bending energies 30
- outer products 518
- overlap elements 82

- pairwise distance corrected Gaussian (PDDG) approximation 123–4
- paramagnetic spin–orbit (PSO) operator 287
- paramagnetism 335

- parameterization
 - accuracy/generalizability 53–5, 71–2
 - basis sets 199, 212
 - combined 56
 - coordination compounds 58–62
 - density functional theory 238, 247–8
 - electronic structure methods 118–25, 130
 - force field methods 51–62
 - generic parameters 57–8
 - missing parameters 54–6
 - parameter reduction in force fields 57–8
 - redundant variables 56–7
 - relativistic effects 62
 - sequential 55–6
 - universal force fields 62
 - validation of force fields 68–9
- parameterized configuration interaction (PCI-X) method 221
- parametric method number 3 (PM3) 122–4, 125–7, 130–1
- parametric method number 5 (PM5) 123–4
- Pariser–Pople–Parr (PPP) method 49–50, 118
- partial charge models 43–4
- partial least squares (PLS) 557–9
- particle mesh Ewald (PME) method 467
- partition functions 427–8
- partitioned rational function optimization (P-RFO) 404
- Pauli equation 281
- PAW *see* projector augmented wave
- PB *see* Poisson–Boltzmann
- PBE *see* Perdew–Burke–Ernzerhof; Poisson–Boltzmann equation
- PCA *see* principle component analysis
- PCI-X *see* parameterized configuration interaction
- PCM *see* polarizable continuum model
- PDDG *see* pairwise distance corrected Gaussian
- penalty function 407–8
- pentuple zeta (PZ) basis sets 195
- Perdew–Burke–Ernzerhof (PBE)
 - basis sets 225
 - density functional methods 249, 255–6
 - dipole moment convergence 357–8
 - geometry convergence 353
 - vibrational frequency convergence 360–1, 374
- Perdew–Kurth–Zupan–Blaha (PKZB)
 - functional 252
- Perdew–Wang (PW) formula 247, 249, 255–6
- perfect pairing (PP) 275
- pericyclic reactions 505–6
- periodic boundary conditions 464–8
- periodic systems 113–15
- PES *see* potential energy surfaces
- PGTO *see* primitive Gaussian type orbitals
- phase space 428
- photochemical reactions
 - qualitative theories 499, 500
 - transition state theory 423–4
- physicist's notation 95–6
- Pipek–Mezey localization scheme 306–8
- PKZB *see* Perdew–Kurth–Zupan–Blaha
- plane wave basis functions 211–12
- PLS *see* partial least squares
- PM3 *see* parametric method number 3
- PM5 *see* parametric method number 5
- PME *see* particle mesh Ewald
- PMF *see* potential of mean force
- points-on-a-sphere (POS) models 61
- Poisson–Boltzmann equation (PBE) 478
- Poisson–Boltzmann (PB) methods 478–9
- Polak–Ribiere (PR) method 385
- polarizability 43–7, 317
- polarizable continuum model (PCM) 483
- polarizable embedding 75
- polarization 195, 196–8, 206, 228
- polarization consistent basis sets 207–8
- polarization propagators (PP) 344–5
- Pople style basis sets 202–4
- population analysis 293–304
 - atoms in molecules method 299–303
 - basis functions 293–6
 - electron density 299–304
 - electrostatic potential 296–9
 - generalized atomic polar tensor charges 304, 311
 - Hirshfeld atomic charges 303–4, 311
 - Mulliken 294–6, 310–12
 - Stewart atomic charges 304, 311
 - Voronoi atomic charges 303, 311
- POS *see* points-on-a-sphere
- potential energy surfaces (PES) 11, 19–21
 - electronic structure methods 80, 84–5
 - force field methods 22
 - simulation techniques 459–60, 469–70
- potential of mean force (PMF) 464
- potentials 4–5
- Powell method 388
- PP *see* perfect pairing; polarization propagators; pseudopotential
- PPP *see* Pariser–Pople–Parr
- PR *see* Polak–Ribiere
- pre-NAOs 310
- precision 550
