APPLICATIONS OF SUM OF SQUARES DECOMPOSITIONS IN THEORETICAL CHEMISTRY

Martin Gerard Burke

Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

Imperial College London
Department of Chemistry
June 2007

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that the copyright rests with its author.

I hereby declare that the work presented in this thesis is entirely my own except where other publications or people have been cited.

Martin Burke.

Date.

Abstract

A recently developed approach to polynomial optimisation which produces efficient (and in many cases exact) relaxations to NP-hard problems [1], has been applied to two problems within theoretical chemistry. The approach differs fundamentally from the standard approaches in a number of ways: i) By lifting the domain space to a higher dimension and solving in the coefficient space instead of the coordinate space, the problem can be expressed as a convex problem of finding a sum of squares decomposition which is a lower bound to the original function. This is achieved in practice by solving a semidefinite programme for which polynomial time interior point methods have been very successful. ii) The approach completely avoids sampling or deforming the original function, at least one of which is necessary in all other approaches. iii) Proofs of optimality are produced by exploiting the duality properties of convex programming.

The first problem addressed using this approach is that of finding the global minimum of a potential energy surface for a system of molecules. This is a long studied problem using techniques drawn from stochastic and deterministic optimisation. The advantages described above are highlighted. This new approach was demonstrated for a number of model potentials with pathological characteristics, and the performance of the method was compared with sampling approaches. The method was then applied to a small molecular system modelled using a standard force field (Amber94).

The second problem is that of coarse-graining of molecular systems using ellipsoidal particles and potentials. Such potentials can be based on an approximation to the minimum distance between the particles [2]. Noting the equivalence between a sum of squares polynomial and the functional form of the ellipsoidal particle, coarse-grained molecular dynamics was performed whereby a conic programme was solved at each step in order to calculate the exact minimum distance and this was used to calculate the interparticle forces. To the author's knowledge, this is the first application of coarse-grained molecular dynamics in which the true minimum distance is employed.

To my parents

Acknowledgements

My deepest gratitude must go to my supervisor, Dr. Sophia Yaliraki first for giving me the opportunity to work in her group during both my Masters and PhD studies and second for her guidance and support over the last five years. I am also extremely grateful to my parents, Stephen (deceased) and Mary Burke, for their love, encouragement and support and to my two brothers Conal and Stephen.

I am also very grateful to Claire, my girlfriend for her support and understanding of long hours spent in the lab. Thanks to all the members of Sophia's group, Ella, Joao, Leonid, and Tim as well as past members and members of the Gould group, for their help and support in my research as well as their very genuine friendship. It has been a pleasure to work with such fantastic colleagues.

I must also thank Professor Pablo Parrilo, Dr. Mauricio Barahona and Dr. Antonis Papachristodoulou for helpful discussions on the topics of sum of squares decompositions and convex optimisation.

Finally I wish to acknowledge the financial support I received from the Biotechnology and Biological Sciences Research Council, U.K. and the Office of Naval Research, U.S.A.

Contents

Contents **6**

List of Tables **10**

List of Figures **11**

Publications **13**

1 Introduction **14**

1.1 Energy landscapes background 14

1.1.1 Free and potential energy landscapes 15

1.1.2 Calculating potential energy surfaces 16

1.2 The Born-Oppenheimer Approximation 18

1.3 Categorising and describing energy landscapes 19

1.3.1 Stationary points and their significance 19

1.3.2 Steepest descent pathways 20

1.4 Quantification and topological considerations for stationary points . . 21

1.4.1 Topological categorisation 24

1.4.2 Disconnectivity graphs 24

1.5 Approaches for sampling the potential energy landscape 26

1.5.1 Identifying local minima 27

1.5.2 Identifying transition states 28

1.5.3 Sampling thermodynamic properties on the PES 30

1.5.4 Calculating thermodynamic properties from a PES minima database 31

1.6 The global optimisation problem 33

1.6.1 Complexity studies of global optimisation on the PES 33

1.6.2 Potential energy landscapes and NP-hardness 38

1.7 State of the art approaches for global optimisation on a PES 39

1.7.1 MD based annealing 41

1.7.2 Hypersurface deformation methods 42

1.7.3 Genetic algorithms 46

1.7.4 Hybrid algorithms . . . 47
1.7.5 Deterministic methods . . . 48
1.8 Conclusions and objectives . . . 51

2 Sum of Squares Optimisation 54
2.1 Introduction . . . 54
2.2 Convex Optimisation . . . 55
2.2.1 Preliminaries . . . 55
2.2.2 Optimisation problems . . . 59
2.2.3 Duality theory for convex optimisation . . . 62
2.2.4 Optimality criteria . . . 64
2.2.5 Extending Lagrangian duality to generalised inequalities . . . 66
2.3 Global optimisation and global nonnegativity . . . 68
2.3.1 Polynomials, nonnegative polynomials and sum of squares . . . 68
2.3.2 The positive semidefinite and sum of squares cones . . . 69
2.3.3 Relaxation of the global nonnegativity condition . . . 71
2.3.4 Sum of squares decompositions . . . 72
2.3.5 The dual problem in sum of squares . . . 75
2.3.6 Sum of squares liftings . . . 77
2.4 Computational algebraic geometry . . . 77
2.4.1 The complex case . . . 78
2.4.2 Hilbert's Nullstellensatz . . . 79
2.4.3 Real algebraic geometry . . . 80
2.4.4 Alternative sum of squares representations . . . 83
2.5 Putting it all together . . . 84
2.6 Sum of squares representations of rational polynomials . . . 85
2.6.1 Jibetean's approach . . . 86
2.6.2 Change of variables . . . 87
2.7 Conclusions . . . 87

3 Sum of squares for energy landscape problems 89
3.1 Introduction . . . 89
3.2 Software for sum of squares decompositions . . . 90
3.3 Model energy landscape problems . . . 91
3.3.1 Challenging topological features of the PES . . . 91
3.3.2 Multi-well potentials . . . 92
3.3.3 Golf-course-like potential . . . 95
3.3.4 Müller potential . . . 95
3.4 Symmetric polynomials . . . 96
3.5 Polynomial representation of a classical force field . . . 100

3.5.1 The Amber94 force field . . . 101
3.5.2 Ethane as a case study . . . 102
3.5.3 Mapping the terms of the potential energy functional into polynomial form. . . . 104
3.6 SOS programme for ethane model . . . 111
3.6.1 Definition of the potential energy functional . . . 112
3.6.2 Adding constraints . . . 114
3.6.3 Running the SDP . . . 115
3.6.4 Post-processing and results . . . 116
3.6.5 Improving the SOS representation . . . 118
3.7 A trigonometric basis . . . 118
3.7.1 SOS decomposition for the torsional potential in a trigonometric basis . . . 119
3.7.2 SOS decomposition for the remainder of the potential using a trigonometric basis . . . 119
3.8 Discussion . . . 122
3.8.1 Lifting . . . 123
3.8.2 Polynomial fitting . . . 124
3.8.3 Interior point methods . . . 124
3.8.4 Dealing with solvation . . . 126
3.8.5 Identifying other minima . . . 126
3.9 Conclusions . . . 127

4 Second order cone programming in coarse grained molecular dynamics 128
4.1 Introduction . . . 128
4.1.1 Single-site anisotropic potentials . . . 129
4.2 Objective . . . 133
4.3 Ellipsoid representations using SOCP . . . 134
4.4 Initial evaluation . . . 136
4.4.1 Test setup . . . 137
4.4.2 Results . . . 138
4.4.3 Conclusion . . . 139
4.5 Molecular dynamics implementation . . . 139
4.5.1 Initialisation . . . 140
4.5.2 Main loop . . . 144
4.6 Model system . . . 149
4.7 Results . . . 150
4.8 General SOS formulation for arbitrary geometries . . . 151

4.9 Discussion 154
4.10 Conclusion 156

Appendices **158**

A How to obtain the global minimum and a SOS decomposition in practice **159**

B Interior point methods **165**
B.1 Unconstrained minimisation 166
B.1.1 Methods for descending 166
B.2 Newton's method for unconstrained minimisation 167
B.2.1 The Newton step 167
B.3 Newton's method for equality constrained minimisation 169
B.3.1 The Newton step 169
B.3.2 Infeasible starts 170
B.4 Interior point methods 171
B.4.1 The central path 171
B.4.2 The path-following method 173
B.4.3 Logarithmic barriers for generalised inequalities 174
B.4.4 Primal-dual interior point methods 176
B.5 Primal-dual interior point methods applied to SDP 178

References **181**

List of Tables

1.1 The number of stationary points of index I for Lennard-Jones clusters of between 4 and 13 atoms. 23

3.1 Summary of SOS optimisation results for the 1D, 10D multi-well, Müller and Golf-like potentials. 97
3.2 Coefficients of polynomials which were fit to the X-CT-CT-X torsional functional when graphed on 4 different abscissas. 112
3.3 Coefficients of fitted polynomials to the ϕ vs. $1/r_{ij}$ curves when graphed on 4 different abscissas. 113
3.4 Mean and standard deviations of the distributions of $\log_{10}$ of the coefficients from the fitted polynomials. 114
3.5 Local minima of total potential energy function. 117

4.1 Outline of MOLSOCP coarse grained molecular dynamics programme. 142
4.2 Reduced units for coarse-grained MD simulation. 143

List of Figures

1.1 Minima, maxima and transition states on a potential energy surface. . 16
1.2 Schematic of the five interaction terms in a typical force field. 18
1.3 Contour plot for a monkey saddle. 22
1.4 A set of monotonic sequences for a LJ_{19} cluster. 25
1.5 Disconnectivity graph for a LJ_{13} cluster. 26
1.6 One-dimensional potential energy surfaces and their corresponding disconnectivity graphs. 27
1.7 Free energy graph for an alanine tripeptide. 32
1.8 The Königsberg bridge problem. 35
1.9 Polynomial versus exponential growth. 36
1.10 Complexity problem classifications. 37
1.11 Nonisocahedral global minima of Lennard-Jones clusters. 40
1.12 Disconnectivity graph for the LJ_{38} cluster. 41
1.13 Diffusion Equation Method hypersurface deformation. 44
1.14 The basin hopping transformation of a PES to a staircase energy landscape. 45
1.15 One-dimensional branch and bound example. 50

2.1 The boundary of the positive semidefinite cone in $\mathbb{S}^2$. 57
2.2 Geometrical interpretation of a convex function. 58
2.3 Geometric interpretation of the optimality condition. 61
2.4 Newton polytope for a sum of squares decomposition. 75
2.5 Schematic flowchart of a SOS optimisation. 85

3.1 Sum of squares optimisation on a 1D multi-well potential. 93
3.2 A 2D projection of a 4-dimensional 16^{th} degree multi-well potential. . 94
3.3 A 2D projection of a 10-dimensional multi-well potential. 94
3.4 Contour plot of a Golf-like landscape. 96
3.5 Contour plot of the original fitted Müller potential and its SOS representation. 97
3.6 The polynomial function $x^4 - 40x^2$. 98
3.7 The polynomial function $x^4 - 40x^2 + y^4 - 40y^2$. 99

3.8 SOS optimisation postprocessing flowchart 100
3.9 CPK representation of Ethane (C_2H_6). 102
3.10 Numbering scheme for the hydrogen molecules in ethane. 103
3.11 Potential energy function for bond stretching between two sp^3 carbons. 104
3.12 Polynomial fit of a torsional potential energy functional. 106
3.13 Plots of ϕ vs. r_{ij} for the hydrogens in ethane. 107
3.14 Plots of ϕ vs. $1/r_{ij}$ and fitted polynomials for the hydrogens in ethane. 109
3.15 van der Waals energies from the hydrogen interactions as a function of the torsional angle ϕ. 110
3.16 Coulombic energies from the hydrogen interactions as a function of the torsional angle ϕ. 111
3.17 Block diagonal positive semidefinite SOS matrix. 116
3.18 CPK representation of propane (C_2H_6). 127

4.1 The elliptic contact function for two ellipsoidal particles. 131
4.2 The directional distance of closest approach for two ellipsoidal particles. 133
4.3 The principal axes of an ellipsoid determine its geometry completely. 135
4.4 Starting configurations for minimum distance accuracy test. 137
4.5 MOSEK quadratic form evaluations at the minimum distance points (configuration 1). 139
4.6 MOSEK quadratic form evaluations at the minimum distance points (configuration 2). 140
4.7 FMINCON quadratic form evaluations at the minimum distance points (configuration 1). 141
4.8 FMINCON quadratic form evaluations at the minimum distance points (configuration 2). 141
4.9 32 ellipsoidal particles with minimum distance vectors. 145
4.10 The force and torque resulting from intermolecular interactions. . . . 146
4.11 Original and shifted Lennard-Jones potential. 147
4.12 Simulation of a binary system of Lennard-Jones anisotropic molecules: initial configuration. 150
4.13 Simulation of a binary system of Lennard-Jones anisotropic molecules: step 500 151
4.14 Simulation of a binary system of Lennard-Jones anisotropic molecules: step 20000. 152
4.15 Trajectory data for the equilibration phase of a simulation of Lennard-Jones anisotropic molecules. 153
4.16 The minimum distance vector $\mathbf{a} - \mathbf{b}$ between two arbitrarily shaped particles. 154

Publications

M. G. Burke and S. N. Yaliraki, *Exploring Model Energy and Geometry Surfaces Using Sum of Squares Decompositions*, Journal of Chemical Theory and Computation, 2, 575-587 (2006).

Chapter 1

Introduction

1.1 Energy landscapes background

Energy landscape theory is a broad subject that encompasses much of modern theoretical chemistry. It seeks to describe and sometimes catalogue the energy (either free or potential energy) of a system of atoms or molecules in its (many) relevant physical states. From this study a better understanding of the thermodynamic and kinetic properties of the system can emerge. The widespread use of energy landscape theory is a relatively recent occurrence but it is now considered to be fundamental to the study of many examples of complex molecular systems. The theory of course is not new but what has led to a rise in interest in the area is the arrival of affordable computing power which has put the study of systems of interest using the energy landscape framework within reach. Examples of problems which have been studied using this approach range from understanding glassy behaviour to protein folding to the structure and dynamics of atomic clusters.

The focus of this chapter is on energy landscape theory itself. It aims to outline the current state of the art in energy landscape theory, detailing where it is valid, and in particular focusing on potential energy surfaces (PES) since these are the subject of interest for much of this thesis. It describes how potential energy landscapes are categorised, what are the salient features and how these can be used to provide information to the chemist. Some time is then spent describing state of the art methods for exploring potential energy surfaces and particular attention is paid to the problem of identifying the global minimum (the lowest point on the surface) since this is the problem which is addressed in chapter 3.

A very comprehensive review of energy landscape theory and methodologies is the recent book "Energy Landscapes" by Wales [3], which provides a very extensive bibliography as well as excellent coverage of the topic. With the exception of computational complexity theory most of what is discussed in this chapter is covered in greater detail for the interested reader.

1.1.1 Free and potential energy landscapes

An important distinction to be made when discussing energy landscape theory is that between free and potential energy landscapes. Although the term energy landscape was first used to describe free energy landscapes [4–10] (low-dimensional functions of one or a few order parameters), the landscape that will be of greatest interest here is the potential energy landscape or potential energy surface (PES) as it will be referred to throughout this work.

Potential energy landscapes

The potential energy surface is a high dimensional function of all the independent degrees of freedom of an atomic or molecular system. The topography and topology of a PES plays a defining role in determining both the structure and dynamics which any system will assume. The topography of a landscape is defined primarily by stationary points on the surface which consist of minima, transition states and for completeness maxima. Minima are points for which every displacement requires an increase in energy and therefore correspond to mechanically stable configurations where the gradient is zero (figure 1.1). The PES of many systems of interest may contain astronomically many minima, and equilibrium thermodynamics is concerned only with the relative energies of the minima and the volume of configurational space for which each minimum is an attractor. The lowest minimum on the PES is called the global minimum and provided it is kinetically accessible this is the region most highly sampled in thermodynamic equilibrium. Minima are often associated with "reactants" and "products" in a chemical reaction.

Transition points and their relative heights above the adjacent minima decide the kinetic behaviour of the system (figure 1.1). These can be compared to saddles between mountains which connect valleys on either side of a mountain range together. Pathways linking minima travel over these saddles since they represent the lowest highest points over which a pathway connecting minimum A from minimum B must cross. The connectivities of the various minima through the different transition pathways defines the topology of the potential energy landscape which as stated earlier defines the dynamics of the system. Thermodynamic properties can be regarded as being properties of an ensemble average of systems sampling the same potential energy surface.

Normally the potential energy U for a system of N atoms is defined as a function of the $3N$ configurational degrees of freedom of the system, therefore the PES is a $3N$-dimensional surface embedded in a $(3N+1)$-dimensional space. This representation of the PES is made under the assumption of the Born-Oppenheimer approximation (see section 1.2).

Numerous heuristic approaches to the study of potential energy surfaces have been proposed in recent years yet a number of outstanding challenges remain. In this work a new methodology for studying atomistic potential energy landscapes is introduced accompanied by a variation of the method which can be coupled to molecular dynamics to explore coarse-grained potential energy surfaces. The methodology introduced has some advantages over current approaches for specific problems in energy landscape theory such as identifying the global minimum on the surface. The method, its advantages, limitations, as well as some applications will be described in detail in the coming chapters.

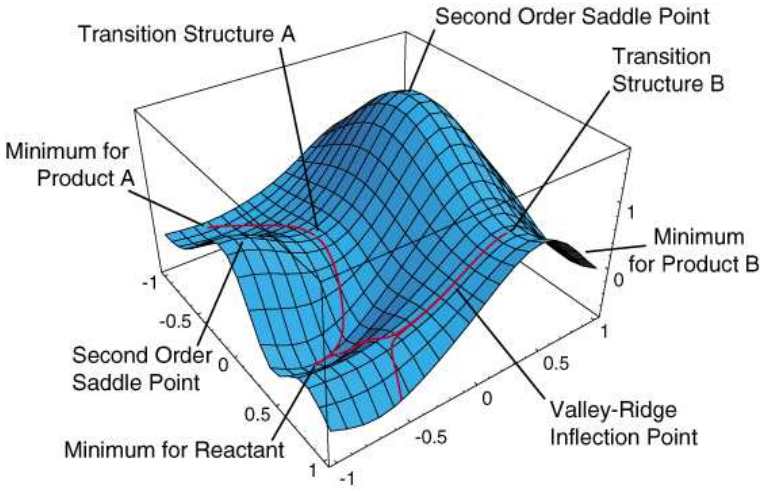


Figure 1.1: Minima, maxima and transition states on a potential energy surface. Adapted from [11].

Free energy landscapes

The free energy surface is a low dimensional function of a number of order parameters which is obtained from the PES by averaging over all the other degrees of freedom of the system for fixed values of the order parameters. Therefore it can be argued that the potential energy surface is a more fundamental object than the free energy surface since it requires no choice of appropriate order parameters to be made [12]. The study of free energy surfaces is an interesting topic in itself since free energy determines thermodynamic equilibrium in real systems. However free energy surfaces are not dealt with in this work and therefore will not be dealt with in detail in this introduction.

1.1.2 Calculating potential energy surfaces

A fundamental question is where do the functional forms and associated parameters that define the PES come from? In fact there are a number of ways that a PES can be

derived. The most rigorous approach is that of quantum mechanical calculations in which approximate solutions to the Schrödinger equation are obtained. These range from Hückel and extended Hückel theory [13], semiempirical approaches, Hartree-Fock self consistent field theory, Møller-Plesset perturbational approaches to electron correlation [14], density functional theory [15] and finally quantum Monte Carlo methods which include electron correlation effects [16]. In all cases in order to produce a surface it is necessary to repeat such calculations over a discrete grid in configurational space and use interpolating functions to fit these points.

Because of the computational expense of such calculations for anything but modest systems a compromise is to employ interatomic or even intermolecular potentials which are based on a combination of fitting to experimental data and quantum mechanical calculations. Such a PES is often termed a force field for historical reasons. In general a force field is a series of terms each representing interactions between increasingly large sets of atoms

$$U = \epsilon^{(0)} + \sum_a \epsilon_a^{(1)} + \sum_a \sum_{b>a} \epsilon_{ab}^{(2)} + \sum_a \sum_{b>a} \sum_{c>b} \epsilon_{abc}^{(3)} + \cdots . \tag{1.1}$$

Here the superscript (n) represents the n^{th}-body interaction contribution. A valence force field is written as a function of bond lengths, bond angles, and torsion (or dihedral) angles. Such force fields contain terms up to 4^{th} body terms. Here the bond length and bond angle terms are often represented as quadratic functions $k_s(r - r_e)^2$ and $k_b(\theta - \theta_e)^2$ where r and θ are the interatomic distances and bond angles respectively and r_e and θ_e are their respective equilibrium values. k_s and k_b are the stretch and bend force constants respectively. Dihedral angles are defined by the angle between two planes. The planes in turn are defined by the positions of the first, second, third and second, third, fourth atoms respectively in a linear sequence of four atoms. The functional form of the improper torsions is usually either a harmonic expression $k_{it}(\phi - \phi_e)^2$ or an expansion $\sum_m k_{it,m}[1 + \cos(m\phi + \gamma)]$ where ϕ is the torsional angle and γ is the phase factor.

Beyond a valence force field, non-bonded interactions usually include electrostatic contributions and dispersion effects. Electrostatic terms usually include contributions from point charge interactions, dipole moments as well as contributions from higher order multipoles. Figure 1.2 provides a schematic description of the five interaction terms.

The first widely used classical force field designed to cover a variety of molecules was the molecular mechanics (MM) force field of Allinger [17, 18] of which MM3 is the current release. There are many others now and some of the more widely used include AMBER [19], CHARMM [20, 21], ECPP [22, 23], GROMOS [24], OPLS [25] and UNRES [26, 27].