- predicted residual sum of squares (PRESS) 554
- predictive correlation coefficients 554
- preservation of bonding 503–4
- PRESS *see* predicted residual sum of squares
- primitive Gaussian type orbitals (PGTO) 200–6
- principal axes of inertia 431
- principle component analysis (PCA)
 - mathematical methods 523
 - statistical methods 556–7, 558–9
- principle propagator 346
- probabilistic equations 6
- projected Møller–Plesset methods 168–9, 364–7

- projected unrestricted Hartree–Fock (PUHF) 152–3, 363–5
projection 534–5
projector augmented wave (PAW) method 225
propagator methods 343–8
pseudo-atoms 39, 59
pseudo-Newton–Raphson methods 388
pseudopotential (PP) 222–4
pseudospectral methods 227–9
PSO *see* paramagnetic spin–orbit operator
PUHF *see* projected unrestricted Hartree–Fock
PW *see* Perdew–Wang
PZ *see* pentuple zeta
- QA *see* quadratic approximation
QCI *see* quadratic configuration interaction
QED *see* quantum electrodynamics
QM/MM *see* quantum mechanics – molecular mechanics methods
QMC *see* quantum Monte Carlo
QSAR *see* quantitative structure–activity relationships
QST *see* quadratic synchronous transit
quadratic approximation (QA) method 387, 404
quadratic configuration interaction (QCI) 176, 215–18
quadratic functions 380–2
quadratic synchronous transit (QST) 395–6
quadruple zeta (QZ) basis sets 195
quadruple zeta valence (QZV) basis sets 205
quadrupole–quadrupole interactions 35, 46
 see also multipoles
qualitative molecular orbital theory 494–7
qualitative theories 487–513
 Bell–Evans–Polanyi principle 506–10
 density functional theory 492–4
 frontier molecular orbital theory 487–92
 Hammond postulate 506–10
 Marcus equation 506–10
 More O’Farrell–Jencks diagrams 510–12
 qualitative molecular orbital theory 494–7
 Woodward–Hoffmann rules 497–506, 508
quality analysis 553–5
quantitative structure–activity relationships (QSAR) 559–61
quantum electrodynamics (QED) 5, 285
quantum mechanics 6–7, 14–19
quantum mechanics – molecular mechanics (QM/MM) methods 74–7, 476
quantum methods 459–60
quantum Monte Carlo (QMC) methods 187–9
quarks 4–5
quaternions 515
QZ *see* quadruple zeta
QZV *see* quadruple zeta valence
- radial distribution functions 470–1
radial functions 16–17
radiative transitions 423
radius of convergence 540
Raman absorption 319
random errors 547–8
random phase approximation (RPA) 346, 347
RASSCF *see* restrictive active space self-consistent field
rational function optimization (RFO) 387, 404
RATTLE algorithm 453
Rayleigh–Schrödinger perturbation theory 161, 321
RCA *see* relaxed constraint algorithm
reaction field model 476–7
reaction path (RP) methods 460–5
reactive energy surfaces 73–4
read/write data function 2
ReaxFF method 73–4
reciprocal cells 113
reduced
 density matrix methods 236–40
 scaling techniques 110–13
redundant variables 56–7
relationship determination 2
relative energies of isomers 374–8
relativistic methods 277–92
 Dirac equation 278–88
 effects of relativity 289–92
 electric potentials 280–4
 equations 6
 four-component calculations 287–8
 geometry convergence 352
 magnetic potentials 282–4
 many-particle systems 284–7
 molecular properties 324
 singularities 288
relaxed constraint algorithm (RCA) 104
renormalized Davidson correction 175
resolution of the identity 180–1, 183, 534–5
resonance energy 273
RESP *see* restrained electrostatic potential
response methods 343–8
restrained electrostatic potential (RESP) 42
restricted Hartree–Fock (RHF) methods 99–100
 configuration interaction 145–53
 dissociation 145–53
 electron correlation methods 133, 145–53, 154, 157, 168–9
 Møller–Plesset perturbation theory 168–9
restricted Møller–Plesset methods 364–7
restricted open-shell Hartree–Fock (ROHF) methods 99–100
 electron correlation methods 133, 150, 152, 168–9, 176
restrictive active space self-consistent field (RASSCF) 155–6
RFO *see* rational function optimization
RHF *see* restricted Hartree–Fock theory
RHF dissociation 363–7

- Rice–Ramsperger–Kassel–Marcus (RRKM) theory 424–5
- ridge optimization 398
- rigid-rotor harmonic-oscillator (RRHO) approximation 429–39
- bimolecular reactions 434–6
 - electronic degrees of freedom 433
 - enthalpy and entropy contributions 429, 433–9
 - rotational degrees of freedom 430–1
 - transition states 436–9
 - translational degrees of freedom 430
 - unimolecular reactions 434, 435–6
 - vibrational degrees of freedom 431–3
- ring critical points 302
- ring-closures 500–4
- ring-opening reactions 508–12
- RK *see* Runge–Kutta
- RMS *see* root mean square
- ROHF *see* restricted open-shell Hartree–Fock theory
- root mean square (RMS) 550
- Roothaan–Hall equations
- electronic structure methods 94, 96–7, 100
 - relativistic methods 289
- rotational degrees of freedom 430–1
- RP *see* reaction path
- RPA *see* random phase approximation
- RRHO *see* rigid-rotor harmonic-oscillator approximation
- RRKM *see* Rice–Ramsperger–Kassel–Marcus
- Rumer basis 272
- Runge–Kutta (RK) algorithm 417, 452
- Rydberg orbitals 310
- Rydberg states 187, 259
- SAC *see* scaled all correlation; spin-adapted configurations
- saddle optimization 397
- saddle points
- coordinate selection 405–6
 - dimer method 405
 - dynamic methods 406–7
 - gradient norm minimization 402–3
 - interpolation methods 394–402
 - local methods 402–6
 - multi-structure interpolation methods 398–401
 - Newton–Raphson methods 403–5
 - one-structure interpolation methods 394–7
 - optimization techniques 381, 394–407
 - transition state theory 422
 - TS modelling 70–4
 - two-structure interpolation methods 397–8
- SAM1/SAM1D *see* semi *ab initio* method 1
- SAS *see* solvent accessible surfaces
- scalar functions 531
- scalar relativistic corrections 281–2
- scalars 514
- scaled all correlation (SAC) 219, 221
- scaled external correlation (SEC) 219, 221
- scaling
- electron correlation methods 184, 189
 - self-consistent field theory 110–13
- SCF *see* self-consistent field
- Schrödinger equation 6–8, 10–12, 15–20
- adiabatic approximation 84–5
 - basis sets 192, 211
 - Born–Oppenheimer approximation 80, 82–6
 - density functional theory 238, 240
 - Dirac equation 278, 280–4
 - electron correlation methods 159–66, 170–4, 178–9, 187–8
 - electronic structure methods 80–92
 - force field methods 22
 - mathematical methods 526–8, 538
 - molecular properties 315–16
 - relativistic methods 277, 278, 280–4
 - rigid-rotor harmonic-oscillator approximation 432
 - self-consistent field theory 87, 92
 - Slater determinants 81, 87–92
 - statistical mechanics 429
- Schwartz inequality 109
- SCRf *see* self-consistent reaction field
- SCSAC *see* small curvature semi-classical adiabatic ground state
- SCVB *see* spin-coupled valence bond
- SD *see* Slater determinants; steepest descent
- SEAM method 71–3
- SEC *see* scaled external correlation
- second-order differential equations 536–8
- second-order perturbation theory 487
- second-order polarization propagator approximation (SOPPA) 347–8
- segmented contraction 201–2
- self-consistent field (SCF) theory
- basis set approximation 96–7
 - convergence 101–4
 - damping 101
 - direct inversion in the iterative subspace 102–3, 104