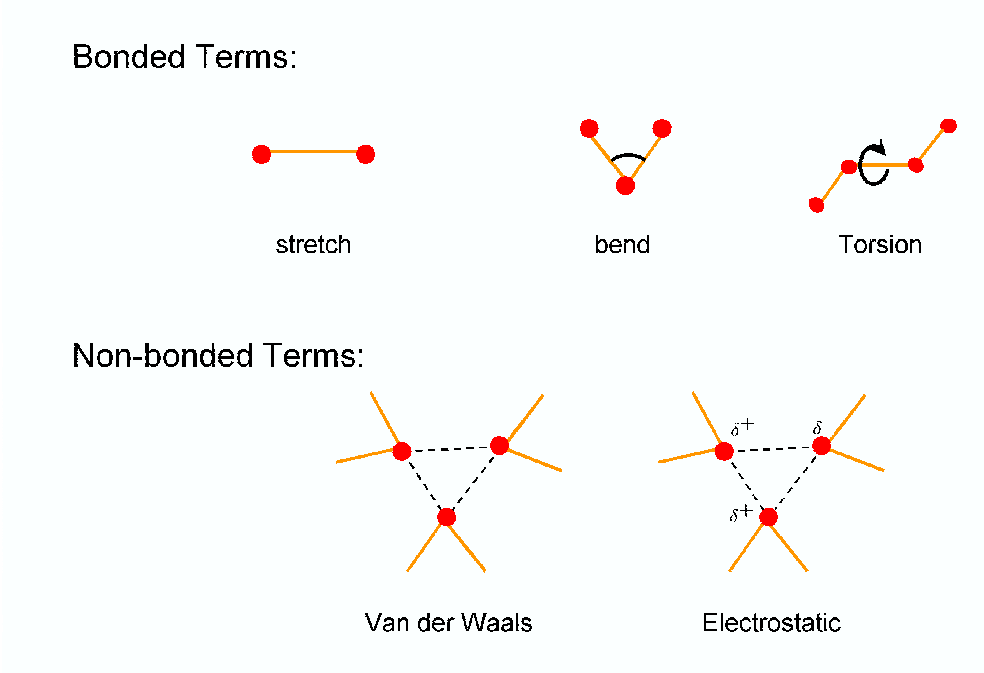


Figure 1.2: A schematic of the five interaction terms in a typical force field. The terms are often subdivided into bonded and non-bonded terms. Atoms are represented as red dots and covalent bonds are represented as orange lines. Non-covalent interactions are represented as dashed lines.

1.2 The Born-Oppenheimer Approximation

The Born-Oppenheimer approximation is central to the concept of a potential energy surface and without it, classical PES's would not exist. The approximation made is to neglect the terms that couple together the nuclear and electronic degrees of freedom. The result is that the Schrödinger equation can be separated into independent nuclear and electronic parts. This makes the construction of partition functions in statistical mechanics much simpler. The physical basis for this approximation is that the nuclei and electrons are so different in mass (the mass of a proton is 1836 times the mass of an electron) that the electron density can adjust almost instantaneously to a change in the position of the nucleus. The result of the Born-Oppenheimer approximation from a energy landscape point of view is to make it possible to define a single PES for each electronic configuration of the system. By far the most common PES studied and the one that will be of interest in this work is the ground state PES.

The Schrödinger equation for a molecule with n electrons, with mass M_e and N nuclei of mass m_t, is given by [12]

$$\left(-\sum_{t=1}^{N} \frac{\hbar^2}{2m_t}\nabla_t^2 - \sum_{i=1}^{n} \frac{\hbar^2}{2M_e}\nabla_i^2 + U(\mathbf{X},\mathbf{x}) \right) \psi(\mathbf{X},\mathbf{x}) = E_{total}\psi(\mathbf{X},\mathbf{x}), \qquad (1.2)$$

where $\mathbf{X}$ and $\mathbf{x}$ are the electronic and nuclear coordinates respectively and $\hbar = h/2\pi$ where h is Planck's constant. The potential energy is defined by

$$U(\mathbf{X},\mathbf{x}) = \frac{e^2}{4\pi\epsilon_0}\left(-\sum_{i,t}\frac{Z_t}{r_{it}} + \sum_{i<j}\frac{1}{r_{ij}} + \sum_{t<s}\frac{Z_t Z_s}{r_{ts}}\right), \tag{1.3}$$

where Z_t is the atomic number of nucleus t and e is the charge on a proton. $r_{it} = |\mathbf{X}_i - \mathbf{x}_t|$ where $\mathbf{X}_i$ and $\mathbf{x}_t$ are the position vectors of electron i and nucleus t and ϵ_0 is the dielectric constant in a vacuum. The Born Oppenheimer approximation gives the total wave function as

$$\psi(\mathbf{X},\mathbf{x}) = \psi_e(\mathbf{X},\mathbf{x})\psi_n(\mathbf{x}), \tag{1.4}$$

where $\psi_e(\mathbf{X},\mathbf{x})$ is the solution to the "electronic Hamiltonian"

$$[\hat{\mathcal{H}} - \hat{K}_n]\psi_e(\mathbf{X},\mathbf{x}) = U_e(\mathbf{x})\psi_e(\mathbf{X},\mathbf{x}), \tag{1.5}$$

where $\hat{\mathcal{H}}$ is the total Hamiltonian operator in (1.2) and $\hat{K}_n$ is the nuclear kinetic operator in (1.2). $\psi_e(\mathbf{X},\mathbf{x})$ is a function of the electronic coordinates $\mathbf{X}$ but also depends parametrically on the nuclear coordinates $\mathbf{x}$. The electronic energy $U_e(\mathbf{x})$ also depends on the nuclear coordinates and it is this variation that defines the PES. Later the subscript e will be dropped and the potential energy surface will be referred to as $U(\mathbf{x})$.

To appreciate how valid the Born-Oppenheimer approximation actually is consider a hydrogen-deuterium diatom. Under the Born-Oppenheimer approximation, the diatom H−D has no dipole moment, where in reality its dipole moment is 10^{-4} Debye, making it a good approximation. The quality of the approximation actually improves for heavier atoms. There are a number of possibilities where the approximation can break down but these involve transitions from the ground state potential energy surfaces and therefore will not be discussed here.

1.3 Categorising and describing energy landscapes

1.3.1 Stationary points and their significance

Stationary points are defined as points where the gradient of the potential vanishes and are the most noteworthy points on the PES. The Taylor expansion of the potential at a stationary point in normal coordinates $Q_1, \ldots, Q_{3N}$ is

$$U(\mathbf{Q}) = \frac{1}{2}\sum_{\alpha=1}^{3N}\omega_\alpha^2 Q_\alpha^2 + \mathcal{O}(Q^3) + \ldots, \tag{1.6}$$

where $\mathbf{Q} = \mathbf{0}$ denotes a stationary point and $U(\mathbf{0}) = \mathbf{0}$. $\mathcal{O}(Q^3)$ denotes terms of degree 3. The curvature of the PES at $\mathbf{Q} = \mathbf{0}$ in the direction α is given by the Hessian eigenvalues ω_α^2. Different types of stationary points can be distinguished by the entries in the Hessian at that point. The number of negative eigenvalues in the Hessian at a stationary point is called the Hessian index (I) or the order of that point. Therefore minima are stationary points of zero order.

A transition state on a PES is defined as a stationary point with a single negative eigenvalue [28] corresponding to a negative force constant, or curvature of the surface and to an imaginary normal mode frequency (ω_α). Stationary points with one or more negative eigenvalues are called saddle points. The term "transition state" is used throughout the literature, often with different meanings, so it is useful to make clear what these definitions are and how they are related. A reaction can be represented as a free energy profile plotted as a function of some reaction coordinate. A unique maximum on this free energy profile defines a transition state for the reaction, in which all other degrees of freedom except the reaction coordinate have been averaged. This definition of a transition state comes from unimolecular rate theory [29–32]. Such a free energy transition state will typically include configurations associated with various local minima on the PES. In certain cases it may be possible to calculate the free energy profile directly from a sampling of the local minima on the PES.

An alternative definition of a transition state is arrived at by considering the dynamical behaviour of a system on a potential energy surface. A region of the PES where the system is equally likely to relax into either a reactants or products configuration is called the seperatrix and can also be used to define a transition state. In the simplest cases this transition state may coincide with the earlier definition of a transition state on a PES. In transition state theory (for prediction of unimolecular rate constants) [29, 30] an important assumption made is that any trajectory originating in the reactants region and which crosses the seperatrix will enter the products region without recrossing the surface. Here the dividing surface passes through the potential energy maximum which corresponds to the earlier definition of a transition state. However in variational transition state theory, the position of the dividing surface is adjusted until the flux through it is minimised and this minimised flux is used to calculate the rate constant [33, 34]. It has been demonstrated that for rates calculated at conditions of constant temperature, the seperatrix of minimum flux is located at the position of maximum Gibbs free energy [34–36].

1.3.2 Steepest descent pathways

Steepest descent pathways have been used to divide the PES into basins of attraction, i.e. regions of configurational space which map to a local minimum through a

steepest descent minimisation. This aids in a discretisation of the PES which will be described in section 1.4. The definition of a steepest descent pathway is one that follows the negative of the gradient vector. The path can be expressed in terms of a parameter s by the first order differential equation [12]

$$\frac{d\mathbf{x}}{ds} = \frac{-g(\mathbf{x})}{\|g(\mathbf{x})\|_2}, \tag{1.7}$$

where $g(\mathbf{x})$ is the gradient vector and s is the integrated path length. A steepest descent path is uniquely specified by a single point with $g(\mathbf{x}) \neq 0$. At a true transition state there are only two solutions to equation 1.7 which decrease the potential energy for a small finite displacement from the stationary point. However zero Hessian eigenvalues excluding those relating to overall translation and rotation can result in unusual saddles such as the monkey saddle for cubic surfaces (figure 1.3).

For any stationary point of a non-linear molecule there are $2(3N-6)$ principal steepest descent paths corresponding to positive and negative displacements along the $3N-6$ eigenvectors with nonzero eigenvalues. There are infinitely many steepest descent paths which are combinations of the principal steepest descent paths, however all steepest descent paths approach a minimum parallel to the principal direction with the smallest eigenvalue for which the path has a non-zero projection. For all but the most modest systems, using steepest descent methods to catalogue stationary points is computationally inefficient.

1.4 Quantification and topological considerations for stationary points

It is instructive to consider how the number of stationary points scales with the number of atoms in a system. For all but the smallest systems representing or categorising these stationary points becomes a challenging task in itself. However techniques do exist which facilitate this and these are described in section 1.4.2. Consider a system consisting of N_A atoms of type A, N_B atoms of type B etc., for which the Hamiltonian is invariant to permutations of equivalent nuclei as well as to inversion of all coordinates through the origin. The number of permutation-inversion isomers of a specific structure could be as large as

$$2N_A!N_B!\ldots, \tag{1.8}$$

but if the system possesses symmetries this will be reduced. For a molecule whose point group contains o members there will be o permutation-inversion operations

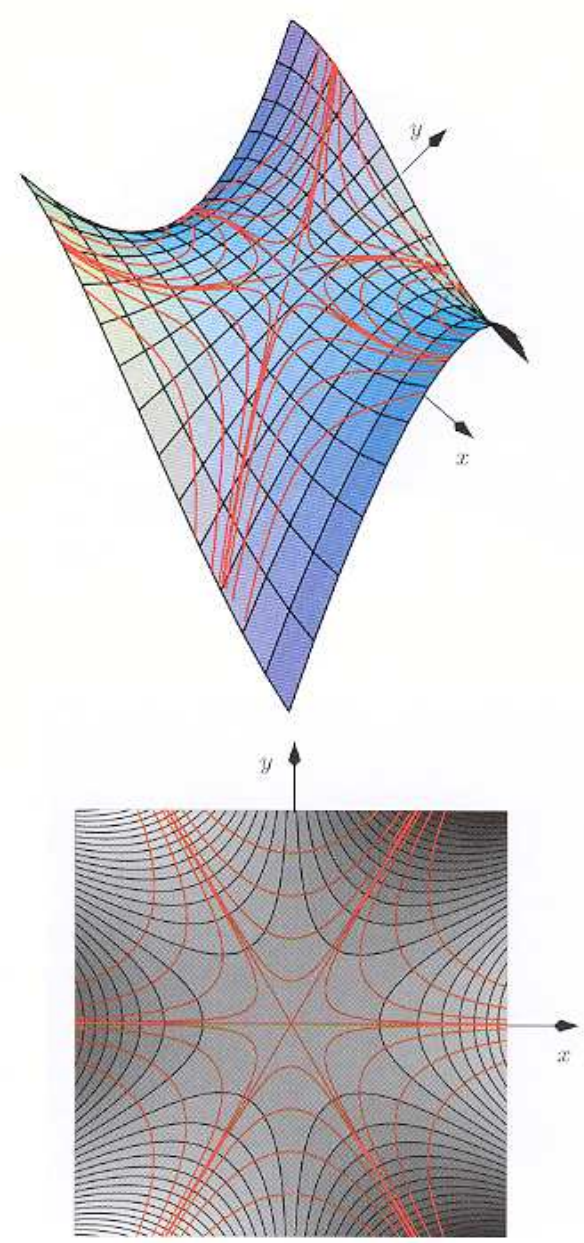


Figure 1.3: Contour plot for a monkey saddle (generated using the cubic polynomial function $x^3 - 3xy^2$) showing steepest descent paths (red curves) and zero Hessian eigenvalue directions, for example, the ±y directions [12]. The shading is darkest for low energy.

(operations that either permute the positions of identical atoms or invert their positions through the centre of mass). The number of permutation-inversion isomers is [12]

$$n = \frac{2N_A!N_B!\cdots}{o}. \tag{1.9}$$

For example in a Lennard-Jones cluster of 7 particles (LJ_7), i.e. a cluster of 7 particles interacting through a Lennard-Jones potential there are 4 local minima, however the number of permutational isomers of these minima ranges from 504 for

the most symmetrical geometry to 5040 for the least symmetrical [37]. Note that enantiomers are counted within the sets of permutation-inversion isomers.

There is however no general rule defining how the number of stationary points on the PES scales with the number of atoms in the system. Both empirical data [38] and simple theories [39–41] indicate an exponential increase (as a function of the number of atoms, N) in the number of minima (n_{min}) on the PES. A widely used functional form is

$$n_{min}(N) = e^{\alpha N}, \tag{1.10}$$

where α is a system dependent constant. Similar results have been obtained for the number of transition states (n_{trans}) [42] where the functional form is

$$n_{trans}(N) = Ne^{\alpha N}. \tag{1.11}$$

Therefore this suggests that the ratio n_{trans}/n_{min} grows linearly with increasing N. Evidence for this has been observed for small Lennard-Jones clusters. Table 1.1 shows the number of different stationary point structures of a given index I for Lennard-Jones clusters of $4-13$ atoms [42].

	Number of Atoms									
I	4	5	6	7	8	9	10	11	12	13
0	1	1	2	4	8	21	64	170	513	1505
1	1	2	3	12	42	165	635	2423	7868	25653
2	2	4	13	44	179	867	4074	17109	*27957*	*88079*
3	1	6	24	98	494	2820	16407	*47068*		
4	1	6	30	168	1000	6729	46277			
5	0	2	26	190	1458	12093	*97183*			
6	0	1	16	168	1619	16292				
7	0	0	5	101	1334	16578				
8	0	0	1	45	852	13226				
9	0	0	0	11	388	8286				
10	0	0	0	1	125	4053				
11	0	0	0	0	26	1444				
12	0	0	0	0	1	376				
13	0	0	0	0	0	56				
14	0	0	0	0	0	1				
15	0	0	0	0	0	0				

Table 1.1: The number of stationary points of index I for Lennard-Jones clusters of between 4 and 13 atoms [42]. The numbers shown in italics are probably weak lower bounds.

Combining the equation for the number of permutational isomers with the estimate for the number of local minimum structures (1.10) gives an estimate for the

total number of minima on a homonuclear PES as

$$n_{min}(N) = 2N!e^{\alpha N}. \tag{1.12}$$

1.4.1 Topological categorisation

The first work on formally describing the topological relationships between stationary points on the PES was the work by Kunz and Berry on monotonic sequences. These are sequences of local minima and their connecting transition states in which the potential energy at subsequent minima monotonically decreases [43, 44]. Monotonic sequences can be grouped into sets which converge to a particular minimum and such sets are known as basins. An example of a basin is shown in figure 1.4 where a cluster of 19 argon atoms was modelled as a LJ_{19} cluster. The barrier heights are typically large compared to the successive change in potential energy, however this is not always the case. The argon cluster monotonic sequences are described as being sawtooth-like while monotonic sequences for a $(KCl)_{32}$ cluster are described as being staircase-like in profile [12]. This is useful in helping to identify good clusters which rapidly seek some local minimum, however monotonic sequences are not very useful for the study of global properties of the PES.

Another related concept is that of a basin of attraction [44] or catchment region [45] which are terms used to describe the region of configurational space where the steepest descent pathways converge to a particular minimum. An alternative description of basins of attraction was given by Stillinger and Weber when they coined the term inherent structures [39] for local minima obtained by quenching the kinetic energy while sampling using molecular dynamics trajectories.

1.4.2 Disconnectivity graphs

A more global representation of the topology of the PES was devised by Becker and Karplus who used a graph theoretical approach to represent the multidimensional connectivity of the minima and transition points as a graph, which they called a disconnectivity graph [47]. The first application was to the tetrapeptide isobutyryl-$(ala)_3$-NH-methyl(IAN) based on a database of minima which had already been constructed by Czerminski and Elber [48]. In the disconnectivity graph approach the ordinate defines the energy U. At any given energy U the minima can be put into disjoint sets known as superbasins such that all members of a particular superbasin set are accessible from each other through any transition state without having to exceed the energy U. Therefore if the highest transition state on the lowest energy pathway connecting two minima exceeds U then the two minima are in different sets. When U lies below the second lowest minimum, the only superbasin remaining defines the global minimum. As U is increased more superbasins appear

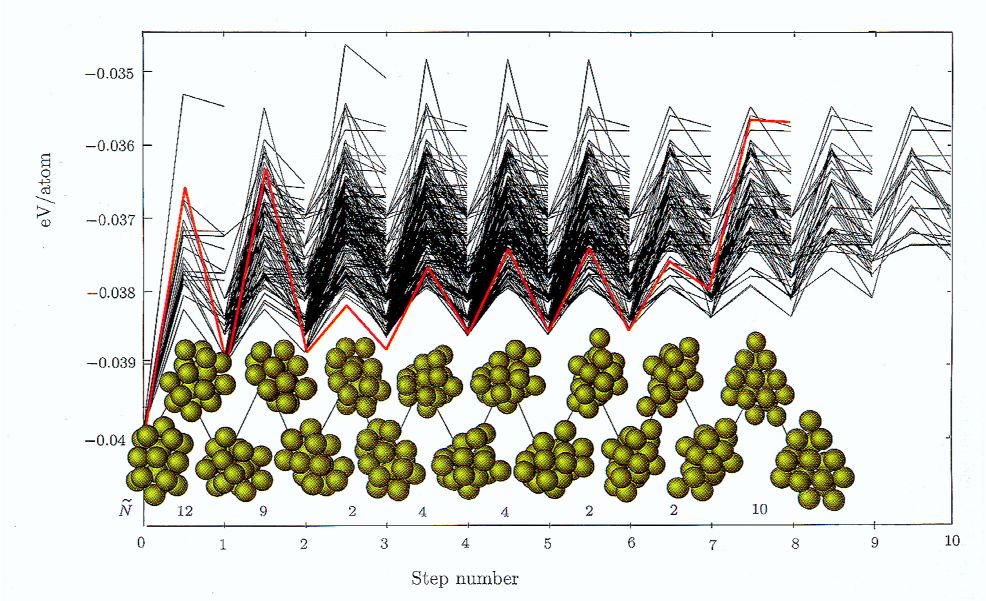


Figure 1.4: A set of monotonic sequences for an LJ_{19} cluster. The structures shown correspond to the highlighted monotonic sequence. Notice that in this case the barriers are typically large relative to the change in energy between successive minima. $\tilde{N}$ is the number of atoms which participate in each rearrangement [46].

and disappear again as the energy barrier separating them is exceeded. In the limit $U \to \infty$ there again is only one superbasin containing all local minima. A disconnectivity graph for a cluster of 13 atoms interacting via a Lennard-Jones potential is given in figure 1.5.

Constructing a disconnectivity graph first involves building a database of minima and transition points. Then a superbasin analysis is performed at discrete energy levels, and each superbasin is represented by a point on a horizontal axis at that particular energy. Varying the frequency of energies at which the superbasin analysis is performed will affect the coarse-graining of the disconnectivity graph. If a large gap ΔU is used the graph will loose much of the topological information leading to a more coarse-grained representation of the landscape.

There are three common disconnectivity graph motifs which are recognisable in many of the disconnectivity graphs reported in the literature. All three are called after trees which isn't surprising given their structure. Schematic examples are given in figure 1.6 along with the 1 dimensional potential energy surface that each graph represents. The palm tree (shown in figure 1.6(a)) is typical of a system which is good at finding its global minimum since the barriers to be overcome are relatively small when compared to the ΔU between the highest and lowest minimum. This is also the type of PES that is likely to produce a free energy surface with a single well defined free energy minimum. Figure 1.6(b) shows a willow tree which involves higher barriers which can easily form kinetic traps depending on the initial

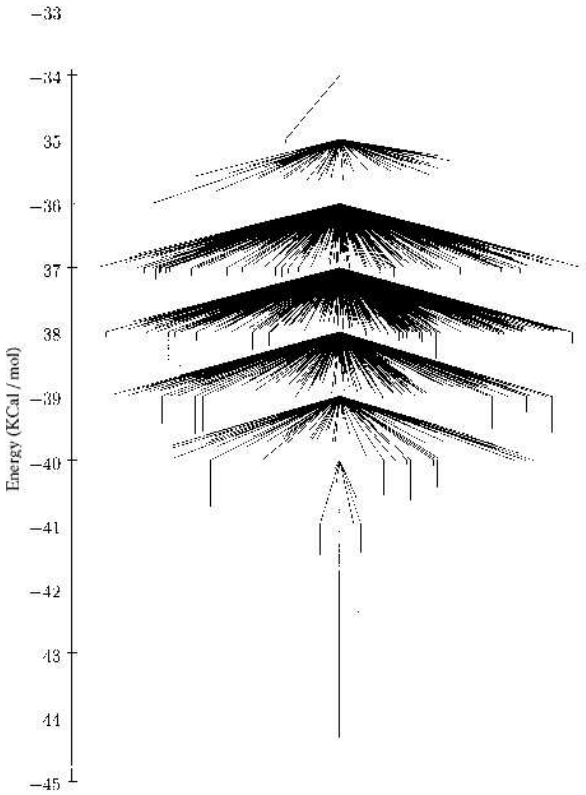


Figure 1.5: Disconnectivity graph for a cluster of 13 identical atoms interacting via a Lennard-Jones potential. All 1467 minima are shown [49].

configuration. However there is still a single clear global minimum. At the other extreme to the palm tree is the banyan tree (figure 1.6(c)) which has many low lying minima and the barrier heights far exceed that of the differences between the minima. Most real landscapes will be a combination of the above three with a common example being that of a PES with multiple palm tree topologies in different regions of configurational space. Typically such landscapes can have different free energies associated with these different regions of configurational space.