 - direct minimization techniques 103–4
 - direct SCF 108–10
 - electron correlation methods 152, 181–2
 - electronic structure methods 86–7, 92, 96–7, 100–13
 - extrapolation 101
 - initial guess orbitals 107
 - level shifting 101–2
 - minimum HF energies 105–7
 - molecular properties 322, 342
 - processing time 110
 - qualitative theories 505
 - reduced scaling techniques 110–13
 - Slater determinants 92
 - symmetry 104–5

- techniques 100–13
- valence bond theory 273–5
- self-consistent Hückel methods 128
- self-consistent reaction field (SCRf) models 481–4
- self-penalty walk (SPW) optimization 398–9
- semi *ab initio* method 1 (SAM1/SAM1D) 124–5, 127
- semi-empirical electronic structure methods 115–18
 - advantages 129–31
 - limitations 129–31
 - parameterization 118–25, 130
 - performance 125–7
- separation of variables 8–12
- separability theorem 525
- sequential parameterization 55–6
- SHAKE algorithm 453, 458
- shifting function approach 465–6
- short-range solvation 475–6
- similarity transformations 171, 522
- simple harmonic expansions *see* Taylor expansions
- simple Hückel theory 128–9
- Simplex method 383
- simulated annealing 413
- simulation techniques 445–86
 - Born/Onsager/Kirkwood models 480–1, 483
 - Car–Parrinello methods 457–9, 476
 - constrained sampling methods 463–4
 - continuum models 476–84
 - direct methods 455–7
 - extracting information from simulations 468–72
 - free energy methods 472–5
 - Langevin methods 455, 476
 - non-Born–Oppenheimer methods 463
 - non-natural ensembles 450, 454–5
 - periodic boundary conditions 464–8
 - Poisson–Boltzmann methods 478–9
 - potential energy surfaces 459–60
 - quantum methods 459–60
 - reaction path methods 460–5
 - self-consistent reaction field models 481–4
 - solvation methods 475–84
 - thermodynamic integration 473–5
 - thermodynamic perturbation 472–3
 - time-dependent methods 450–64
 - see also* molecular dynamics; Monte Carlo methods
- SINDO *see* symmetric orthogonalized intermediate neglect of differential overlap
- single value decomposition 520, 524
- singlet instability 106
- singularities 288
- size consistency 153
- size extensivity 153
- Slater determinants (SD) 18
 - electron correlation methods 135–7, 139, 145, 163–5, 178
 - electronic structure methods 81, 87–92, 99
 - excited 135–7, 139, 146, 163–5
 - mathematical methods 528–9
 - optimization techniques 380
 - valence bond theory 268, 269, 274
 - wave function analysis 304–5
- Slater type orbitals (STO) 192–4, 202–4
- Slater–Condon rules 140
- Slater–Kirkwood equation 53
- slow growth method 473
- SM *see* string method
- small curvature semi-classical adiabatic ground state (SCSAC) 462–3
- small rings 48–50
- SO *see* strong orthogonality
- solar system 13–14, 18–19
- solvation models
 - Born/Onsager/Kirkwood models 480–1, 483
 - continuum models 476–84
 - Poisson–Boltzmann methods 478–9
 - self-consistent reaction field models 481–4
 - simulation techniques 475–84
- solvent accessible surfaces (SAS) 477–8, 480
- SOPPA *see* second-order polarization propagator approximation
- SOS *see* sum over states
- special relativity 279
- specific reaction parameterization (SRP) 456–7
- specific solvation 475–6
- sphere optimization 396–7
- spherical
 - harmonic functions 16–17
 - polar systems 514–15
- spin contamination 148–53
- spin-adapted configurations (SAC) 139
- spin-coupled valence bond (SCVB) theory 270–5
- spin-Zeeman term 331–2, 333, 335, 337
- spinors 287–8
- split valence basis sets 195, 205
- spread 549