Becker and Karplus also presented disconnectivity graphs where the ordinate is temperature (in kT) [47]. The topological analysis has also been carried out in terms of free energies, using approximations to the densities of states for calculating the partition function of each minimum [50–52].

1.5 Approaches for sampling the potential energy landscape

The importance of the role of stationary points in energy landscape theory has been made clear by now but little has been said about how to identify and catalogue these points. Surveying the PES is commonly described as geometry optimisation and usually involves the following:

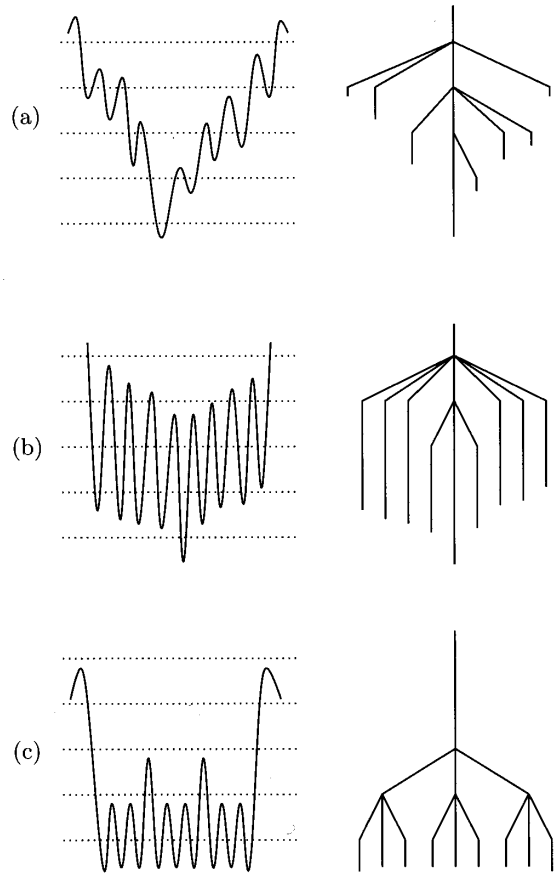


Figure 1.6: One-dimensional potential energy surfaces and their corresponding disconnectivity graphs. (a) Palm tree: Low barriers with minima dropping monotonically downhill leading to a well defined global minimum. (b) Willow tree: Larger downhill barriers but still a single obvious global minimum. (c) Banyan tree: Multiple low lying minima separated by large barriers. On the PES, the energies at which the superbasin analysis was performed is indicated by the dotted lines [12].

1. Identifying local minima on the PES.
2. Identifying transition states connecting the minima as well as higher index saddles.
3. Finding the global minimum.

1.5.1 Identifying local minima

Many different algorithms are available for finding the local minima, and include methods based on first and second derivative information such as Newton's method. Others use just first derivative information since evaluating and inverting the Hessian can be very expensive both in terms of time and memory requirements. Other methods simply use function evaluations without any derivative information. Methods that use an approximation of the inverse Hessian are known as quasi-Newton

methods [53,54] and a popular one for large systems is the limited memory Broyden-Fletcher-Goldfarb-Shanno method of Nocedal (LBFGS) [55–60]. The LBFGS method approximates the inverse Hessian using gradient information from a prespecified number of previous steps. An important advantage is that storage requirements scale as $\mathcal{O}(N^2)$ instead of $\mathcal{O}(N^3)$ as is the case with conventional Newton's method. Conjugate gradient methods have also been used [61, 62].

1.5.2 Identifying transition states

Finding transition states is more difficult since it is a maximisation problem in one direction while being a minimisation in all others. The search for transition states usually begins with a database containing a sample of minima on the PES which is then extended by identifying connecting transition states and pathways resulting in an extension of the database of minima as pathways leading to further minima are identified. A number of algorithms for identifying transition states have been proposed and these are commonly categorised as being either single ended or double ended [12].

Methods commencing from one minimum

Single ended algorithms move from a single starting minimum in a given direction toward a transition state. Most single ended methods are variants on the eigenvector following approach. As the name suggests eigenvector following uses an eigenvector decomposition of the Hessian matrix at the starting minimum to determine the direction with the highest probability of leading to a transition state. Since individual molecules and systems with periodic boundary conditions contain three zero eigenvalues, inverting the Hessian for the Newton step is made impossible. A shifting procedure can be used to make the zero eigenvector non-zero and moves in these directions can be ignored [63]. Typically the geometry optimisation is carried out in a coordinate system which uses the Hessian eigenvectors as a basis. In such a basis transformation the point $\mathbf{x}$ is transformed to the point $\mathbf{w}$. A Newton step (see appendix B.2) in the direction α in such a coordinate system is

$$\Delta x_{Newton,\alpha} = -\frac{g_\alpha(\mathbf{w})}{\epsilon_\alpha^2}, \tag{1.13}$$

where g_α is the gradient and ϵ_α^2 is the eigenvalue in the direction α. Moving in directions with $\epsilon_\alpha^2 < 0$ leads to a reduction in energy while moving in directions with $\epsilon_\alpha^2 > 0$ leads to an increase in energy. For large systems (containing hundreds or thousands of atoms) calculating the Hessian and indeed diagonalising it may become prohibitively expensive. If formation of the Hessian is possible then iterative

techniques can be used to identify the eigenvectors with the smallest and largest nonzero eigenvalues. A hybrid eigenvector-following with minimisation scheme can be applied in which a maximisation is carried out along the direction with the smallest non-zero eigenvalue whilst minimisation is carried out in the other directions [64]. Other approaches based on similar ideas include reduced gradient following [65–70], distinguished coordinate methods [48,71–79] and the conjugate peak refinement method [80]. If it is not feasible or efficient to calculate the Hessian a variational approach can be used to identify the smallest eigenvector and its corresponding eigendirection.

Methods commencing from two minima

Double ended methods normally start with two known minima and an attempt is made to find a connecting transition state [48, 80–88]. Early work on double ended methods was pioneered by Elber and Karplus [82], Pratt [81] and Chandler [89]. Although there cannot be a guarantee that the two minima in question are connected by a single transition state, it is always possible to find a connecting pathway if intermediate minima and transition states are considered. The nudged elastic band method of Jónson [90–93] is one of the most successful and has recently been incorporated into the Amber molecular modelling package. The principle of the method is to consider a series of configurations (called images) at regularly spaced intervals along the pathway between the two minima. The regularity of the spacing is maintained by applying attractive springs between adjacent pairs of images. Because the spring forces should only control the spacing of the images only the component of the force tangent to the path at image i is considered

$$F_i^s = k(\|\mathbf{x}(i+1) - \mathbf{x}(i)\| - \|\mathbf{x}(i) - \mathbf{x}(i-1)\|)\hat{t}_i, \tag{1.14}$$

where k is a force constant and $\hat{t}_i$ is the normalised tangent vector to the path at point $\mathbf{x}(i)$.

For calculating the gradient forces only the forces orthogonal to the path are considered, so the total force on image i is defined as

$$F_i = F_i^s - \nabla U(\mathbf{x}(i)) + (\nabla U(\mathbf{x}(i)) \cdot \hat{t}_i)\hat{t}_i, \tag{1.15}$$

where the third term removes components of the real force parallel to the path. Recently further improvements to the method have been made which allows images to climb uphill in energy [91] and now the end points can be dynamically adjusted allowing the user to focus on a particular region of interest.

1.5.3 Sampling thermodynamic properties on the PES

Calculating thermodynamic properties requires a knowledge of a Boltzmann weighted sum of a great many minima on the PES which forms a realistic representation of the partition function. Approaches for calculating such a sum usually involve sampling. Methods for sampling the landscape in order to calculate thermodynamic properties are usually based on molecular dynamics or Monte Carlo simulation approaches and sophisticated techniques have been developed that allow sampling to be carried out efficiently while still producing a thermodynamically representative sample of minima. These methods include histogram methods and biased sampling methods. Thermodynamic properties of interest include the density of states, $\Omega(E)$ (the volume of phase space in which the system has energy E), for micro-canonical ensembles and the partition function, $Z(T)$, for canonical ensembles which give access to a variety of other thermodynamic properties.

Histogram methods [94–97] discretise the potential energy probability distribution for a particular temperature T with $n(T_i, U_j)$ being the number of data points from a simulation with potential energy in the range $[U_j, U_j + \Delta U]$. Since $n(T_i, U_j)$ is therefore proportional to the probability that the system has potential energy in this range at temperature T, it can be shown that

$$\frac{n(T_i, U_j)}{\sum_j n(T_i, U_j)} = \frac{\Omega_c(U_j)\Delta U e^{-U_j/kT}}{\sum_{U_j} \Omega_c(U_j)\Delta U e^{-U_j/kT}}, \tag{1.16}$$

where the sum j is over the different bins in the histogram. Therefore the density of states can be obtained up to a constant multiplier. The multiple histogram approach in which histograms are constructed over a range of temperatures allows the density of states and partition function to be obtained using a least squares fit, provided that overlapping ranges of potential energy are included.

By combining multiple histogram methods with biased sampling techniques, averages for observables can be efficiently sampled using the following biasing scheme for an observable O,

$$\begin{aligned}
\langle O(T) \rangle &= \frac{\int O e^{-U/kT} d\mathbf{x}}{\int e^{-U/kT} d\mathbf{x}} && (1.17)\\
&= \frac{\int O e^{W/kT} e^{-(U+W)/kT} d\mathbf{x}}{\int e^{-(U+W)/kT} d\mathbf{x}} \frac{\int e^{-(U+W)/kT} d\mathbf{x}}{\int e^{W/kT} e^{-(U+W)/kT} d\mathbf{x}} && (1.18)\\
&= \frac{\langle O e^{W/kT} \rangle_W}{\langle e^{W/kT} \rangle_W}, && (1.19)
\end{aligned}$$

where the W subscript denotes that the average has been performed for the biased potential $U + W$.

The free energy difference between a pair of states may be calculated using a form

of biased sampling called umbrella sampling [98] and biased sampling and histogram methods combined can be used to calculate free energy as a function of one or more order parameters [99].

Multiple simulation techniques have also been used to improve the sampling of minima on the PES by overcoming high potential barriers. The main idea is to run a number of simulations for identical systems at different temperatures and allowing the systems to interchange configurations in a way that maintains the appropriate ensemble. The jump-walking algorithm was the first to use such an approach to improve sampling [100, 101]. This approach used information from a previous saved run at a higher temperature to cross barriers between multiple low energy regions of configurational space. At regular intervals a trial move to a configuration from the saved high temperature simulation is attempted with the probability

$$P(\mathbf{x} \rightarrow \mathbf{x}') = \min\left(e^{(\frac{1}{kT'}-\frac{1}{kT})(U(\mathbf{x}')-U(\mathbf{x}))}, 1\right), \tag{1.20}$$

where the prime denotes the high temperature run.

A parallel version of the jump-walking algorithm is the parallel tempering algorithm, also known as the multiple Markov chain or replica exchange algorithm [100–112]. Here the moves are attempted between instantaneous configurations on parallel trajectories at different temperatures. A significant advantage of this approach over jump-walking is the removal of the need to store and access multiple trajectories.

1.5.4 Calculating thermodynamic properties from a PES minima database

Disconnectivity graphs have been produced with free energy as the ordinate. The free energy is approximated for each minimum on the PES. Using a harmonic approximation to the vibrational density of states the canonical partition function can be written as

$$Z_\alpha(T) = \frac{n_\alpha e^{-\beta U_\alpha}}{(\beta h \bar{\nu}_\alpha)^\kappa}, \tag{1.21}$$

where n is defined in equation 1.9, U_α is the potential energy for minimum α, $\bar{\nu}_\alpha$ is the geometric mean vibrational frequency and κ is the number of vibrational degrees of freedom. The RRKM (Rice-Ramsperger-Kassel-Marcus) rate constant out of minimum α via transition state $\dagger$ with potential energy $U^\dagger$ and geometric mean frequency $\bar{\nu}^\dagger$ is

$$k_\alpha^\dagger(T) = \sigma \frac{\bar{\nu}_\alpha^\kappa}{\bar{\nu}^{\dagger(\kappa-1)}} e^{-(U^\dagger - U_\alpha)/kT}. \tag{1.22}$$

The free energy for the minima and transition states are

$$A_\alpha(T) = -kT \ln Z_\alpha,$$
$$A^\dagger(T) = -kT \ln Z^\dagger, \tag{1.23}$$

leading to a relationship between the free energy difference and the rate constant

$$A_\alpha(T) - A^\dagger(T) = kT \ln(hk_\alpha^\dagger(T)/kT). \tag{1.24}$$

Using this definition of $A^\dagger$ the free energy of the transition state may be lower than that of a connected minimum so an arbitrary constant f_c is added to all transition state free energies. Superbasin analysis described earlier can then be used to generate a free energy disconnectivity graph. An example free energy disconnectivity graph for an alanine tripeptide is shown in figure 1.7.

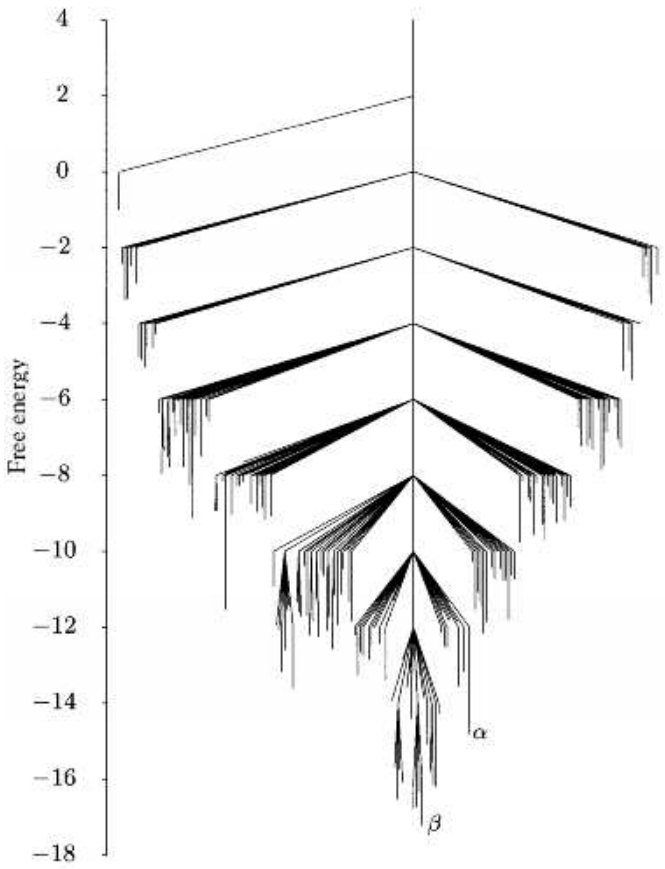


Figure 1.7: Free energy disconnectivity graph for Ac(ala)$_3$NHMe at $298K$ with a transition state shift (f_c) of $kT \ln(h \times (5 \times 10^{12} s^{-1})/kT)$. Only the 193 all-trans minima are displayed. The energy scale is in kcal/mol [113].

1.6 The global optimisation problem

The global optimisation problem on a PES is an important problem since it gives important information about equilibrium structures. In addition for sampling the PES it is very useful to be able to start at the global minimum if it is known. This allows for the possibility of sampling for low temperatures in the global minimum tree on the disconnectivity graph which will be of central importance when attempting to predict thermodynamic properties from the PES. The problem of global optimisation is an interesting mathematical problem in itself and has produced a vast body of research throughout the literature in areas as diverse as the design of computer chips, finance and protein structure prediction [12]. Many of the approaches proposed in the literature are specific to a single or at most a few of the applications and as such are poorly transferable between disciplines. The methods reviewed here can be regarded as unbiased since no user expertise has been used to bias the search toward a suitable region of the configurational space. These methods are therefore more general than biased methods. Examples of biased methods include the extensive use of databases of known protein structures which are used to bias the search of conformational space in the protein structure prediction problem.

An important and useful concept when dealing with global optimisation problems is that of computational complexity theory. This theory allows categorisation of such an optimisation problem based on worst case bounds on the amount of time required to identify the global minimum. The following is a brief review of some of the relevant points in computational complexity theory and how it applies to global optimisation on a PES.

1.6.1 Complexity studies of global optimisation on the PES

The literature on energy landscape theory and global optimisation of energy landscapes contains many references to computational complexity theory. It is useful therefore to briefly summarise some of the principal ideas of this theory. However it is important to bear in mind that computational complexity theory refers to worst case situations because it is completely general.

Background complexity theory

If a problem is to be solved with the aid of a computer then it must be possible to solve it using an algorithm, i.e. a precisely stated procedure or set of instructions that can be applied in the same way to all instances of the problem [114]. In principle therefore any problem for which an algorithm can be devised can be solved mechanically.

A famous problem in computational complexity theory is the Königsberg bridge

problem (see figure 1.8). The most obvious approach to solving this problem is to exhaustively list all the possible paths that cross all the bridges and to eliminate all that cross any bridge more than once. This approach is known as exhaustive search. However Leonhard Euler devised another method of solving the Königsberg bridge problem. He showed that a suitable solution will exist if the graph meets the two following conditions:

- It must be possible to move from any point on the graph to any other point by following the lines of the graph.
- With two possible exceptions every point on the graph must be at the junction of an even number of lines.

Euler's method is far more efficient for solving such a problem than an exhaustive search approach. A way of evaluating the usefulness of such approaches is to consider how the size of the problem increases with increasing degrees of freedom in the problem statement. As the size of the graph in the Königsberg bridge problem increases the lists of possible routes produced in an exhaustive search grows exponentially in the number of vertices and edges. Exponential growth can take the form 2^N where N represents the number of vertices and edges. Functions with even worse growth include N^N and $N!$. These functions are categorised as having exponential growth.

Another category is that of polynomial growth, where the problem size scales as a polynomial function of the degrees of freedom of the problem. An important distinction between exponential and polynomial growth is that exponential type growth will always overtake and remain greater than polynomial type growth for large enough N even if for small N the polynomial is greater than the exponential (see figure 1.9). The exact functional form of the polynomial or exponential only affect at what size (N) the polynomial growth is exceeded by the exponential growth.

The mathematical properties of algorithms was studied in the $1930's$ by the mathematician A. M. Turing. As part of his research he devised a conceptual machine called a Turing machine [115]. The machine was an automaton with an infinite supply of paper tape marked with square regions. The Turing machine was capable of performing just four operations, i) it could move the tape a single square; ii) it could place a mark in the current square; it could erase a mark in the current square if one was already present and finally iii) it could stop. Turing proved that this machine could solve any problem for which an algorithm was devised. Turing also demonstrated that there are problems that cannot be solved by Turing machines, and the resulting conclusion is that there are problems which cannot be solved by any automatic computer (i.e. the category for which no algorithm can be written even in principle).

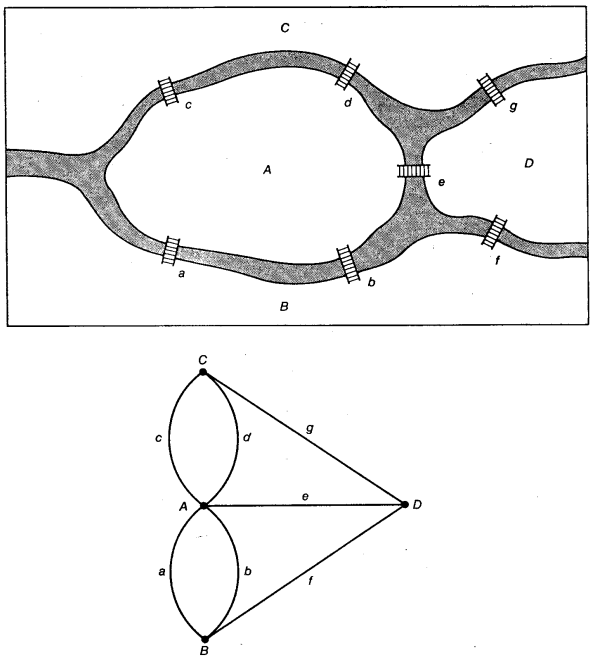


Figure 1.8: The Königsberg bridge problem. The problem was first stated in terms of a walk through a park in the 18^{th} century city of Königsberg (top). The problem is to decide whether it is possible to walk through the park using a route that crossed each of the bridges only once. The problem can be represented as a graph in which the vertices represent the two river banks and the two islands and the edges represent the bridges connecting them (bottom). By counting the number of lines meeting at each point it can be shown that the Königsberg park does not have such a path [114].

An important result of Turing's work was to divide all possible problems into two classes. The first is those problems for which algorithms cannot be written in a formal sense and remain permanently unsolvable. In certain instances these problems may be solved through luck or insight, but a general method which works for every problem in this category will never be found. The second category which contains all other problems in mathematics and logic can be solved by algorithms. This class can be further divided into two subgroups: those for which efficient, (polynomial time) algorithms exist and those for which only non-efficient (exponential time) algorithms are known.

An example of an exponential time problem was addressed by William Rowan Hamilton in the 19^{th} century. Hamilton's problem is to decide whether a graph has a path that reaches each point exactly once. Although sounding initially very similar to Euler's problem this is a much more difficult problem, and any such path is now called a Hamiltonian path. Hamilton's problem can be solved by exhaustive search however no efficient algorithm has been found.

Problems which are known to have polynomial time solutions such as Euler's problem are members of the class P (polynomial). However Hamilton's problem is a member of another class NP (nondeterministic polynomial). Problems in com-

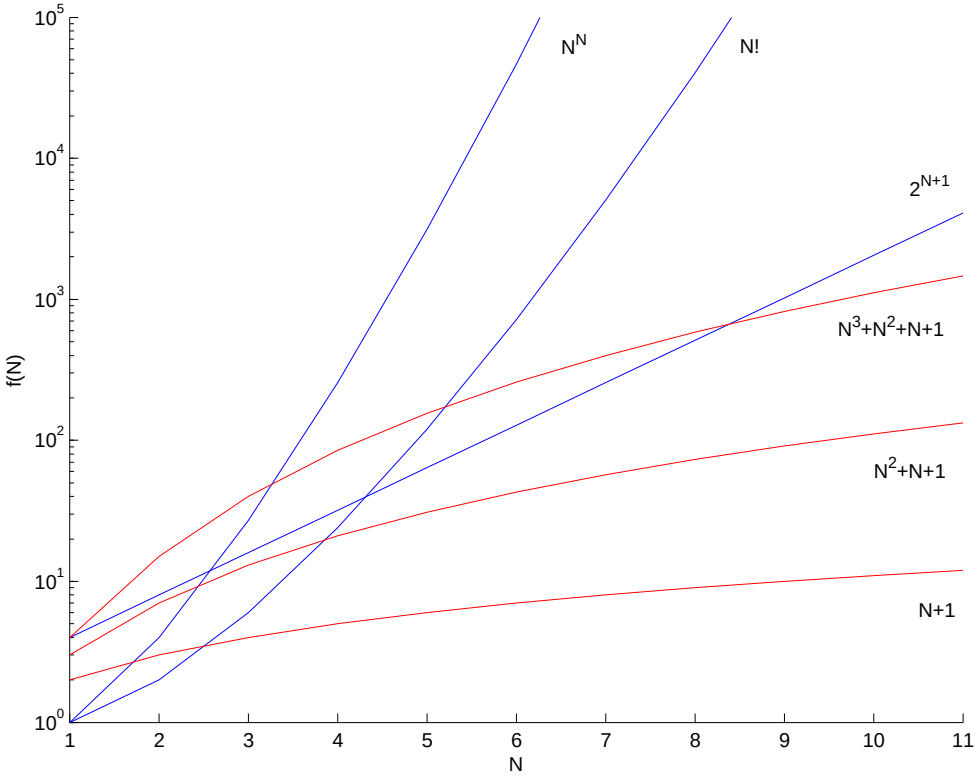


Figure 1.9: Polynomial (red) versus exponential growth (blue). The graph shows the growth in the execution time of an algorithm (f(N)) as the size of the problem (N) increases. Note that for small problems exponential algorithms may be quicker than polynomial algorithms, and there is a threshold problem size at which a given polynomial algorithm becomes more efficient than an exponential algorithm.

putational complexity theory are described as decision problems, since they can be described as an infinite set of questions with yes/no answers. Problems in the class P can have their yes/no questions answered efficiently. However problems in NP need not have an efficient means of answering their decision problems, but what is required is an efficient way of verifying that the answer is correct. These verifications are called proofs or certificates and must be polynomial. For example in the case of a graph there is still no efficient way to verify whether it has a Hamiltonian path but if one is found it can be verified in polynomial time simply by listing the edges in the path in the order in which they are traversed. Therefore NP can be regarded as the class of decision problems that can be solved by guessing certificates. If the NP problem instance under consideration has a yes answer, then with some luck it is possible to discover the required certificate by making a sequence of guesses. However if the answer is no, guessing cannot yield an answer any faster than exhaustive search would.

Turing devised a method of guessing by introducing the non-deterministic Turing machine. This is identical to a Turing machine except that at certain points in its operation it has a choice of more than one move to make next. A non-deterministic

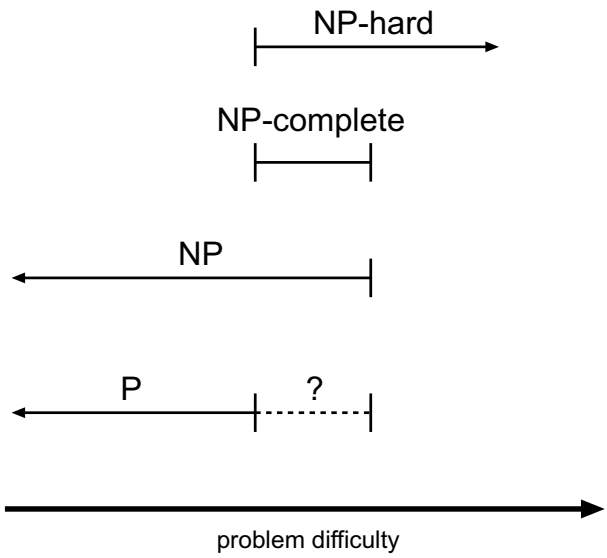


Figure 1.10: Complexity problem classifications. Decision problems with known polynomial time algorithms are classed as P problems. Examples include linear programmes. The class NP are those problems that can be solved by a non-deterministic Turing machine. The NP-complete class of problems is a subset of NP with the property that if any one of these problems could be solved with an efficient algorithm then all other problems in NP could be solved efficiently. The identification of a polynomial time algorithm for an NP-complete problem would mean that the sets P and NP are identical as depicted by the dotted line. NP-hard problems are those problems which are at least as hard as NP-complete.

Turing machine will ultimately answer yes to a decision problem if there exists some sequence of choices which can lead it to a yes conclusion. These problems exactly define the class NP. Therefore it is from the non-deterministic Turing machine that the category non-deterministic-polynomial time gets its name. The class NP includes the class P as a subset.

A problem A in computational complexity theory is polynomially reducible to a problem B if there exists a polynomial algorithm which converts problem A to problem B. A problem A is NP-complete if every other problem in NP is polynomially reducible to A. Therefore any NP-complete problem is as difficult as all other problems in NP.

The class NP-hard contains problems at least as hard an any NP-complete problems but which have not been shown to be in NP. It is therefore a more general class and contains the class of NP-complete problems. The relationship between NP, NP-complete and NP-hard can be summed up in the following relation [116]

$$\text{NP-complete} = \text{NP} \cap \text{NP-hard}. \tag{1.25}$$

NP-hard problems are exponential time problems for which there is no polynomial time verification procedure that the solution obtained is correct.

The theory of NP-completeness was developed for discrete decision problems like

Hamilton's problem. An optimisation problem is not immediately recognisable as a decision problem. Therefore optimisation problems such as the famous travelling salesman problem are not addressed within the NP-complete framework. The travelling saleman problem states: "Given a set of roads between y cities, find the shortest route that visits all of them". However this can be transformed into a decision problem by repeatedly asking the question "is there a route shorter than d?" for decreasing d, where d is a distance. Each of these decision problems is NP-complete, therefore the optimisation problem is NP-hard. The complexity classifications are categorised in terms of their difficulty in figure 1.10.

1.6.2 Potential energy landscapes and NP-hardness

There have been a number of attempts in the literature to address important questions relating to the scaling of minima on a potential energy surface using computational complexity theory but two problems in particular have been studied by many groups, namely protein folding and spin glasses.

Barahona applied computational complexity theory to Ising spin glass models [117]. He showed that finding the ground state of an Ising spin glass in $3D$ is NP-hard. In addition he demonstrated that the problem of computing the magnetic partition function of a $3D$ Ising spin glass is also NP-hard. Prior to this Bieche et al. [118] demonstrated the existence of a polynomial algorithm for the $2D$ Ising spin glass problem.

Much of the results for protein folding have been established by discretising the problem into a combinatorial optimisation problem which is best suited to treatment using computational complexity theory. Ngo and Marks showed that a typical statement of the protein folding problem using an empirical potential energy landscape is an NP-hard problem [119]. They achieved this through proving reducibility between two problems. The first which they called "diamond lattice path" (DLP) is an abstract problem based on the conformational geometry of alkanes. They then proved a reduction of a known NP-complete problem called "partition" to DLP. They further demonstrated reductions from DLP to two more general problems "endpoint-constrained polymer structure prediction" and "polymer structure prediction" which can be considered cases of potential energy minimisation for polymers or proteins. However in an effort to maintain generality an important consideration that was neglected by Ngo and Marks is the effect of hydrophobicity and compactness on the number of conformations accessible to a protein chain.

Unger and Moult [120] demonstrated reducibility from an NP-complete problem "optimal linear arrangement" to a discretised protein structure prediction problem "decision protein folding" in which only the positions of the C_α atoms were considered. Istrail and Hart subsequently proved that for any reasonable lattice the

protein-structure prediction problem is NP-hard [121].

Finding the global minimum on a PES is therefore a very challenging problem. Not only is there no known polynomial algorithm to locate the global minimum but there is also no way of verifying that the minimum is in fact the global minimum without exhaustively evaluating all the minima on the PES. All of the methods to be described in the following section are exponential time methods, although it must be emphasised that some perform very well in practice.

1.7 State of the art approaches for global optimisation on a PES

Computational complexity theory provides a worst case scenario for the complexity of global optimisation problems. In reality the problem of finding the global minimum at least for some simple systems in chemistry such as Lennard-Jones clusters seems to be tractable using sophisticated sampling algorithms and modern computing resources. However in light of the last section, proving that the global minimum has been identified is a far more challenging problem.

The goal of this section is to review the methodologies which are currently state of the art for global optimisation within the framework of energy landscape theory. These methodologies can be divided into two broad categories, stochastic and deterministic approaches. With the exception of one of the deterministic methods, all are heuristics and share the common property that they require sampling of points on the PES. How that sampling is carried out is primarily what differentiates the methods. Because the sampling is stochastic in nature these algorithms are often labelled as stochastic approaches.

The most straight-forward stochastic approaches are those based upon simulated annealing in which statistical mechanical techniques were first used to solve combinatorial optimisation problems. Another category are the hypersurface deformation approaches in which the non-convexity of the PES is incrementally reduced leaving a simpler problem to solve. The category of evolutionary algorithms take inspiration from the process of evolution in nature. Finally deterministic methods such as the αBB branch and bound algorithm [122] provides certificates that the global minimum is not in certain regions of configurational space. In the following paragraphs a description of each category is made. Their advantages and shortcomings are discussed. In many cases the algorithms have been tested on Lennard-Jones clusters as these have become an informal benchmark of the quality of a particular sampling algorithm.

The potential energy surfaces of Lennard-Jones (LJ) clusters have been explored with some of the following algorithms for clusters as large as 309 atoms [123–134].

Most LJ clusters have global minima based on a Mackay icosahedron [135] and with the exception of a few unusual cases, have minima that are reasonably easy to locate with the best of the sampling algorithms currently available. For magic number sizes (sizes at which a complete icosahedral cluster can be formed) the global minimum is a complete Mackay icosahedron, with intermediate sizes typically having a Mackay icosahedron with a partially completed outer layer. For other sizes the icosahedral structure must compete with other low-lying structures which can be based upon a truncated octahedron, Marks decahedron, or a Leary tetrahedron, for example, (see figure 1.11). In these cases the PES typically has a disconnectivity graph with

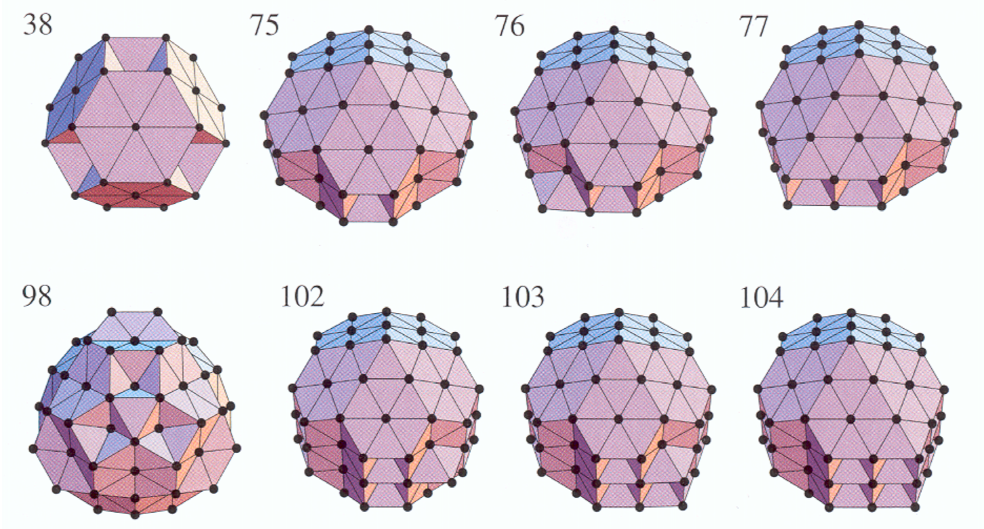


Figure 1.11: Nonisocahedral global minima of Lennard-Jones clusters with $N < 110$. For $N = 38$, the global minimum is a truncated octahedron, for $N = 75-77$ and $102-104$, a Marks decahedron, and for $N = 98$, a Leary tetrahedron [12].

double palm tree topology with the separate minima lying at the bottom of their respective funnels. The liquid state and therefore the vast majority of local minima are based on polytetrahedral packing which is closer to an icosahedral structure than to any of the other structures. Therefore finding these non-icosahedral global minima using sampling techniques can prove difficult even for small systems such as LJ_{38} where the global minimum is a face-centred-cubic truncated octahedron (see figure 1.12). The problem is made more difficult by the relative sizes of the competing funnels. Because the non-icosahedral funnel contains typically significantly less minima than the icosahedral funnel, the icosahedral funnel benefits from having greater configurational entropy. Interestingly it has been found that applying a compressive force to the cluster increases the likelihood of locating such non-icosahedral global minima. This is because compression biases the sampling away from the more spherical icosahedral structures [136]. In fact it took a number of years for the cluster community to discover the truncated octahedron (which was discovered due to improved sampling routines and increasing computer power). Prior to this

the isocahedral structure was believed to be the global minimum. Perhaps there is yet another yet lower energy structure remaining to be discovered, though this is unlikely at least in this case.

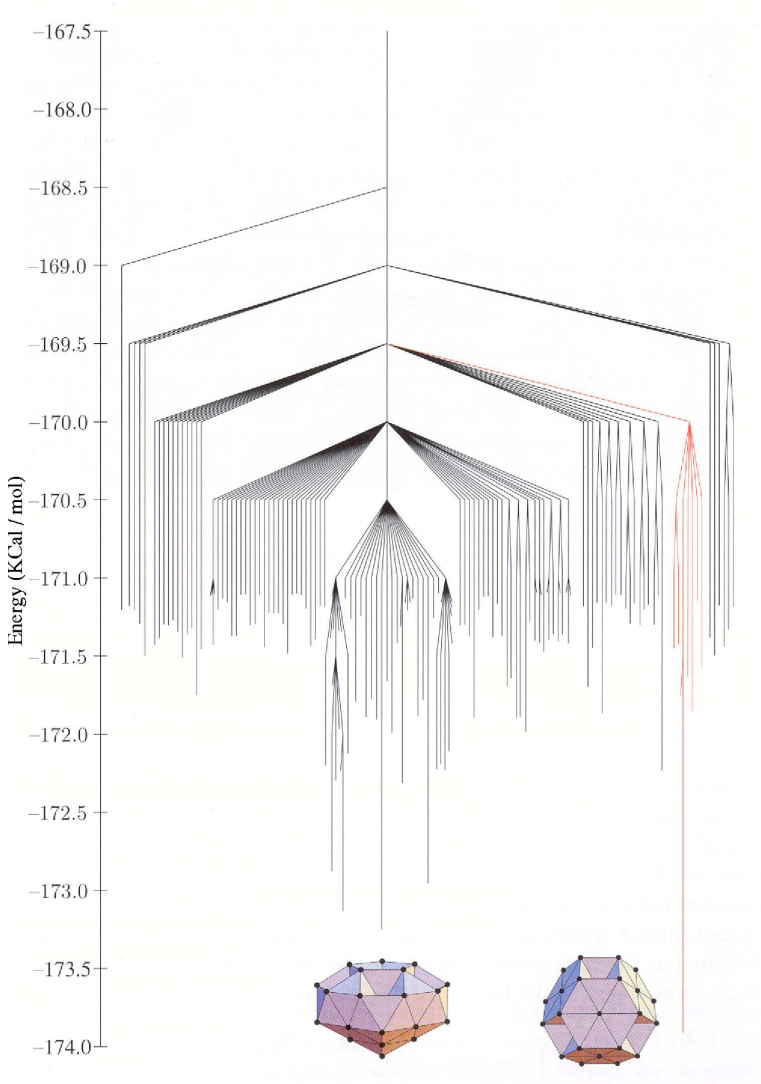


Figure 1.12: The lowest 200 minima on the disconnectivity graph for the LJ_{38} cluster. The truncated octahedron global minimum lies at the bottom of a narrow funnel (red) [137].

1.7.1 MD based annealing

Simulated annealing

The simulated annealing algorithm of Kirkpatrick, Gelatt and Vecchi in 1983 was probably the first significant attempt at global optimisation of combinatorial problems based on statistical mechanical principles [138] and therefore it should be a good candidate algorithm for finding the global minimum on a PES. It is based on the principle of annealing where a metal is heated to a high temperature and cooled slowly allowing for a more crystalline state to be formed. In simulated annealing the

temperature of the system of interest is raised and then slowly cooled in the hope that the free energy minimum tends toward the global minimum on the potential energy surface as the cooling progresses. It is the slow cooling that prevents the system from leaving equilibrium. Sampling is carried out using a Metropolis sampling scheme [139] and a necessary and sufficient condition for the simulated annealing algorithm to arrive at the global minimum is for the temperature to decrease logarithmically with time [140–147]. This makes the algorithm very slow and alternative compromise approaches have been suggested. However there is a risk that because of rapid cooling the system may become trapped in the wrong side of a barrier that is increasing relative to kT as the temperature is reduced. This is particularly so when PES's with multiple palm tree (funnels) or banyan tree disconnectivity graphs are being studied.

1.7.2 Hypersurface deformation methods

The central theme in hypersurface deformation methods is the successive relaxation of the potential surface to a more convex surface. The effect of this is to reduce the number of minima on the PES thereby reducing the likelihood of getting trapped behind barriers as the temperature of the system is reduced. However deforming the surface almost always makes it impossible not to deform the global minimum, without knowing its location a priori, and therefore it is often necessary to perform some reverse mapping procedure once the relaxation has been completed. Both the distance scaling method and the diffusion equation method require a reverse mapping procedure.

Diffusion equation method

To understand the deformation in the diffusion equation method (DEM) it is useful to consider the one dimensional case. Here the potential is deformed as

$$U^{(1)}(x) = U(x) + \beta U''(x) \quad \text{for } \beta > 0, \tag{1.26}$$

which if repeated N times yields

$$U^{(N)}(x) = \left(1 + \beta \frac{d^2}{dx^2}\right)^N U(x) \quad \text{for } \beta > 0. \tag{1.27}$$

As a result of such a transformation the inflection points remain the same while regions where the function is convex increase and regions where the function is concave reduce toward the value at the inflection point. β is then chosen to be

$\beta = t/N$, in which case equation 1.27 is transformed to

$$\exp\left(t\frac{d^2}{dx^2}\right)U(x), \tag{1.28}$$

where $t > 0$ is a parameter. The exponential can be represented by the Taylor series

$$1 + A + \frac{1}{2!}A^2 + \frac{1}{3!}A^3 + \cdots, \tag{1.29}$$

where $A = t(d^2/dx^2)$. The effect of the successive operation of the exponential on $U(x)$ is that high frequency fluctuations in the PES flatten more pronouncedly than lower frequency fluctuations. However the Taylor series (equation 1.29) can be divergent for certain $U(x)$ and sufficiently large t, but when it does converge it can be verified that its sum is a solution of the diffusion equation

$$\frac{\partial^2\hat{U}}{\partial x^2} = \frac{\partial\hat{U}}{\partial t}, \tag{1.30}$$

where t represents time and $\hat{U}(x,0) = U(x)$. The solution of equation 1.30 is followed instead of applying the exponential operator.

For t sufficiently large the function $\hat{U}(x, t_{max})$ will hopefully have a single minimum. However its position may well be different to the global minimum of $U(x)$. This minimum is then used as a starting point for a reversing procedure. Here the deformation is incrementally reversed in the series $\{\hat{U}(x, t_{max}), \hat{U}(x, t_{max} - \Delta), \hat{U}(x, t_{max} - 2\Delta), \ldots, \}$ where Δ is a reversing increment and a minimisation is carried out using the minimum point from the previous surface as a starting point. Hopefully when the reversal is complete the current configuration will lie at the global minimum. The deformation and reversal procedure is summarised in figure 1.13.