- SPW *see* self-penalty walk
- SRP *see* specific reaction parameterization
- standard deviation 549
- standard model 5
- starting condition 3
- state correlation diagrams 497–9, 501–3
- state-averaged multi-configuration self-consistent field 187
- state-selected configuration interaction 159
- stationary orbits 13
- statistical mechanics 426–9
- statistical methods 547–59
 - correlation between many data sets 553–9
 - correlation between two data sets 550–3
 - elementary measures 549–50

- errors 547–9
- illustrative example 558–9
- multiple linear regression 555–6, 558–9
- multiple-descriptor data sets 553–5
- partial least squares 557–9
- principle component analysis 556–7, 558–9
- quality analysis 553–5
- steepest descent (SD) method 383–4
- step control 386–7
- step-and-slide optimization 398
- steric energy 50
- Stewart atomic charges 304, 311
- STO *see* Slater type orbitals
- STO-nG basis sets 202, 204
- stochastic dynamics 455
- stochastic methods 411–12
- stockholder atomic charges 303–4, 311
- STQN *see* synchronous transit-guided quasi-Newton
- stretch energies 26–7, 64–5
- string method (SM) 401
- string theory 5
- strong orthogonality (SO) condition 275
- structural units *see* functional groups
- substitution reactions
 - qualitative theories 489, 509
 - rigid-rotor harmonic-oscillator approximation 435–6
- sum over states (SOS) methods 321
- Sun–Earth system 12–13, 16
- supercell approach 212
- superoperators 345–7, 530
- superposition errors 225–7
- switching function approach 465–6
- symmetric orthogonalized intermediate neglect
 - of differential overlap (SINDO) approximation 118
- symmetrical orthogonalization 533
- symmetry 104–5
- symmetry-breaking phenomena 106–7
- synchronous transit-guided quasi-Newton (STQN) 395–6
- system description 3
- systematic errors 547–8
- Tao–Perdew–Staroverov–Scuseria (TPSS)
 - functional 252, 253
- Taylor expansions
 - density functional theory 234, 249
 - force field methods 24–8, 47, 58, 59
 - mathematical methods 526, 539–40, 541
 - simulation techniques 451
 - statistical methods 553
- TDDFT *see* time-dependent density functional theory
- TDHF *see* time-dependent Hartree–Fock
- tensors 516
- TF *see* Thomas–Fermi
- TFD *see* Thomas–Fermi–Dirac
- theoretical chemistry
 - chemistry 19–21
 - classical mechanics 6, 12–14
 - definitions 1–2
 - dynamical equations 3, 4, 5–12
 - fundamental forces 4–5
 - fundamental issues 2–3
 - quantum mechanics 6–7, 14–19
 - system description 3–4
- thermal decomposition 425–6
- thermal reactions
 - qualitative theories 499
 - transition state theory 423–4
- thermodynamic
 - cycles 474–5
 - integration 473–5
 - perturbation 472–3
- Thomas–Fermi (TF) theory 234
- Thomas–Fermi–Dirac (TFD) model 234, 247
- time-dependent density functional theory (TDDFT) 346
- time-dependent Hartree–Fock (TDHF) 346
- torsional energies 30–4, 42–3, 48, 57, 63
- total energy convergence 354–6
- TPSS *see* Tao–Perdew–Staroverov–Scuseria
- trans* effect 59, 60
- trans* isomerism 32
- transferability 43
- transition state theory (TST) 421–6
 - dynamical effects 425–6
 - Rice–Ramsperger–Kassel–Marcus theory 424–5
 - rigid-rotor harmonic-oscillator approximation 436–9
 - variational 438
 - see also* frontier molecular orbital theory
- transition structures *see* saddle points
- translational degrees of freedom 430
- transmission coefficient 422
- transoid* configurations 41
- triple zeta plus double polarization (TZ2P) 196–7
- triple zeta plus polarization (TZP) 214, 225
 - geometry convergence 354
 - vibrational frequency convergence 360