Distance scaling method

In the distance scaling method (DSM) which has been applied principally to Lennard-Jones clusters [128], a molecular dynamics simulation is performed using a deformed potential energy surface. The standard potential energy function $U(x)$ is deformed using a deformation parameter $t > 0$ to $U(x,t)$, such that $U(x,0) = U(x)$. For sufficiently large $t = t_{max}$, a single basin containing the global minimum of $U(x)$ is hoped for. Following the deformation again a reversing procedure is applied, in which $U(x, t_{max} - \delta)$, $0 < \delta \ll 1$, is minimised with respect to x using the minimum of $U(x, t_{max})$ as the starting position. One of the key differences between the DEM and DSM methods is that the minimum of the Lennard-Jones potential does not change throughout the deformation but the range of the potential does

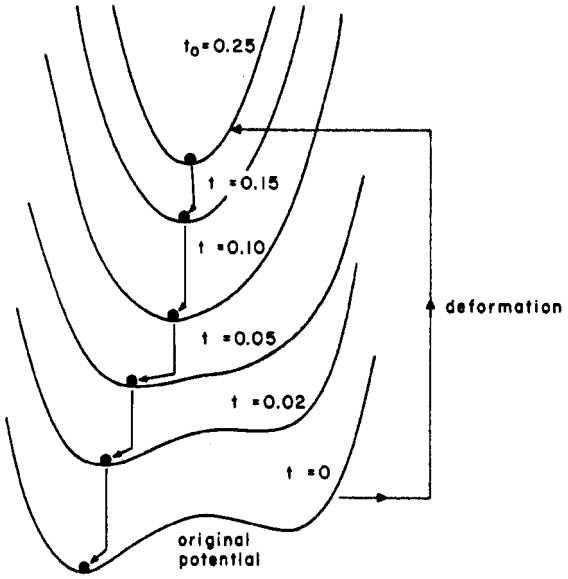


Figure 1.13: Deformation of the original surface followed by the reversing procedure in the diffusion equation method. Deformation $t_{max} = 0.25$ results in a curve with a unique minimum. This is followed by a reversal procedure in which at each step minimisation relaxes the system into the nearest minimum [148].

by necessity since the deformed potential remains an underestimator of the original pairwise Lennard-Jones function. The Lennard-Jones potential is expressed in the scaled distance r' where

$$r' = \frac{r + r_0 t}{1 + t}, \tag{1.31}$$

and r_0 is the equilibrium distance of the unmodified Lennard-Jones potential. Therefore as t increases the Lennard-Jones potential become progressively more convex but the minimum remains in the same position. The DSM has succeeded in locating the global minimum of the LJ_{38} cluster but has not been as successful for other larger clusters with non-icosahedral global minima [12].

Basin hopping

A hypersurface deformation algorithm that has been considerably more effective than either DEM or DSM methods is the basin hopping algorithm of Wales which has been used to identify global minima for many Lennard-Jones clusters [123, 131, 149–151]. In the basin hopping algorithm a discrete sampling of the entire basin in which the current configuration lies is made and the potential energy at all these points is mapped to that of its local minimum through a steepest descent minimisation according to

$$\tilde{U}(\mathbf{x}) = \min_{\text{steepest descent}} \{U(\mathbf{x})\}. \tag{1.32}$$

The landscape is therefore transformed from a rugged topography to one of a multidimensional staircase (figure 1.14).

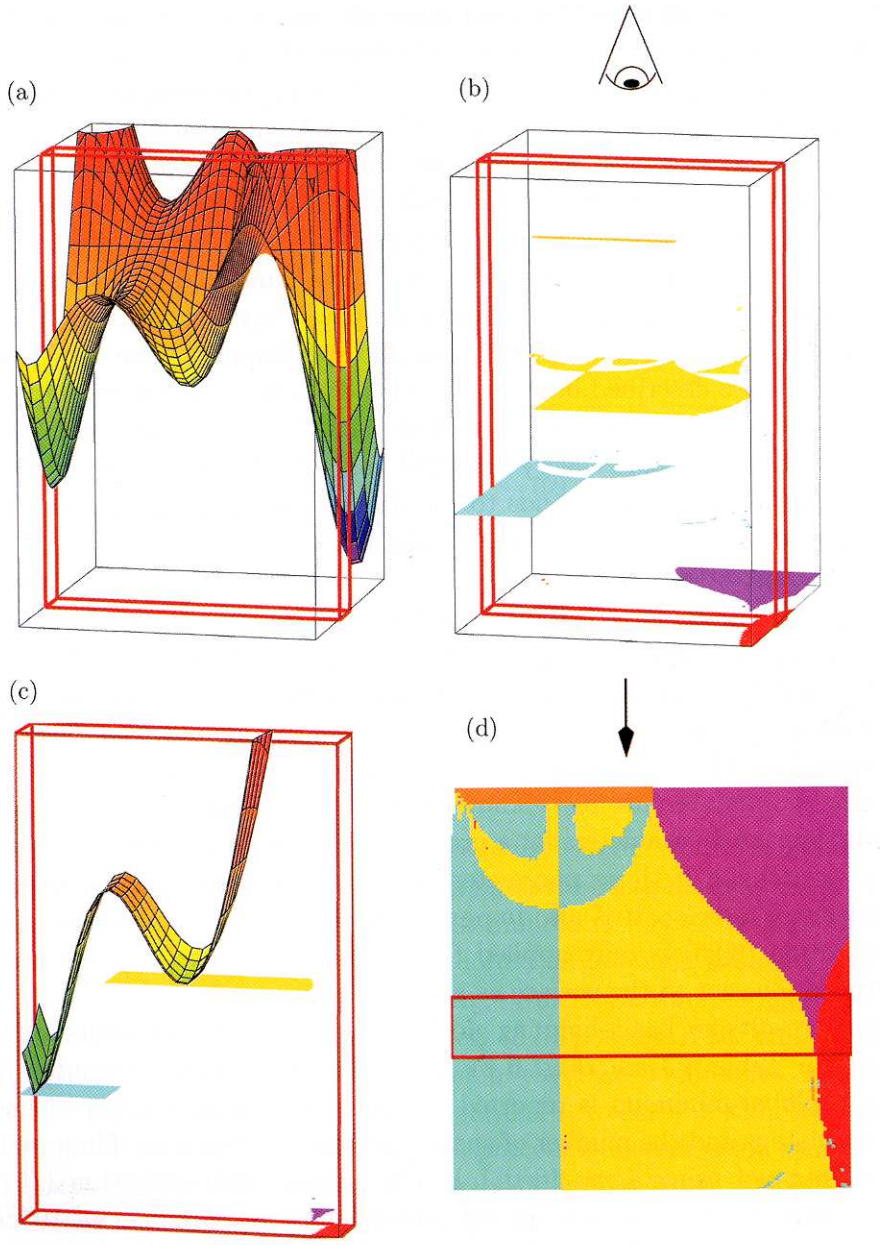


Figure 1.14: The basin hopping transformation of a PES to a staircase energy landscape. The original $2D$ surface is shown in (a). The transformed surface is shown in (b) where the colours of the "stairs" correspond to energy as in the original surface (a). Image (c) shows a section through the surface as highlighted in red in (b). The view of the transformed surface from above is given in (d) [152].

The basin to basin transition is governed by a Monte-Carlo-like sampling scheme in which steps are taken by adding a perturbation to the current coordinates which lie in minimum *old* with minimum potential energy U_{old} and carrying out a minimisation to obtain potential energy U_{new}. The step is accepted if $U_{old} \geq U_{new}$ or $U_{old} < U_{new}$ and $\exp((U_{old} - U_{new})/kT) > R$ where R is a random number in the interval $[0, 1]$.

There are two possibilities once the move has been decided: i) the coordinates are set to that of the local minimum it is in or ii) the coordinates are allowed to vary continuously.

An important effect of the basin hopping mapping to a staircase landscape is to remove transition regions between adjacent basins. This has the effect of accelerating dynamics on the transformed PES. Notice that atoms can effectively pass through each other without hitting infinite energy barriers. However the reason why basin hopping is particularly effective at locating nonicosahedral global minima is due to thermodynamic considerations [153, 154]. On the original PES the occupation probabilities of the competing structures as a function of temperature have little overlap. However as a result of the transformation of the PES, the occupation probabilities are broadened meaning that there is an increased chance of transitioning, for example, from the icosahedral cluster funnel, after the free energy global minimum has switched from the icosahedral to the non-icosahedral structure.

In addition to clusters the basin hopping algorithm has also been applied to peptides, polymers and glasses [123, 155–165].

The method also has some significant disadvantages. The fact that the PES has been transformed to such an extent means that all dynamical information is lost. Additionally being a sampling approach no guarantees can be produced about the global minimum.

Monte Carlo with minimisation

The Monte Carlo with minimisation algorithm of Li and Scheraga [166, 167] is equivalent to the basin hopping algorithm where the updated coordinates are set to those of the local minimum.

1.7.3 Genetic algorithms

The idea for genetic algorithms was derived from the process of natural selection which drives evolution in the natural world [168, 169]. In natural selection a member of a species has a particular genetic code which results in the expression of that code's phenotype (the organism). The relative fitness of the different phenotypes determines which organisms get to pass on their genes to their progeny thereby ensuring the survival of their geneline.

It is easiest to describe genetic algorithms in terms of an application such as atomic cluster geometry optimisation. In cluster optimisation, initially, the phenotype is the 3 dimensional configuration of the cluster which is expressed in a binary string "gene" which is acted upon by particular operators [170]. Among these operations are crossover operations, a binary operation in which two genes are broken at the same point and the ends swapped such that the first part of gene 1 is connected

to the second part of gene 2 and vice versa. Mutation is a unary operation in which a binary flip is performed at randomly selected points in a gene. At the start of an optimisation a "population" of structures is established along with their associated genes. The fitness of a particular "organism" is determined by its potential energy, the lower the energy the fitter the organism. Using potential energy to select the fitter configurations, mating is permitted between fit configurations while the less fit configurations are discarded. This is the traditional genotype algorithm.

In fact it has been demonstrated that it is computationally more efficient not to use binary encoded genes and apply the mutation and crossover operators directly to the structures [171, 172]. This type of algorithm became known as a phenotype algorithm [173].

The first application of genetic algorithms to PES minimisation was performed by Deaven and Ho where they succeeded in locating new global minima for LJ_{88} and LJ_{98} clusters [131, 174]. Subsequent studies have used different variations on the mutation and crossover operations with some considerable success. Some of these variations were compared by White, Niesse and Mayne [175]. In particular they observed considerable performance difference between traditional genetic algorithms (TGA) and modified genetic algorithms (MGA), where TGA represents algorithms where the coordinates for each atom are encoded in binary numbers with a single-point crossover operation. TGA therefore do not mix distance mutations with angle mutations. MGA algorithms include multipoint crossovers and inversions allowing a mixing of distance and angular variables. Space fixed modified genetic algorithms (SFMGA) refers to modified genetic algorithms in which the gene consists of real numbers representing the positions of the atoms in the cluster. In their SFMGA algorithm they also carried out gradient minimisations from a mutated point in configurational space to its local minimum. In their comparison's they found similar performance between SFMGA and basin hopping.

Other groups have combined a number of genetic algorithm techniques into a single computational tool [176, 177] which also includes different functional forms for defining fitness. Using such an approach genetic algorithms have succeeded in finding the lowest minima found by the basin hopping method [12].

1.7.4 Hybrid algorithms

Approaches that are a combination of various other approaches or that have other purposes have been termed hybrid approaches.

Conformational space annealing

Conformational space annealing (CSA) is a hybrid approach of the build up procedure and genetic algorithms which has been applied to finding the global minimum

on the PES of polypeptides and simple proteins. In the build up procedure the original peptide is broken down into several levels of component peptides for which the exhaustive enumeration of the minima can be performed. Various strategies have been devised to eliminate many of the higher energy minima to deal with the exponential growth of minima.

As in genetic algorithms, the CSA algorithm is started with a random population of conformations whose energies are then minimised. These local minima are referred to as the bank. The conformations are refined in such a way as to maximise the distance between them in conformational space so that a large region of conformational space is sampled. Each conformation in the bank is regarded as representing a group of local minima close to it in conformational space. Retaining an arbitrary N conformations in the bank the following two operations are considered: i) Each conformation may be replaced by a lower energy conformation in its representative group. ii) When a conformation representing a new group is located, then the conformation in the bank with the highest energy may be replaced by this new conformation.

In order to determine whether a new conformation belongs to a current group or to a new group, a cutoff distance D_{cut} in conformational space is specified which defines the size of each group. At the beginning of the run D_{cut} is large enough so that the N groups cover the whole of relevant conformational space. As the algorithm progresses D_{cut} is successively reduced and a progressively smaller region of conformational space is sampled.

CSA has been applied to met-enkephalin [178] and a 20 residue portion of melittin [179]. It has also been used in a hierarchical approach to protein structure prediction [180] which was entered in the CASP3 project [181].

Sampling methods

The thermodynamic sampling methods described in section 1.5.3 can also be used to locate the global minimum on a PES since the global minimum must be at least as low as the lowest minimum sampled during a simulation. Methods such as jump-walking have been used to identify global minima for Lennard-Jones and water clusters [38] while parallel tempering has been applied to small peptides [182].

1.7.5 Deterministic methods

An important difference between stochastic and deterministic global optimisation methods is that deterministic global optimisation methods have in-built mathematical guarantees that the global minimum will be located if the algorithm is allowed to run to completion. One of the best known deterministic global optimisation algorithms is the α Branch and Bound algorithm of Androulakis, Maranas and

Floudas [122]. Other methods include convex global underestimator algorithm of Phillips, Rosen and Walke [183] which was developed for potential energy surfaces with a single funnel and the terminal repeller unconstrained subenergy tunnelling (TRUST) approach of Barhen, Protopopescu and Reister [184].

Branch and bound

The branch and bound method is a deterministic global optimisation algorithm suitable for general continuous optimisation problems in which non-convexities in the objective or constraints exist. The main idea of the approach is to seek upper and lower bounds to the global minimum. These bounds are successively improved through iterative partitioning of the problem domain. The lower bounds are obtained as a set of solutions to special convex lower bounding functions, which can be constructed by augmenting the potential U with separable quadratic functions for example,

$$L = U + \sum_i \alpha_{\phi_i} (\phi_i^L - \phi)(\phi_i^U - \phi), \tag{1.33}$$

in the case of a potential considering only dihedral angles ϕ_i in a peptide or protein. Here ϕ_i^L and ϕ_i^U are the lower and upper bounds on the dihedral angle. The α_{ϕ_i} represent non-negative parameters which must be greater or equal to the negative one-half of the minimum eigenvalue of the Hessian of U over the entire domain. Various approaches have been developed to estimate these parameters. The upper bound can be obtained by a local minimisation of U. Once these bounds converge the exact global minimum has been identified.

The partitioning of the domain involves successive subdivision of the original hyper-rectangle by bisecting at the midpoint of its longest side. To ensure that the lower bounds are nondecreasing, the half with the infimum of the lower bounds is further bisected. A one dimensional example is given in figure 1.15. The branch and bound method has been used to study small peptides [185–188], and clusters [189] and is very effective in systems small enough for it to be computationally effective. A particularly attractive aspect of the approach is the proof of optimality for the purported global minimum. In addition to obtaining global minima, branch and bound also allows for the identification of all local minima and transition states within a range ΔU of the global minimum. Nevertheless despite these positive aspects, branch and bound remains an exponential time algorithm. In addition it can struggle to find the global minimum on PES's with very localised but deep wells.

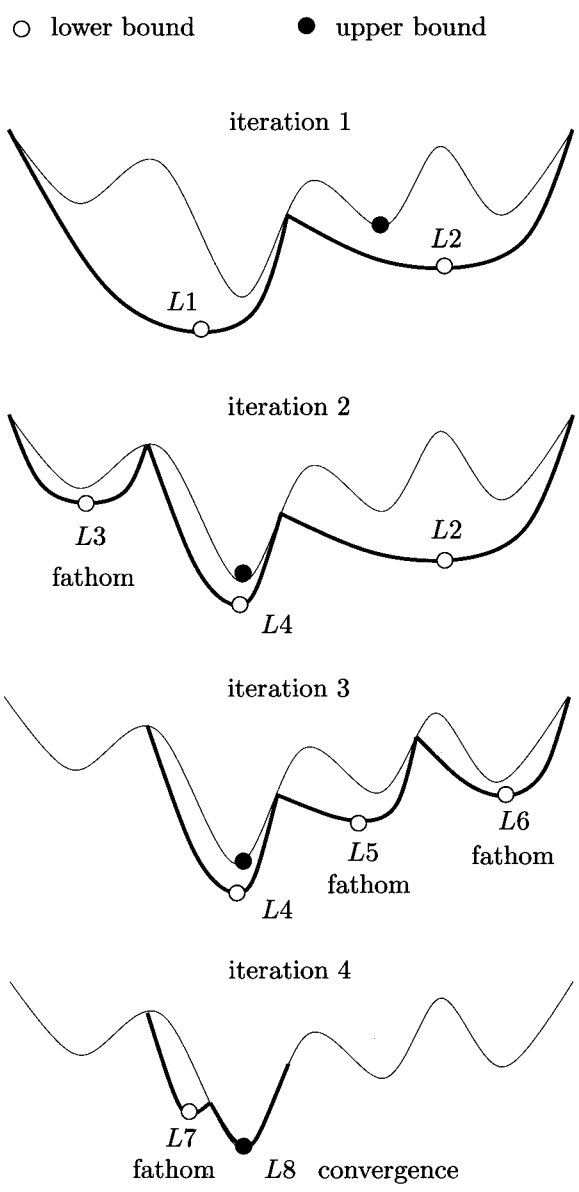


Figure 1.15: One-dimensional branch and bound example. Iteration 1: The overall domain is bisected and two convex lower bounding functions are generated and minimised to obtain two lower bounds L_1 and L_2 for their respective regions. An initial upper bound is obtained for the entire domain. Iteration 2: Since $L_1 < L_2$, region 1 is further bisected into regions 3 and 4, while region 2 from the 1^{st} iteration is stored. Lower bounds for these regions are obtained (L_3, L_4). A new upper bound is generated for region 2. Since L_3 is greater than the upper bound for region 2, region 3 can be discarded (fathomed). Iteration 3: The region with the next lower bound is bisected into regions 5 and 6 and lower bounds obtained. Since L_5 and L_6 are both greater than the current upper bound the entire region 2 is discarded. Iteration 4: Region 4 is bisected giving regions 7 and 8 for which lower bounds are obtained. Region 7 can be immediately discarded and since the upper and lower bounds for region 8 coincide, the global minimum has been obtained [185].

Convex global underestimator

The convex global underestimator method [183] fits a quadratic function to a set of points sampled from the PES. The minimum of this quadratic function then becomes the centre of another smaller box in conformational space in which further

sampling takes place and another quadratic function is fitted through these points. The approach will work for PES's with a single funnel, but is unlikely to be effective for surfaces with more complex topologies.

1.8 Conclusions and objectives

In this chapter the main components of energy landscape theory for studying potential energy surfaces have been introduced. A number of important topics were discussed,

- The physical basis of the landscape.
- What are the important features of the landscape (stationary points).
- How energy landscapes are quantitatively described.
- Approaches for studying potential energy landscapes i.e. methods for:
    1. finding minima and transition states,
    2. extracting thermodynamic information.

A key problem in energy landscape theory was also discussed namely that of finding the global minimum of the energy surface, a global optimisation problem. This problem is fundamental to energy landscape theory since any set of configurations which does not contain the global minimum cannot be regarded as representative of an equilibrium ensemble. Therefore any density of states or partition function and resulting thermodynamic quantities obtained from such a set must be questionable.

The methods described in section 1.7 are the current state of the art approaches for performing global optimisation on potential energy surfaces. In practise many of them have been thoroughly tested and seem to work quite well at least on small systems. However as has been pointed out, the global optimisation problem in general and specifically examples of global optimisation on a PES have been shown to be NP-hard. The methods described are effectively heuristic approaches which lack any guarantees in term of their computational complexity. With the exception of the branch and bound algorithm they also lack any way to prove that the global minimum has been found / sampled.

The most effective way to prove global optimality is to seek tight lower bounds on the global minimum while at the same time producing tight upper bounds which can both be progressively refined. Once these bounds are equal to each other optimality has been provably achieved. Ideally a method for finding the global minimum on a PES should have two characteristics:

1. The ability to produce progressively tighter upper and lower bounds on the global minimum, thereby verifiably obtaining global optimality.

2. A polynomial computational complexity bound on the time taken to obtain the global minimum.

Of the methods reviewed only the branch and bound method provides (1) and no method provides (2). Since the problems of interest are NP-hard, there is no algorithm that satisfies (1) and (2) (unless P=NP). This paints a rather gloomy outlook for the prospects of success in adressing the problem of provably identifying the global minimum on a PES. However there is an example of a famous discrete NP-hard maximisation problem, MAXCUT, for which excellent lower bounds have been obtained by using a relaxation approach. In their seminal paper, Goemans and Williamson [190] proved that a lower bound, guaranteed to be no more than 12.2% below the optimal value, can be obtained through a convex relaxation which is solvable as a semidefinite programme. An equally exciting development with possibly direct applications to global optimisation on a PES is a very interesting relaxation in a similar vein by Parrilo [191].

The approach applies to a specific class of optimisation problems known as polynomial optimisation problems, in which the objective and constraints are represented as polynomial functions. Polynomial optimisation problems are NP-hard for polynomials of degree greater than 4. It has been shown that to bound a given function from below, one approach is to find a certificate of nonnegativity for the function. In general it is not possible to find such a certificate. For the specific case of polynomial optimisation problems this novel approach instead seeks a sum of squares decomposition of the function which acts as a certificate of nonnegativity. It therefore addresses the problem of providing a certificate of non-negativity by relaxing the problem to one of finding a sum of squares decomposition. A sequence of sum of squares decompositions may be built up providing progressively tighter lower bounds on the global minimum.

An ingenious aspect of the approach is the way that the sum of squares decompositions are obtained. The global optimisation of the polynomial function is mapped to a problem of finding a suitable sum of squares decomposition. It has been shown that finding such a decomposition can be achieved by solving a semidefinite programme, an advanced optimisation problem belonging to the category of optimisation problems known as convex optimisation problems. Semidefinite programmes (and convex optimisation programmes in general) have been proved to have polynomial time complexity. In addition, note that the highly non-convex problem of finding the global minimum has been mapped to a convex problem.

Convex optimisation programmes are comprised of 2 separate but related problems called primal and dual programmes. Typically the primal programme is a minimisation and the dual programme is a maximisation. The solution to the primal is always greater than the solution to the dual therefore if they are equal then both problems are optimal. In the case of a semidefinite programme resulting from a reformulation of a polynomial optimisation problem as a sum of squares decomposition, the primal semidefinite programme provides the lower bound on the function while the dual problem provides the location of the optimal point.

In summary, excellent bounds on an NP-hard problem can be obtained by solving a sequence of polynomial time problems and this involves mapping from a highly non-convex problem to a convex problem in higher dimensions (convex optimisation problems will be discussed in chapter 2). This approach therefore possesses characteristic (1) in the list above while (2) is asymptotically achieved since solving a series of semidefinite programmes is still an optimisation problem with polynomial time complexity.

The main objective of this thesis is to apply this methodology (which will be referred to as sum of squares optimisation) to the problem of finding the global minimum on a potential energy surface. In particular the application of the method to a commonly used classical force field (Amber94) for a simple chemical system is described in chapter 3. This chapter places emphasis on practical aspects of the implementation of the method to energy landscape theory and difficulties encountered as well as approaches to overcome them are discussed.

A second objective of this thesis is to demonstrate that the sum of squares decomposition approach can also be used to solve a problem in coarse-grained molecular dynamics, namely the distance of closest approach of ellipsoidal representation of two molecules. It is shown that the problem reduces to solving a particular category of convex optimisation problem namely a second order cone programme. This work is described in detail in chapter 4.

These are the primary objectives of this thesis, however before these are addressed, a detailed introduction to the sum of squares decomposition approach is required. This is given in chapter 2. This chapter describes the relationship between global nonnegativity and sum of squares as well as providing a brief and basic introduction to algebraic geometry, the mathematical discipline upon which sum of squares optimisation is built.

Chapter 2

Sum of Squares Optimisation

2.1 Introduction

This chapter describes the theory behind a radically new approach to global optimisation of potential energy surfaces. It is a fundamentally different approach to those described in section 1.7 in that no sampling of configuration space is performed. Instead a semidefinite programme (see section 2.2.2) is formulated and solved. Because of the duality properties of convex optimisation problems in general and semidefinite programming problems in particular it is possible to obtain proofs or certificates of optimality like those described in section 2.2.4. The theory is largely the work of two people, Parrilo [1, 191], and Lasserre [192], who independently built on the earlier work of Shor [193]. The approach has deep theoretical foundations in the subject of algebraic geometry, a mathematical discipline which studies the relationship between sets of polynomial functions and the sets defined by the solutions of these sets.

A relaxation is performed, but one which does not deform the potential energy surface in the way that hypersurface deformation methods do. Instead it seeks a function which provides a lower bound to the function to be minimised (potential energy functional). This lower bound function is a sum of squares (SOS) and may even be identical to the function it is bounding. The definition of a sum of squares will be given in section 2.3.1. One of the powerful aspects of the method is that it allows this bound to be as tight as required, with the gap to the minimiser asymptotically approaching zero.

Finding the sum of squares decomposition is achieved by solving a semidefinite programme (SDP), a category of convex optimisation programme. Convex optimisation programmes have a special property known as duality which provides certificates that the optimal solution has been achieved. This duality also has an analogue in algebraic geometry and this correspondence becomes very important in solving a sum of squares optimisation problem. It is not possible to explain this in detail before describing the theory but at its most basic level it can be understood in the

following way: By solving a primal SDP, a lower bound on the global minimum is produced, while by solving its corresponding dual SDP the minimiser (coordinates in configurational space) of the global minimum is obtained.

This chapter begins by introducing some convex optimisation theory (section 2.2). The level of detail is only that necessary to understand how sum of squares optimisation problems are formulated and solved. The interested reader is referred to the excellent books of Boyd and Vandenberghe [194] and of Ben-Tal and Nemirovski [195] for a complete coverage of the topic.

Having completed the background on convex optimisation it is possible to discuss sum of squares theory. This begins with some remarks (section 2.3) about the relationship between global optimality and global nonnegativity. What is meant by a sum of squares is defined next (section 2.3.1). This makes it possible to introduce the relaxation which is fundamental to the approach (section 2.3.3). Section 2.4 introduces some important results in algebraic geometry which are necessary to obtain an understanding of how an optimisation problem which involves constraints such as equation 2.15 can be addressed. Algebraic geometry also provides the framework for the construction of the hierarchy of lower bound approximations to the global minimum. The next section (2.5) brings everything that has gone before together into a coherent algorithm for dealing with energy landscape minimisation problems. It describes using a flowchart how a problem which is defined by polynomial functions is systematically mapped into a problem of optimisation over sets of positive semidefinite matrices, i.e. an SDP. Finally section 2.6 introduces some modifications to the SOS approach due to Jibetean [196] which allows optimisation to be carried out over rational polynomials. This may be particularly useful for dealing with non-bonded dispersion interactions since the functional form of the empirical potential for dispersion interactions is often approximated by a rational polynomial.

In this chapter a number of theorems are stated accompanied by the relevant citation. For the sake of brevity and because this is not a mathematical dissertation the proofs are omitted. However the interested reader may refer to the cited references to see the original proofs.

2.2 Convex Optimisation

2.2.1 Preliminaries

Before delving into the theory of how to solve convex optimisation problems it is first of all necessary to introduce some terminology which makes the explanation of many concepts more straightforward.

Convex sets and cones

Convex sets A set C is convex if the line segment through any two points in C also lies in C i.e. if for any $\mathbf{x}_1, \mathbf{x}_2 \in C$ and any $0 \geq \theta \leq 1$,

$$\theta \mathbf{x}_1 + (1-\theta)\mathbf{x}_2 \in C, \tag{2.1}$$

where $C \in \mathbb{R}^n$. The convex hull of a set C denoted $\mathbf{conv}C$ is the set of all convex combinations of points in C, i.e.:

$$\mathbf{conv}C = \{\theta_1 \mathbf{x}_1 + \ldots + \theta_k \mathbf{x}_k \mid \mathbf{x}_i \in C,\ \theta_i \geq 0,\ i = 1 \ldots k,\ \theta_1 + \ldots \theta_k = 1\}. \tag{2.2}$$

This is the smallest convex set that contains C.

Convex cones A set C is a cone if for every $\mathbf{x} \in C$ and $\theta \geq 0$, $\theta \mathbf{x} \in C$. Such a set is called a cone. A cone is convex if for any $\mathbf{x}_1, \mathbf{x}_2 \in C$ and $\theta_1, \theta_2 \geq 0$, we have

$$\theta_1 \mathbf{x}_1 + \theta_2 \mathbf{x}_2 \in C. \tag{2.3}$$

Examples

The second order cone The second order cone is a cone based on the Euclidean norm:

$$C = \left\{(\mathbf{x}, t) \in \mathbb{R}^{n+1} \mid \|\mathbf{x}\|_2 \leq t\right\}. \tag{2.4}$$

The positive semidefinite cone The set of symmetric $n \times n$ matrices is denoted

$$\mathbb{S}^n = \left\{\mathbf{X} \in \mathbb{R}^{n \times n} \mid \mathbf{X} = \mathbf{X}^T\right\}, \tag{2.5}$$

which is a vector space of dimension $n(n+1)/2$. A subset of this is the set of positive semidefinite matrices denoted

$$\mathbb{S}^n_+ = \{\mathbf{X} \in \mathbb{S}^n \mid \mathbf{X} \succeq 0\}. \tag{2.6}$$

The set of positive semidefinite matrices also forms a convex cone.

Generalised inequalities and proper cones

A cone $K \subseteq \mathbb{R}^n$ is called a proper cone if K is convex, closed, solid (non-empty interior) and pointed. A pointed cone means that if $\mathbf{x} \in K$ and $-\mathbf{x} \in K$ then

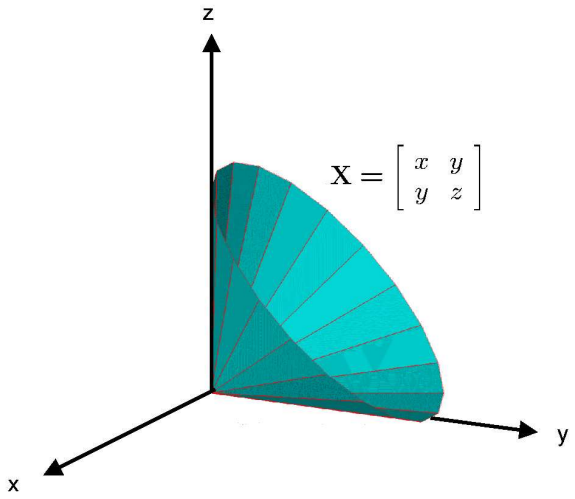


Figure 2.1: The boundary of the positive semidefinite cone in $\mathbb{S}^2$.

$\mathbf{x} = \mathbf{0}$. A proper cone allows the definition of a generalised inequality (a partial ordering on $\mathbb{R}^n$ which has many of the properties of the ordering on $\mathbb{R}$ with which everyone is familiar). The proper cone K has the partial ordering

$$\mathbf{x} \preceq_K \mathbf{y} \iff \mathbf{y} - \mathbf{x} \in K. \tag{2.7}$$

A strict partial ordering is defined as:

$$\mathbf{x} \prec_K \mathbf{y} \iff \mathbf{y} - \mathbf{x} \in \mathbf{int} K, \tag{2.8}$$

where **int** refers to the interior of a set.

Examples

The nonnegative orthant and componentwise inequality $\mathbb{R}^n_+$ is a proper cone and its associated generalised inequality is the componentwise inequality between vectors i.e. $\mathbf{x} \preceq \mathbf{y}$ means $x_i \leq y_i$ for all $i = 1, \ldots, n$.

The positive semidefinite cone and matrix inequality $\mathbb{S}^n_+$ is also a proper cone. Its generalised inequality $\mathbf{X} \preceq \mathbf{Y}$ means that $\mathbf{Y} - \mathbf{X} \in \mathbb{S}^n_+$. Strict generalised inequality means that the resulting matrix lies in the interior of the positive semidefinite cone, i.e. it is a positive definite matrix. Normally the subscript K is dropped since the cone in question is obvious from the context. This convention will be adopted here also.

Dual Cones

If K is a cone then the dual cone of K is the set

$$K^* = \left\{ \mathbf{y} \mid \mathbf{x}^T \mathbf{y} \geq 0 \quad \forall \mathbf{x} \in K \right\}. \tag{2.9}$$

This cone is also convex even when K is not. A cone that is its own dual is called self dual. Both the non-negative orthant and positive semidefinite cones are self dual. Dual cones also have generalised inequalities defined on them but in the case of self dual cones both generalised inequalities are the same. For the purposes of this chapter all cones discussed are self-dual.

Convex functions

The graph of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$\{(\mathbf{x}, f(\mathbf{x})) \mid \mathbf{x} \in \mathbf{dom} f\}, \tag{2.10}$$

where $\mathbf{dom} f$ is the domain of f. A function f is convex if

$$f(\theta \mathbf{x} + (1-\theta)\mathbf{y}) \leq \theta f(\mathbf{x}) + (1-\theta) f(\mathbf{y}), \tag{2.11}$$

where $\mathbf{x}, \mathbf{y} \in \mathbf{dom} f$ and $0 \leq \theta \leq 1$. Geometrically this corresponds to the requirement that any line segment which intersects the graph of $f(\mathbf{x})$ at $\mathbf{x}$ and $\mathbf{y}$ must lie above the graph of $f(\mathbf{x})$ within this interval (see figure 2.2). This must hold for all $\mathbf{x}$ and $\mathbf{y}$ within the domain of f. Furthermore a function f is convex if and only if it remains convex when restricted to any line segment that intersects its domain.

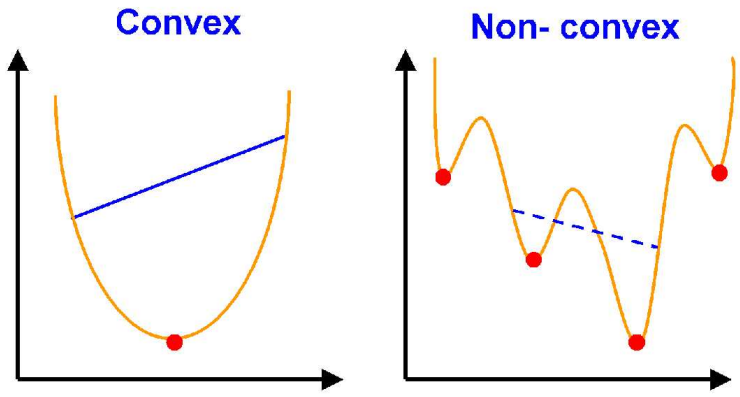


Figure 2.2: Geometrical interpretation of a convex function. Any line segment (blue) which intersects the graph of $f(\mathbf{x})$ (orange curve) must lie above the graph of f within this interval. A convex function has a single global minimum whereas a non-convex function may have many local minima.

Sublevel sets and epigraphs

Let α be a scalar. The α-sublevel set of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$C_\alpha = \{\mathbf{x} \in \mathbf{dom} f \mid f(\mathbf{x}) \leq \alpha\}. \tag{2.12}$$

The epigraph of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$\mathbf{epi} f = \{(\mathbf{x}, t) \mid \mathbf{x} \in \mathbf{dom} f, f(\mathbf{x}) \leq t\}. \tag{2.13}$$

Convexity and generalised inequalities

$f : \mathbb{R}^n \to \mathbb{R}^m$ is K-convex if for all $\mathbf{x}, \mathbf{y}$, and $0 \leq \theta \leq 1$,

$$f(\theta\mathbf{x} + (1-\theta)\mathbf{y}) \preceq_K \theta f(\mathbf{x}) + (1-\theta) f(\mathbf{y}), \tag{2.14}$$

where $K \subseteq \mathbb{R}^m$ is a proper cone.

2.2.2 Optimisation problems

The standard form of a general optimisation problem is defined as:

$$\begin{array}{lll} \text{minimise} & f_0(\mathbf{x}), & \\ \text{subject to} & f_i(\mathbf{x}) \leq 0, & i = 1, \ldots, m, \\ & h_i(\mathbf{x}) = 0, & i = 1, \ldots, p, \end{array} \tag{2.15}$$

where $f_0, \ldots, f_m : \mathbb{R}^n \to \mathbb{R}$. $f_0(\mathbf{x})$ is commonly referred to as the objective function, $f_i(\mathbf{x})$ are the inequality constraints and $h_i(\mathbf{x})$ are the equality constraints. The set of points for which the functions are defined is called the domain of the optimisation problem,

$$\mathcal{D} = \left(\bigcap_{i=1}^{m} \mathbf{dom}\ f_i\right) \bigcap \left(\bigcap_{i=1}^{p} \mathbf{dom}\ h_i\right). \tag{2.16}$$

while the set of points which satisfy the constraint equations is known as the feasible set.

The optimal value p^* of the problem is defined as

$$p^* = \inf\{f_0(\mathbf{x}) \mid f_i(\mathbf{x}) \leq 0,\ i = 1, \ldots, m,\ h_i(\mathbf{x}) = 0,\ i = 1, \ldots, p\}. \tag{2.17}$$

A feasible point $\mathbf{x}$ with $f_0(\mathbf{x}) \leq p^* + \epsilon,\ (\epsilon > 0)$ is called ϵ- suboptimal.

It should be pointed out that there are many different manifestations of optimisation problems which look different from the problem presented in equation 2.15.

These include maximisation problems, problems where the objective has been transformed to a constraint and problems involving slack variables (additional variables introduced to convert inequality constraints to equality constraints). These are all readily transformed into the standard form and for further information on how to achieve this the reader is referred to [194, 195]. A brief description of just one of these formats will be of value later on and that is the expression of a problem in its epigraph form. The epigraph form for equation 2.15 is

$$\begin{aligned} \text{minimise:} \quad & t, \\ \text{subject to:} \quad & f_0(\mathbf{x}) - t \leq 0, \\ & f_i(\mathbf{x}) \leq 0, \quad i = 1, \ldots, m, \\ & h_i(\mathbf{x}) = 0, \quad i = 1, \ldots, p, \end{aligned} \tag{2.18}$$

where $t \in \mathbb{R}$. The principal advantage for the epigraph transformation (2.18) is that the objective is now a linear function. The advantage of this is that any convex optimisation problem with a linear objective function must have its optimal solution on the boundary of the feasible set.

Convex optimisation problems

A convex optimisation problem is a specialisation of equation 2.15 in which three additional specifications are made:

1. The objective is convex.
2. The inequality constraints are convex.
3. The equality constraints are affine.

The consequence of points 2 and 3 is that the feasible set of the problem is convex. The convex optimisation problem can therefore be represented as

$$\begin{aligned} \text{minimise} \quad & f_0(\mathbf{x}), \\ \text{subject to} \quad & f_i(\mathbf{x}) \leq 0, \quad i = 1, \ldots, m, \\ & \mathbf{a}_i^T \mathbf{x} = b_i, \quad i = 1, \ldots, p. \end{aligned} \tag{2.19}$$

A fundamental property of convex optimisation problems is that any locally optimal point is also a globally optimal point. It is therefore possible to define an optimality criterion for any convex optimisation problem for which $f_0(\mathbf{x})$ is differentiable. Let X be the feasible set

$$X = \{\mathbf{x} \mid f_i(\mathbf{x}) \leq 0,\ i = 1, \ldots, m,\ h_i(\mathbf{x}) = 0,\ i = 1, \ldots, p\}\,. \tag{2.20}$$

$\mathbf{x}$ is optimal if and only if $\mathbf{x} \in X$ and

$$\nabla f_0(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) \geq 0 \quad \forall \mathbf{y} \in X. \tag{2.21}$$

The geometrical interpretation of this optimality criterion is that if $\nabla f_0(\mathbf{x}) \neq 0$ then $-\nabla f_0(\mathbf{x})$ defines a supporting hyperplane to the feasible set at $\mathbf{x}$ (see figure 2.3).

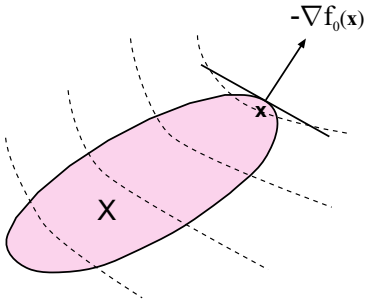


Figure 2.3: Geometric interpretation of the optimality condition (equation 2.21). The dashed lines are level curves. $\mathbf{x}$ is an optimal point since $-\nabla f_0(\mathbf{x})$ defines a supporting hyper-plane to X at $\mathbf{x}$.

Convex optimisation problems with generalised inequality constraints

The standard form of convex optimisation problem (equation 2.19) also applies where the inequality constraints are generalised inequalities defined for vector quantities:

$$\begin{aligned} \text{minimise} \quad & f_0(\mathbf{x}), \\ \text{subject to} \quad & f_i(\mathbf{x}) \preceq 0, \quad i = 1, \ldots, m, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{aligned} \tag{2.22}$$

where $f_0 : \mathbb{R}^n \to \mathbb{R}$, $K_i \subseteq \mathbb{R}^{k_i}$ are proper cones, and $f_i : \mathbb{R}^n \to \mathbb{R}^{k_i}$ are K_i-convex. Also $A \in \mathbb{R}^{p \times n}$ and $\mathbf{b} \in \mathbb{R}^p$. The three properties described above for regular convex optimisation problems also hold for the generalised case.

Example convex optimisation problems

Second-order cone programme A second order cone programme (SOCP) is defined as:

$$\begin{aligned} \text{minimise:} \quad & \mathbf{f}^T\mathbf{x}, \\ \text{subject to:} \quad & \|\mathbf{A}_i\mathbf{x} + b_i\|_2 \leq \mathbf{c}_i^T\mathbf{x} + d_i, \quad i = 1, \ldots, m, \\ & \mathbf{F}\mathbf{x} = \mathbf{g}. \end{aligned} \tag{2.23}$$

where $\mathbf{f}, \mathbf{x} \in \mathbb{R}^n$, and $\mathbf{x}$ is the optimisation variable, $\mathbf{A}_i \in \mathbb{R}^{n_i \times n}$, and $\mathbf{F} \in \mathbb{R}^{p \times n}$. A constraint of the form $\|\mathbf{A}_i\mathbf{x}+b_i\|_2 \leq \mathbf{c}_i^T\mathbf{x}+d_i$ constrains the function $(\mathbf{A}\mathbf{x}+\mathbf{b}, \mathbf{c}^T\mathbf{x}+d)$ to lie in the second order cone in $\mathbb{R}^{k+1}$.

Semidefinite programme The standard form SDP has linear equality constraints

$$\begin{aligned} \text{minimise:} \quad & \mathbf{Tr}(\mathbf{C}\mathbf{X}), \\ \text{subject to:} \quad & \mathbf{Tr}(\mathbf{A}_i\mathbf{X}) = b_i, \quad i = 1, \ldots, p, \\ & \mathbf{X} \succeq 0, \end{aligned} \tag{2.24}$$

where $\mathbf{C}, \mathbf{A}_1, \ldots, \mathbf{A}_p \in \mathbb{S}^n$ and $\mathbf{X} \in \mathbb{S}^n_+$. $\mathbf{Tr}(\mathbf{Y})$ represents the trace of the matrix $\mathbf{Y}$.

2.2.3 Duality theory for convex optimisation

Much of what separates convex optimisation from other optimisation categories is related to the difference in their respective duality theories. Lagrangian duality can be applied to any general optimisation problem, but it is within the category of convex optimisation that duality is most useful in that it enables the specification of exact optimality criteria when certain conditions are met.

The Lagrangian and the dual function

Given the optimisation problem in standard form (equation 2.15), the Lagrangian is the objective function augmented with a weighted sum of the constraint functions

$$L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = f_0(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i f_i(\mathbf{x}) + \sum_{i=1}^{p} \nu_i h_i, \tag{2.25}$$

where λ_i and ν_i are the Lagrange multipliers for the i^{th} inequality and equality constraints respectively.

The (Lagrange) dual function $g : \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$ is defined as the minimum of the Lagrangian over $\mathbf{x}$ as a function of $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\mathbf{x} \in \mathcal{D}} \left(f_0(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i f_i(\mathbf{x}) + \sum_{i=1}^{p} \nu_i h_i \right). \tag{2.26}$$

If the Lagrangian is unbounded from below in $\mathbf{x}$, the dual function has the value $-\infty$. Note that the dual function is always concave even when the problem in equation 2.15 is not convex.