- triple zeta (TZ) basis sets 195–7
- triple zeta valence (TZV) basis sets 205
- triplet instability 106
- truncated configuration interaction (CI)
 - methods 143–4
- truncated coupled cluster (CC) methods 172–4
- truncation errors 548–9
- trust radius 386, 404
- TS *see* transition structure
- TST *see* transition state theory
- turnover rule 160–1
- two-centre one-electron integrals 120
- two-centre two-electron integrals 120
- two-electron integrals 82, 116–18

- TZ *see* triple zeta
 TZ2P *see* triple zeta plus double polarization
 TZV *see* triple zeta valence
- UFF *see* universal force fields
 UHF *see* unrestricted Hartree–Fock theory
 ultrasoft pseudopotentials 224
 umbrella sampling 464
 unbound solutions 13, 17
 unit cells 113
 unitary matrices 519
 unitary transformations 526
 united atom approach 64
 universal force fields (UFF) 62
 unrestricted Hartree–Fock (UHF) methods 99–100
 configuration interaction 148–53
 dissociation 148–53, 363–7
 electron correlation methods 133, 148–53, 154, 157, 168–9
 Møller–Plesset perturbation theory 168–9
 spin contamination 148–53
 unrestricted Møller–Plesset methods 168–9, 364–7
 Urey–Bradley force fields 42
- valence bond (VB) theory 268–76
 benzene 272–4
 classical 269–70
 generalized 275
 resonance energy 273
 spin-coupled 270–5
 valence shell electron-pair repulsion (VSEPR) model 61
 van der Waals energies
 force field methods 34–40, 42–3, 52–3, 57, 61, 65–7
 mathematical methods 545
 simulation techniques 471, 476
 surfaces 477–8, 480, 545
 variable metric methods 388
 variance 549
 variational
 principle 570–1
 problem 98–9
 quantum Monte Carlo 188
 variational transition state theory (VTST) 438
 VB *see* valence bond
 vectors 514–15, 517, 532
 velocity Verlet algorithms 452, 453
 Verlet algorithms 8, 451–3, 458
 very fast multipole moment (vFMM) method 111, 467
 vibrational degrees of freedom 431–3
 vibrational frequency convergence 358–61
 Ab initio methods 358–60
 density functional theory 360–1, 374
 problematic systems 373–4
 vibrational normal coordinates 19, 526–8
 von Weizsacker kinetic energy 234
 Voorhis–Scuseria exchange–correlation (VSXC) 251–2, 263
 Voronoi atomic charges 303, 311
 Vosko–Wilk–Nusair (VWN) formula 247
 VSEPR *see* valence shell electron-pair repulsion
 VSXC *see* Voorhis–Scuseria
 exchange–correlation
 VTST *see* variational transition state theory
 VWN *see* Vosko–Wilk–Nusair
- W–H *see* Woodward–Hoffmann
 Wannier orbitals 306
 wave function analysis 293–314
 atoms in molecules method 299–303
 basis functions 293–6
 computational considerations 306–8, 311–12
 critical points 302
 electron density 299–304
 electrostatic potential 296–9, 312
 examples 312–13
 generalized atomic polar tensor charges 304, 311
 Hirshfeld atomic charges 303–4, 311
 localized molecular orbitals 304–8
 natural atomic orbitals 309–12
 natural orbitals 308–9
 population analysis 293–304, 311–12
 Stewart atomic charges 304, 311
 Voronoi atomic charges 303, 311
 wave packages 459
 weak interactions 258
 well-tempered basis sets 198–200
 width (data) 549
 Wigner correction 462
 Wigner intracule 239
 Woodward–Hoffmann (W–H) rules 497–506, 508
 write/read data function 2
- Z-matrix construction 575–82
 Z-vector method 324
 ZDO *see* zero differential overlap
 Zeeman interactions 283–4, 286
 zero differential overlap (ZDO) approximation 116
 zero field splitting (ZFS) 336
 zeroth-order regular approximation (ZORA) method 282
 ZFS *see* zero field splitting
 ZORA *see* zeroth-order regular approximation
 Zwitterbewegung 281