The dual function provides a lower bound on the optimal value p^* of the standard

form optimisation problem (2.15) since for any $\boldsymbol{\lambda} \succeq 0$ and any $\boldsymbol{\nu}$,

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \leq p^*. \tag{2.27}$$

Hence for any feasible $\hat{\mathbf{x}}$

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\mathbf{x} \in \mathcal{D}} L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \leq L(\hat{\mathbf{x}}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \leq f_0(\hat{\mathbf{x}}). \tag{2.28}$$

The dual problem

Since the dual function gives a lower bound on p^* as a function of $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ it would be useful to find the best lower bound. This results in a second optimisation problem

$$\begin{aligned} \text{maximise:} \quad & g(\boldsymbol{\lambda}, \boldsymbol{\nu}), \\ \text{subject to:} \quad & \boldsymbol{\lambda} \succeq 0. \end{aligned} \tag{2.29}$$

This is the dual problem to the problem in equation 2.15 which is commonly termed the primal problem. Notice that the dual problem is always a convex optimisation problem since it is a maximisation where the objective is concave and the constraint is convex.

Weak duality

The property that $d^* \leq p^*$ termed weak duality is an important property that holds even when the primal problem is nonconvex. Weak duality holds even when the primal objective is $-\infty$ or the dual objective is $+\infty$ in which cases the dual is $-\infty$ and the primal is $+\infty$ respectively. This means that when the dual objective is $+\infty$, the primal problem is infeasible and when the primal objective is $-\infty$, the dual problem is infeasible. The difference between the optimal primal and optimal dual values, $p^* - d^*$ is called the duality gap. The duality gap is always non-negative. If it is zero, i.e. $p^* - d^* = 0$ the problem has strong duality. This means that the lower bound on the primal provided by the dual problem is tight.

Slater's condition

Strong duality is a very useful property as it gives a proof or certificate that the solution to the primal problem is optimal. In addition in cases where the dual problem is easier to solve than the primal because of its inherent convexity, solving the dual will provide the exact answer to the primal problem. Strong duality does not hold for optimisation problems in general. However in the case of convex optimisation problems where certain other criteria known as constraint qualifications are met, strong duality is obtained. A sufficient constraint qualification for strong duality in

convex optimisation problems of the type being considered in this work is known as Slater's condition. Slater's condition states that strong duality holds if there exists an $\mathbf{x}$ in the interior of the domain such that

$$\begin{aligned} f_i(\mathbf{x}) &< 0, \quad i = 1, \ldots, m, \\ \mathbf{A}\mathbf{x} &= \mathbf{b}. \end{aligned} \tag{2.30}$$

In words this means that if there exists a strictly feasible point for a convex primal problem then strong duality holds.

2.2.4 Optimality criteria

Certificates of optimality

The fact that the dual function evaluated at any dual feasible point $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ is always lower than the primal function evaluated at any primal feasible point $\hat{\mathbf{x}}$ means that the dual function evaluated at any dual feasible point, even if it is not optimal, gives a lower bound on p^*. In other words the dual function evaluated at the dual feasible point provides a proof or certificate that $p^* \geq g(\boldsymbol{\lambda}, \boldsymbol{\nu})$. If $\mathbf{x}$ is primal feasible and $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ is dual feasible then

$$f_0(\mathbf{x}) - p^* \leq f_0(\mathbf{x}) - g(\boldsymbol{\lambda}, \boldsymbol{\nu}), \tag{2.31}$$

which establishes that $\mathbf{x}$ and $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ are ϵ-suboptimal, where $\epsilon = f_0(\mathbf{x}) - g(\boldsymbol{\lambda}, \boldsymbol{\nu})$. Therefore if the duality gap is zero this provides a certificate that $\mathbf{x}$ is primal optimal and $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ is dual optimal.

Complementary slackness

If $\mathbf{x}^*$ is a primal optimal point and $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ is a dual optimal point then

$$\begin{aligned} f_0(\mathbf{x}^*) &= g(\boldsymbol{\lambda}, \boldsymbol{\nu}) && (2.32) \\ &= \inf_{\mathbf{x}} \left(f_0(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i f_i(\mathbf{x}) + \sum_{i=1}^{p} \nu_i h_i \right) && (2.33) \\ &\leq f_0(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i f_i(\mathbf{x}) + \sum_{i=1}^{p} \nu_i h_i && (2.34) \\ &\leq f_0(\mathbf{x}^*). && (2.35) \end{aligned}$$

From this it can be concluded that the last two inequalities are in fact equalities. Two important points can be made from this. The first is that $\mathbf{x}^*$ minimises the

Lagrangian function $L(\mathbf{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ at the optimal dual points. The second is that

$$\sum_{i=1}^{m} \lambda_i^* f_i(\mathbf{x}^*) = 0. \tag{2.36}$$

Since each term in the sum is greater or equal to zero this leads to the complementary slackness condition

$$\lambda_i^* f_i(\mathbf{x}^*) = 0, \quad i = 1, \ldots, m. \tag{2.37}$$

When strong duality holds, complementary slackness must hold for any primal-dual optimal pair $\mathbf{x}^*$, $(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$.

Karush-Kuhn-Tucker optimality conditions

If it is now assumed that the objective and constraint functions are differentiable then the same is true for the Lagrangian function $L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu})$. Since it has been shown above that any optimal point $\mathbf{x}^*$ is a minimiser for the Lagrangian over $\mathbf{x}$, the gradient of the Lagrangian must be zero at this point,

$$\nabla f_0(\mathbf{x}^*) + \sum_{i=1}^{m} \lambda_i^* \nabla f_i(\mathbf{x}^*) + \sum_{i=1}^{p} \nu_i^* \nabla h_i(\mathbf{x}^*) = 0. \tag{2.38}$$

This is the final optimality condition required to produce what are known as the Karush-Kuhn-Tucker (KKT) optimality conditions. These are a combination of the constraints from the original problem, the nonnegativity constraint on the λ_i, the complementary slackness condition and the Lagrangian gradient condition

$$\begin{aligned}
f_i(\mathbf{x}^*) &\leq 0 \quad i = 1, \ldots, m, \\
h_i(\mathbf{x}^*) &= 0 \quad i = 1, \ldots, p, \\
\lambda_i^* &\geq 0 \quad i = 1, \ldots, m, \\
\lambda_i^* f_i(\mathbf{x}^*) &= 0 \quad i = 1, \ldots, m, \\
\nabla f_0(\mathbf{x}^*) + \textstyle\sum_{i=1}^{m} \lambda_i^* \nabla f_i(\mathbf{x}^*) + \sum_{i=1}^{p} \nu_i^* \nabla h_i(\mathbf{x}^*) &= 0.
\end{aligned} \tag{2.39}$$

Any optimisation problem for which strong duality holds (and which has differentiable objective and constraint functions) must satisfy the KKT conditions at an optimal point. Hence the KKT conditions are necessary conditions for optimality. In the case of a convex optimisation problem the KKT conditions are both necessary and sufficient. Many algorithms for convex optimisation including interior point methods are in effect methods for solving the KKT equations.

2.2.5 Extending Lagrangian duality to generalised inequalities

The situation is quite similar to that for scalar functions but there are some subtle differences.

The Lagrangian and the dual function

For an optimisation problem with generalised inequalities

$$\begin{array}{rl} \text{minimise} & f_0(\mathbf{x}), \\ \text{subject to} & f_i(\mathbf{x}) \preceq_{K_i} 0, \quad i = 1, \ldots, m, \\ & h_i(\mathbf{x}) = 0, \quad i = 1, \ldots, p, \end{array} \tag{2.40}$$

the Lagrange multipliers are vectors i.e. $\boldsymbol{\lambda}_i \in \mathbb{R}^{k_i}$. The Lagrange dual is therefore

$$\begin{array}{l} L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = \\ f_0(\mathbf{x}) + \boldsymbol{\lambda}_1^T f_1(\mathbf{x}) + \cdots + \boldsymbol{\lambda}_m^T f_m(\mathbf{x}) + \nu_1 h_1(\mathbf{x}) + \cdots + \nu_p h_p(\mathbf{x}), \end{array} \tag{2.41}$$

where $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_m)$ and $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_p)$. The dual function is defined as

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\mathbf{x} \in \mathcal{D}} \left(f_0(\mathbf{x}) + \sum_{i=1}^{m} \boldsymbol{\lambda}_i^T f_i(\mathbf{x}) + \sum_{i=1}^{p} \nu_i h_i \right). \tag{2.42}$$

As before the dual function is concave, even if the primal is not, and the dual gives lower bounds on the primal objective. One distinction is that the scalar nonnegativity constraints on λ_i must be replaced by dual nonnegativity constraints

$$\boldsymbol{\lambda}_i^* \succeq_{K_i^*} 0 \quad i = 1, \ldots, m, \tag{2.43}$$

where K_i^* denotes the dual cone of K_i. Since all cones being dealt with are self-dual this is immaterial from a practical point of view.

Slater's condition

Slater's condition translates well from the scalar to the vector case. For the problem

$$\begin{array}{rl} \text{minimise} & f_0(\mathbf{x}), \\ \text{subject to} & f_i(\mathbf{x}) \preceq_{K_i} 0, \quad i = 1, \ldots, m, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{array} \tag{2.44}$$

the generalised Slater condition states that if there exists an $\mathbf{x} \in \mathbf{int}\mathcal{D}$ with $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $f_i(\mathbf{x}) \prec_{K_i} 0,\ i = 1, \ldots, m$, then strong duality holds.

Example

Lagrange dual of a semidefinite programme A semidefinite programme can also be expressed in inequality form

$$\begin{array}{rl} \text{minimise:} & \mathbf{c}^T\mathbf{x}, \\ \text{subject to:} & \mathbf{F}(\mathbf{x}) = x_1\mathbf{F}_1 + \cdots + x_n\mathbf{F}_n + \mathbf{G} \preceq 0, \end{array} \tag{2.45}$$

where $\mathbf{G}, \mathbf{F}_1, \ldots, \mathbf{F}_n \in \mathbb{S}^k$ and $\mathbf{c}$ and $\mathbf{x}$ are the vectorised forms of $\mathbf{C}$ and $\mathbf{X}$ respectively. The Lagrangian is

$$\begin{aligned} L(\mathbf{x}, \mathbf{Z}) &= \mathbf{c}^T\mathbf{x} + \mathbf{Tr}((x_1\mathbf{F}_1 + \cdots + x_n\mathbf{F}_n + \mathbf{G})\mathbf{Z}) \\ &= x_1(c_1 + \mathbf{Tr}(\mathbf{F}_1\mathbf{Z})) + \cdots + x_n(c_n + \mathbf{Tr}(\mathbf{F}_n\mathbf{Z})) + \mathbf{Tr}(\mathbf{G}\mathbf{Z}). \end{aligned} \tag{2.46}$$

The corresponding dual function is then

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \begin{cases} \mathbf{Tr}(\mathbf{G}\mathbf{Z}) & \mathbf{Tr}(\mathbf{F}_i\mathbf{Z}) + c_i = 0 \quad i = 1, \ldots, n \\ -\infty & \text{otherwise,} \end{cases} \tag{2.47}$$

where $\mathbf{Z} \in \mathbb{S}^k$ has been introduced as a dual variable. Hence the dual problem is given by

$$\begin{array}{rl} \text{maximise:} & \mathbf{Tr}(\mathbf{G}\mathbf{Z}), \\ \text{subject to:} & \mathbf{Tr}(\mathbf{F}_i\mathbf{Z}) + c_i = 0, \quad i = 1, \ldots, n, \\ & \mathbf{Z} \succeq 0. \end{array} \tag{2.48}$$

The dual SDP in standard form is given by

$$\begin{array}{rl} \text{maximise:} & \mathbf{b}^T\mathbf{y}, \\ \text{subject to:} & \sum_{i=1}^p \mathbf{y}_i\mathbf{A}_i \preceq \mathbf{C}. \end{array} \tag{2.49}$$

Optimality conditions

The optimality conditions are almost identical to those for the scalar case so only the results will be presented here.

Complementary slackness The complementary slackness condition (2.37) is replaced by the vector form

$$\boldsymbol{\lambda}_i^{*T} f_i(\mathbf{x}^*) = 0, \quad i = 1, \ldots, m. \tag{2.50}$$

KKT optimality conditions The Karush-Kuhn-Tucker conditions are

$$\begin{aligned} f_i(\mathbf{x}^*) &\preceq_{K_i} 0, \quad i = 1, \ldots, m, \\ h_i(\mathbf{x}^*) &= 0, \quad i = 1, \ldots, p, \\ \boldsymbol{\lambda}_i^* &\succeq_{K_i^*} 0, \quad i = 1, \ldots, m, \\ \boldsymbol{\lambda}_i^{*T} f_i(\mathbf{x}^*) &= 0, \quad i = 1, \ldots, m, \\ \nabla f_0(\mathbf{x}^*) + \textstyle\sum_{i=1}^m Df_i(\mathbf{x}^*)^T \boldsymbol{\lambda}_i^* + \sum_{i=1}^p \nu_i^* \nabla h_i(\mathbf{x}^*) &= 0. \end{aligned} \tag{2.51}$$

2.3 Global optimisation and global nonnegativity

Having laid the foundations of convex optimisation theory it is possible to describe the SOS decomposition method itself and how it can be applied to global optimisation. The problem of global optimisation is defined as

$$\text{minimise:} \quad f(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{R}^n. \tag{2.52}$$

and is an NP-hard problem in general. In energy landscape theory the function $f(\mathbf{x})$ could represent a PES. Any scalar λ is a lower bound on $f(\mathbf{x})$ if

$$f(\mathbf{x}) - \lambda \geq 0, \quad \forall \mathbf{x} \in \mathbb{R}^n. \tag{2.53}$$

It is possible to find the global minimum of (2.52) if a sequence of $\lambda^{(i)}$ can be found that converges to the greatest lower bound of (2.52) as $i \to \infty$. This problem can be written as an optimisation problem.

$$\begin{aligned} &\text{maximise:} \quad \lambda, \\ &\text{subject to:} \quad f(\mathbf{x}) - \lambda \geq 0, \quad \forall \mathbf{x} \in \mathbb{R}^n. \end{aligned} \tag{2.54}$$

This problem is no easier than (2.52) since a fundamental obstacle remains: how to decide when $f(\mathbf{x}) - \lambda \geq 0$. It involves deciding whether $f(\mathbf{x}) - \lambda$ is globally non-negative for a sequence of monotonically increasing λ. Deciding the global non-negativity of a general function is a well known decision problem and one for which no answer is known. However restricting the type of function to polynomial functions allows progress to be made in this respect.

2.3.1 Polynomials, nonnegative polynomials and sum of squares

At this point it is necessary to introduce some definitions:

monomial: A monomial in $x_1, \ldots, x_n$ is a product of the form

$$x_1^{\alpha_1} \cdot x_2^{\alpha_2} \cdots x_n^{\alpha_n}, \tag{2.55}$$

where all the exponents $\alpha_1, \ldots \alpha_n$ are nonnegative integers. The total degree of the monomial is the sum $\alpha_1 + \cdots + \alpha_n$.

polynomial: A polynomial f in $x_1, \ldots x_n$ is a linear combination of monomials.

$$f = \sum_{\alpha} a_\alpha x^\alpha, \tag{2.56}$$

where x^α is a monomial of degree α. The discussion here is restricted to real polynomials so $a_\alpha \in \mathbb{R}$. The total degree of f is the maximum $|\alpha|$ with nonzero coefficient a_α.

homogeneous form A form is a polynomial in which all the monomials have the same total degree. A homogeneous form of degree d satisfies $f(\lambda x_1, \ldots, \lambda x_n) = \lambda^d f(x_1, \ldots, x_n)$.

positive semidefinite form Let $H_d(\mathbb{K}^n)$ denote the set of homogeneous forms of degree d in n variables with coefficients in the field $\mathbb{K}$. Each form $f \in H_d(\mathbb{K}^n)$ can be associated with a vector of dimension $\binom{n+d-1}{n-1}$ of its coefficients. Let e be an even integer. A form $f \in H_e(\mathbb{K}^n)$ is positive semidefinite (PSD) if $f(x_1, \ldots, x_n) \geq 0$ for all $(x_1, \ldots, x_n) \in \mathbb{R}^n$. A form is positive definite if $f(x_1, \ldots, x_n) = 0$ implies $x_j = 0$ for $1 \leq j \leq n$.

sum of squares form: A form is a sum of squares (SOS) if it can be written as a sum of squares of polynomials

$$f^{SOS} = \sum_{i=1}^{r} (p_i(\mathbf{x}))^2 \geq 0, \tag{2.57}$$

where the $p_i(\mathbf{x})$ are polynomials.

The properties of being PSD and SOS are maintained under dehomogenisation of a form to a polynomial and are conversely preserved when a polynomial is homogenised to a form of even degree [197]. The set of polynomials with real coefficients (denoted $\mathbb{R}^d[x_1, \ldots, x_n]$) of degree d and dimension n is a vector space of dimension $\binom{n+d}{d}$.

2.3.2 The positive semidefinite and sum of squares cones

The set of PSD polynomials forms a convex cone ($\mathcal{P}$) within this vector space. The set of positive definite polynomials constitutes the interior of $\mathcal{P}$. The set of sum of squares polynomials also forms a convex cone (Σ) within this vector space. These

cones are an algebraic analogue of the cone introduced in section 2.2.1 and possess the following properties: Given a cone $\mathcal{R}$ which is a subset of $\mathbb{R}[x_1, \ldots, x_n]$,

1. $a, b \in \mathcal{R} \Rightarrow a + b \in \mathcal{R}$.
2. $a, b \in \mathcal{R} \Rightarrow a \cdot b \in \mathcal{R}$.
3. $a \in \mathbb{R}^d[x_1, \ldots, x_n] \Rightarrow a^2 \in \mathcal{R}$.

A fact that underpins the entire SOS decomposition approach is that the sum of squares cone is a subset of the cone of nonnegative polynomials. Work on this problem was ongoing before the famous mathematical congress which took place in Paris in 1900 when David Hilbert produced his famous list of 23 unsolved problems. Specifically his 17^{th} problem asked whether every positive semidefinite polynomial is decomposable as a sum of squares of rational functions. The proof that this is in fact true was produced in 1927 by Emil Artin [198]. A review of the mathematical developments surrounding Hilbert's 17^{th} problem and later the relationship between the SOS and PSD cones is given in [197].

It is this relationship between $\mathcal{P}$ and Σ that is of practical importance. Since not all PSD polynomials are SOS it is clear that $\mathcal{P} \supset \Sigma$. For further examination it is useful to consider this relationship as a function of both degree d and the number of variables n of the polynomial [199]. It is well known that for forms with $d = 2$ the two cones are identical. Two other cases where the cones are proved to be identical are a) bivariate forms ($n = 2$) and b) ternary quartics ($n = 3, d = 4$). Dehomogenisation from forms to polynomials leads to the following three categories a) quadratic polynomials, b) univariate polynomials and c) quartic polynomials in two variables. Many other cases exist but an interesting paper by Blekherman [200] established that if the degree d is fixed and $d > 4$ then the relative volume of $\mathcal{P}$ occupied by Σ goes to zero with increasing n. It has been shown by Parrilo [1, 191] that establishing that a polynomial is a member of Σ is much more straightforward than establishing whether it is a member of $\mathcal{P}$ and this chapter explains how use can be made of this to efficiently solve many global optimisation problems.

Examples of PSD polynomials that are not SOS

A number of examples of polynomials that are PSD but not SOS (denoted $\Delta_{n,d}$) have appeared in the literature over the past fifty years. The most famous is the Motzkin polynomial [201]. For $n \geq 3$ Motzkin proved that

$$(t_1^2 + \cdots + t_{n-1}^2 - nu^2)t_1^2 \cdots t_{n-1}^2 + u^{2n} \in \Delta_{n,2n}. \tag{2.58}$$

The most well known case is where $n = 3$

$$M(x, y, z) = (x^2 + y^2 - 3z^2)x^2y^2 + z^6. \tag{2.59}$$

Another example is the Robinson form [197]

$$\begin{aligned} R(x, y, z) = x^6 + y^6 + y^6 - (x^4y^2 + x^2y^4 + x^4z^2 + x^2z^4 + y^4z^2 + y^2z^4) \\ +3x^2y^2z^2. \end{aligned} \tag{2.60}$$

Others include the Choi-Lam form [197] and that of Lax-Lax and Schmüdgen [197].

2.3.3 Relaxation of the global nonnegativity condition

A particular type of global optimisation problem where objective and constraint functions $f(\mathbf{x})$ are polynomials is known as a polynomial optimisation problem. Equation 2.54 can be regarded as a polynomial optimisation problem if $f(\mathbf{x})$ is restricted to $f(\mathbf{x}) \in \mathbb{R}^d[x_1, \ldots, x_n]$. This problem is only relevant if f is of even degree, since otherwise it is unbounded from below. Notice that an optimisation problem in this form (epigraph form) can be viewed as a sequence of decision problems:

$$\text{is} \quad f(\mathbf{x}) - \lambda^{(i)} \geq 0, \quad \text{for all } \mathbf{x}, \tag{2.61}$$

for a monotonically increasing sequence $\lambda^{(i)}$. The largest λ in the sequence for which $\mathbf{x}$ is feasible is a greatest lower bound on $f(\mathbf{x})$. This is the same as asking whether $f(\mathbf{x}) - \lambda^{(i)}$ belongs to the cone of non-negative polynomials for a given $\lambda^{(i)}$. Although the problem has been recast in a different way, the decidability problem is no easier.

Shor's relaxation

A very effective relaxation to this problem was proposed by N. Z. Shor [193] in which the non-negativity condition is replaced by a sum of squares condition. This equates to testing whether $f(\mathbf{x}) - \lambda$ is in the sum of squares cone rather than the cone of nonnegative polynomials. Because the goal is to produce a lower bound for $f(\mathbf{x})$ it is possible to write the relaxed form of equation 2.54 as

$$\begin{aligned} &\text{maximise:} \quad \lambda \\ &\text{subject to:} \quad f(\mathbf{x}) - \lambda \text{ is } f^{SOS} \quad \forall \mathbf{x} \in \mathbb{R}^n. \end{aligned} \tag{2.62}$$

where f^{SOS} is a sum of squares function. The problem therefore becomes one of finding such a sum of squares decomposition.

2.3.4 Sum of squares decompositions

One way to test whether a polynomial of degree $2d$ in n variables has a sum of squares decomposition is to examine whether it can be written in Gram matrix form. Let $\mathbf{z}$ be a vector of all monomials in the n variables of degree less than or equal to d, then f has a SOS decomposition if and only if there exists a positive semidefinite matrix $\mathbf{Q}$ such that the following equality holds

$$f = \mathbf{z}^T \mathbf{Q} \mathbf{z}, \tag{2.63}$$

where $\mathbf{z}$ is defined as

$$\mathbf{z} = [1, x_1, x_2, \ldots, x_n, x_1 x_2, \ldots, x_n^d]. \tag{2.64}$$

The dimension of $\mathbf{z}$ is $\binom{n+d}{d}$ and therefore the dimension of $\mathbf{Q}$ is $\binom{n+d}{d} \times \binom{n+d}{d}$. The formalisation of this method as the "Gram matrix" method was produced in [202]. It should be noted that the matrix $\mathbf{Q}$ is not unique since the entries in $\mathbf{z}$ are not algebraically independent. The problem of finding a suitable $\mathbf{Q}$ matrix was addressed first by Powers and Wörmann [203]. However a major development was the independent realisation by Parrilo [191] and Lassere [192] that finding such a sum of squares decomposition is equivalent to solving a semidefinite programme in standard primal form (equation 2.24). The following example (from [199]) illustrates Gram matrix decomposition and its relationship to a semidefinite programme.

Example Given the quartic form

$$f = 2x^4 + 2x^3 y - x^2 y^2 + 5y^4, \tag{2.65}$$

the monomial vector can then be defined as $\mathbf{z} = [x^2, y^2, xy]^T$. Note that if a polynomial $f(\mathbf{x})$ is homogeneous of degree $2d$ then its monomial basis in the Gram matrix decomposition can be restricted to the monomials of degree d. The Gram matrix decomposition can be represented as

$$\begin{bmatrix} x^2 \\ y^2 \\ xy \end{bmatrix}^T \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \begin{bmatrix} x^2 \\ y^2 \\ xy \end{bmatrix}. \tag{2.66}$$

which can be expanded as

$$q_{11} x^4 + q_{22} y^4 + (q_{33} + 2q_{12}) x^2 y^2 + 2q_{13} x^3 y + 2q_{23} x y^3. \tag{2.67}$$

This expansion shows that there is a linear relationship between the coefficients in (2.65) and the entries in $\mathbf{Q}$:

$$q_{11} = 2,\ q_{22} = 5,\ q_{33} + 2q_{12} = -1,\ 2q_{13} = 2,\ 2q_{23} = 0. \tag{2.68}$$

A possible solution is

$$\mathbf{Q} = \begin{bmatrix} 2 & -3 & 1 \\ -3 & 5 & 0 \\ 1 & 0 & 5 \end{bmatrix}. \tag{2.69}$$

The number of squares in the SOS decomposition is given by the rank of $\boldsymbol{Q}$ which is 2 and the decomposition itself is obtained from a Cholesky factorisation of $\boldsymbol{Q}$:

$$\mathbf{Q} = \mathbf{L}^T\mathbf{L}, \quad \mathbf{L} = \frac{1}{\sqrt{2}}\begin{bmatrix} 2 & -3 & 1 \\ 0 & 1 & 3 \end{bmatrix}. \tag{2.70}$$

Therefore the SOS decomposition is

$$f^{SOS} = \frac{1}{2}(2x^2 - 3y^2 + xy)^2 + \frac{1}{2}(y^2 + 3xy)^2. \tag{2.71}$$

If the goal is to find a lower bound λ on this function the problem becomes

$$\begin{aligned} &\text{maximise:} && \lambda \\ &\text{subject to:} && 2x^4 + 2x^3y - x^2y^2 + 5y^4 - \lambda \text{ is a SOS.} \end{aligned} \tag{2.72}$$

Now let $\mathbf{A}$ be a positive semidefinite real matrix such that $f(x) = \mathbf{z}^T\mathbf{A}\mathbf{z}$ and let $\mathbf{E}_{11}$ denote a matrix of identical dimensions to $\mathbf{A}$ with zeros everywhere except the top-left entry which contains a 1. Now define $\mathbf{Q}$ as $\mathbf{Q} = \mathbf{A} - \lambda\mathbf{E}_{11}$. The problem is now in the form of a semidefinite programme. A detailed explanation of how such an optimisation is carried out in practise is given in appendix A.

□

The equivalence between the problem (equation 2.72) and the primal SDP (equation 2.24) can be seen as follows: $\mathbf{C} := -\mathbf{E}_{11},\ \mathbf{X} := \mathbf{Q}$ and $\mathbf{Tr}(\mathbf{A}_i\mathbf{X}) = b_i,\ i = 1, \ldots, p$ are defined by equation 2.68. The size of the resulting SDP is polynomial in either n or d if the other is fixed and the solution to an SDP can be obtained in polynomial time. Therefore testing whether $f \in \Sigma$ is answerable in polynomial time.

Sparse monomial vectors and the Newton polytope

One of the obstacles to sum of squares optimisation becoming a practically useful global optimisation method is how the size of the monomial basis vector $\mathbf{z}$ scales

with both increasing polynomial degree $2d$ and number of variables n. The binomial scaling $\binom{n+d}{d}$ described earlier is in fact a worst case scenario. Often the SOS decomposition can be achieved with a much sparser monomial vector. Given a function $f = \sum_{\alpha \in M} c_\alpha x^\alpha$, the set of monomials $M \subset \mathbb{N}^n$ is called the frame of f (**frame**(f)). Further the Newton polytope of f is defined as the convex hull of **frame**(f)

$$\mathbf{new}(f) = \mathbf{conv}(\mathbf{frame}(f)). \tag{2.73}$$

The Newton polytope can be used to choose an appropriate minimal monomial basis vector $\mathbf{z}$ using the following result [204]: If f has a SOS decomposition $f = \sum_{i=1}^{t} g_i^2$, then

$$\mathbf{new}(g) \subseteq \frac{1}{2}\mathbf{new}(f). \tag{2.74}$$

The squares in a SOS decomposition can only contain the monomials

$$\mathbf{new}(\frac{1}{2}f) \cap \mathbb{N}^n. \tag{2.75}$$

Example (from [204]) Consider the problem of finding a SOS decomposition for

$$f = 4x^4y^6 + x^2 - xy^2 + y^2. \tag{2.76}$$

The monomial basis can be restricted to

$$\mathbf{new}(f) \cap \mathbb{N}^n = \{x^2y^3, xy^2, xy, x, y\}. \tag{2.77}$$

The full monomial basis of degree 5 has 21 monomials. This is illustrated in figure 2.4.

□

Other techniques to reduce the monomial basis vector

Another instance where the monomial vector scales as less than $\binom{n+d}{d}$ is the case where the function is seperable into a sum of functions each of which is independent of the variables of the other functions. In this case the monomial vector need only contain the monomials for each seperate SOS decomposition, without including any monomials with cross terms between the variables of the seperable functions.

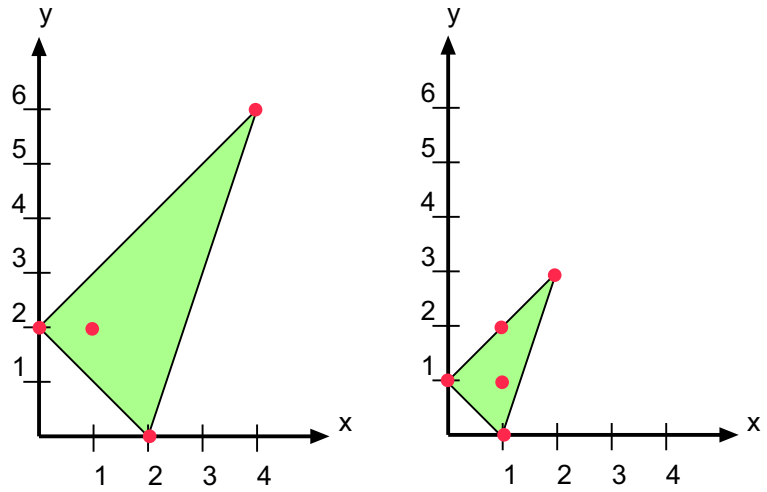


Figure 2.4: Finding a SOS representation for the polynomial function $f = 4x^4y^6 + x^2 - xy^2 + y^2$ shown on the left as red dots. The Newton polytope of f is shown in green. The Newton polytope for the sum of squares decomposition is shown on the right and $\mathbf{new}(\frac{1}{2}f)\cap\mathbb{N}^n = \{x^2y^3, xy^2, xy, x, y\}$ (shown as red dots). (Adapted from [204]).

2.3.5 The dual problem in sum of squares

In section 2.2.1 the concept of a dual cone was introduced and two relevant examples, the second order cone and the cone of positive semidefinite matrices were described. The definition of a dual cone K^* to a cone K given in equation 2.9 and reproduced here is

$$K^* = \{\mathbf{y} \mid \langle \mathbf{x}, \mathbf{y} \rangle \geq 0, \quad \forall \mathbf{x} \in K\}, \tag{2.78}$$

where the operation $\mathbf{x}^T\mathbf{y}$ has been rewritten in inner product form. This can be used to produce a dual cone for the cone of sum of squares. Returning to the representation of a sum of squares function in Gram matrix form, a function $f(\mathbf{x})$ can be rewritten as [199]

$$f(\mathbf{x}) = \mathbf{z}^T\mathbf{Q}\mathbf{z} = \langle \mathbf{z}\mathbf{z}^T, \mathbf{Q} \rangle \geq 0. \tag{2.79}$$

The definition of the matrix $\mathbf{z}\mathbf{z}^T$ is identical to the definition of a dual cone above. The matrix $\mathbf{z}\mathbf{z}^T$ has rank one and if $f(\mathbf{x})$ is SOS then $\mathbf{z}\mathbf{z}^T$ must be positive semidefinite. The entries of $\mathbf{z}\mathbf{z}^T$ are

$$\mathbf{z}\mathbf{z}^T = \begin{bmatrix} z_1^2 & z_1z_2 & \dots & z_1z_n \\ z_1z_2 & z_2^2 & \dots & z_2z_n \\ \vdots & \vdots & \ddots & \vdots \\ z_1z_n & z_2z_n & \dots & z_n^2 \end{bmatrix}, \tag{2.80}$$

where the z_i are the entries in the monomial vector $\mathbf{z}$. Since the z_i are not independent, there are multiplicative relationships between the entries of $\mathbf{z}\mathbf{z}^T$ which can be represented in the form $z_i z_j = z_k^2$. Now if $\mathbf{z}\mathbf{z}^T$ is replaced by a matrix

$$\mathbf{W} = \begin{bmatrix} w_{11} & w_{12} & \dots & w_{1n} \\ w_{21} & w_{22} & \dots & w_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ w_{1n} & w_{2n} & \dots & w_{nn} \end{bmatrix}, \tag{2.81}$$

then the multiplicative relationships are linearised to conditions of the form $w_{ij} = w_{kl}$. An example will clarify this.

Example Continued from earlier example [204].
In this case since $\mathbf{z} = [x^2, y^2, xy]$, $\mathbf{z}\mathbf{z}^T$ is defined as follows

$$\mathbf{z}\mathbf{z}^T = \begin{bmatrix} x^4 & x^2y^2 & x^3y \\ x^2y^2 & y^4 & xy^3 \\ x^3y & xy^3 & x^2y^2 \end{bmatrix}, \quad \text{and} \quad \mathbf{W} = \begin{bmatrix} w_{11} & w_{12} & w_{13} \\ w_{12} & w_{22} & w_{23} \\ w_{13} & w_{23} & w_{33} \end{bmatrix}, \tag{2.82}$$

and so the constraint $z_1 z_2 = z_3^2$, becomes $w_{12} = w_{33}$. Using this to eliminate w_{12}, the inner product $\langle \mathbf{W}, \mathbf{Q} \rangle$ is

$$\begin{aligned} \langle \mathbf{W}, \mathbf{Q} \rangle &= \begin{bmatrix} w_{11} & w_{33} & w_{13} \\ w_{33} & w_{22} & w_{23} \\ w_{13} & w_{23} & w_{33} \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \\ &= w_{11}q_{11} + w_{33}(q_{33} + 2q_{12}) + 2q_{13}w_{13} + w_{22}q_{22} + 2w_{23}q_{23} \\ &= 2w_{11} + 2w_{13} - w_{33} + 5w_{22}. \end{aligned} \tag{2.83}$$

□

Returning to the formulation of the SOS optimisation problem as an SDP, it has been shown [1] that there is a dual SDP which optimises over the dual cone to the sum of squares cone. By relaxing the rank one constraint on $\mathbf{W}$ an SDP in standard dual form is obtained where $\mathbf{W} := \mathbf{C} - \sum_{i=1}^{m} y_i \mathbf{A}_i \succeq 0$. If the relaxation is exact the solution matrix $\mathbf{W}^*$ will also be of rank one and the components of the optimal point $\mathbf{x}^*$ can be obtained from the powers of the appropriate entries of $\mathbf{W}^*$.

It is therefore very attractive to use the primal-dual interior point methods described in appendix B to solve SOS optimisation SDPs since both the primal and dual problems are of intrinsic value beyond certifying optimality. In summary, the primal and dual problems each give important information about an optimal point in a SOS optimisation problem:

1. The variable λ in the primal sum of squares SDP gives a lower bound on the objective function.

2. The entries of the dual matrix $\mathbf{W} := \mathbf{C} - \sum_{i=1}^{m} y_i \mathbf{A}_i \succeq 0$ gives the appropriate powers of the components of the minimiser vector $\mathbf{x}^*$.

2.3.6 Sum of squares liftings

An important concept relating to SOS optimisation is that the objective function need not be a sum of squares in the space of interest ($\mathbb{R}^n$). What is required is that it be a projection of a higher dimensional sum of squares in $\mathbb{R}^{(n+o)}$ where $\mathbb{R}^o$ is the orthogonal complement to $\mathbb{R}^n$. Therefore lifting the problem into a higher dimensional space makes it possible to find a lower bound which is a sum of squares and is exact even when the objective function does not have a SOS decomposition. A way of systematically constructing appropriate sum of squares liftings is required and the tools for this are best understood in the context and language of algebraic geometry.

2.4 Computational algebraic geometry

Algebraic geometry is concerned with the relationship between geometric objects defined by the solutions (intersections) of sets of polynomial functions. The basic geometric object (the solution set) is known as an affine variety and the set of polynomials defining a variety is called its ideal. In fact there is a dual relationship between the geometric set (variety) and the polynomial objects (ideal) that define it. Algebraic geometry is most easily understood in the context of the complex field $\mathbb{C}$ because it is a closed field. However real algebraic geometry (algebraic geometry for the field $\mathbb{R}^n$), although more involved (since the real field is not closed), is more relevant to the optimisation problem described in (2.54). The approach used here is to introduce the concepts and principal results for the complex field where the results are most easily understood and then to extend them to the special case of the real field.

Real algebraic geometry is useful in addressing the optimisation problem posed in equation 2.54 because it provides proofs of global nonnegativity in a much more powerful, yet subtle way than a direct relaxation from a global nonnegativity to a SOS requirement. The approach works by systematically constructing proofs that a given set defined by polynomial equality, inequations and inequality constraints is empty. So in order to prove that equation 2.61 holds, it is equally possible to prove that

$$f(\mathbf{x}) - \lambda^{(i)} < 0, \quad \text{is empty for all } \mathbf{x}. \tag{2.84}$$

Secondly it provides a systematic way of producing higher dimensional sum of squares necessary for the lifting process leading to more elaborate SOS relaxations. The reason why this is advantageous comes from duality. The dual of a primal problem is not exclusively defined by the primal objective and primal feasible sets alone. It is also dependent on the equality and inequality functions used to define the primal problem. In fact in the case where duality is not tight, adding additional redundant constraints can make the dual bound on the primal tighter. Such additional constraints are known as valid inequalities. Algebraic geometry therefore provides tools to systematically produce better bounds on a given polynomial optimisation problem by making use of the properties of the sum of squares cone.

2.4.1 The complex case

Varieties and ideals

The basic concepts are described for the complex field and extensions to real algebraic geometry are presented without delving into detail as it is possible to appreciate the key results based on the results for the complex field.

Affine varieties Let $f_1, \ldots, f_s$ be polynomials in the field k.

$$\mathbf{V}(f_1, \ldots, f_s) = \{(a_1, \ldots, a_n) \in k^n \mid f_i(a_1, \ldots, a_n) = 0,\ \forall i = 1, \ldots, s\} \tag{2.85}$$

is the affine variety defined by $f_1, \ldots, f_s$. In words $\mathbf{V}(f_1, \ldots, f_s)$ is the set of all solutions to the system of equations $f_i(a_1, \ldots, a_n) = 0$.

Ideals A subset $I \subset k[x_1, \ldots, x_n]$ is an ideal if each of the following three statements are true.

1. $0 \in I$.
2. $f, g \in I$ implies $f + g \in I$.
3. $f \in I$ and $h \in k[x_1, \ldots, x_n]$, implies $hf \in I$.

If $f_1, \ldots, f_s \in k[x_1, \ldots, x_n]$ then $\langle f_1, \ldots, f_s \rangle$ is an ideal of $k[x_1, \ldots, x_n]$ and is the ideal generated by $f_1, \ldots, f_s$. There is a famous theorem (The Hilbert Basis Theorem) which states that every ideal $I \subset k[x_1, \ldots, x_n]$ has a finite generating set of polynomials, i.e. $I = \langle g_1, \ldots, g_t \rangle$ for some $g_1, \ldots, g_t \in I$. There are a number of different bases which can be used to construct ideals including a Groebner basis which can be calculated using for example Buchberger's algorithm but this not discussed here [205] . Ideal-variety relationships can be studied using elimination theory and resultants. There is a precedent for geometry optimisation in theoretical

chemistry using such algebraic approaches in the work by Emiris [206] in which the geometry optimisation of cyclic molecules was studied.

Monoid The set $\mathbf{monoid}\{f_1, \ldots, f_s\} \subset \mathbb{R}[x_1, \ldots, x_n]$ is the set of all finite products of the polynomials f_i together with 1.

2.4.2 Hilbert's Nullstellensatz

Another famous theorem of Hilbert, the Nullstellensatz behaves like an algebra-geometry dictionary [205] translating (mapping) any statement about a variety into one about its corresponding ideal and vice versa. This is a one-to-many mapping since a given variety may be defined by more than one ideal.

Theorem 2.4.1 *Nullstellensatz.*
For an algebraically closed field, let $I \subset k[x_1, \ldots, x_n]$ *be an ideal satisfying* $\mathbf{V}(I) = 0$. *Then* $I = k[x_1, \ldots, x_n]$.

□

This theorem provides a strong duality result for polynomials in the complex field since it can be interpreted in the following way

$$1 \in \mathbf{ideal}(f_1, \ldots, f_s) \Leftrightarrow V(I) = \{\emptyset\}. \tag{2.86}$$

The reason for this is as follows. If the constant polynomial 1 is in I then for every polynomial $f \in k[x_1, \ldots, x_n]$ it can be said that $f \in I$ by point 3 in the definition of an ideal. Another way of stating the Nullstellensatz is

If there exist polynomials $h_i \in \mathbb{C}[x_1, \ldots, x_n]$ such that

$$f_1(x)h_1(x) + \cdots + f_s(x)h_s(x) = 1 \Leftrightarrow V(I) = \{\emptyset\}. \tag{2.87}$$

The difficult part is finding or proving the existence of such polynomials h_i. There are two possible ways to find them. Linear algebraic methods can be used to solve a system of linear equations. Alternatively Gröbner basis methods may be employed [205]. This does not hold for real polynomials since the Nullstellensatz conditions are necessary and sufficient only when the algebraic field is closed. Only the sufficiency condition holds in the real case.

Example (Taken from [204]) Consider the polynomials

$$f_1(x) = x_1^2, \quad f_2(x) = 1 - x_1x_2. \tag{2.88}$$

There is no $x \in \mathbb{C}^2$ which satisfies both $f_1(x) = 0$ and $f_2(x) = 0$, i.e.

$$\mathcal{V}\{f_1, f_2\} = \emptyset. \tag{2.89}$$

The Nullstellensatz says that there exists polynomials h_1, h_2 such that

$$1 = h_1(x)f_1(x) + h_2(x)f_2(x). \tag{2.90}$$

Suitable candidates include

$$h_1(x) = x^2, \quad h_2(x) = 1 + x_1x_2. \tag{2.91}$$

This example can also be viewed as a feasibility problem: Given two equality polynomial constraints, h_i, h_j, determine whether a feasible point exists. In section 2.2.4 the concept of certificates of optimality was introduced, it is also possible to have certificates of infeasibility which certify that no solution exists. The Nullstellensatz provides such a certificate of infeasibility in the complex field. By finding any h_i, h_j, (for example equations 2.91) such that equation 2.90 is true an infeasibility certificate for the feasibility problem is produced.

□

The Nullstellensatz is therefore a way of producing certificates of infeasiblity for feasibility problems defined using equality constraints in the complex field. But for global optimisation problems there are two additional requirements: i) it is necessary to be able to handle inequalities and ii) to work in the real field.

2.4.3 Real algebraic geometry

A real variety is defined as

$$\mathcal{V}_{\mathbb{R}}\{h_1, \ldots, h_m\} = \{\mathbf{x} \in \mathbb{R}^n \mid h_i(\mathbf{x}) = 0 \; \forall i = i, \ldots, m\}, \tag{2.92}$$

where $h_1, \ldots, h_m \in \mathbb{R}[x_1, \ldots, x_n]$.

The Real Nullstellensatz

There is an equivalent theorem to Hilbert's Nullstellensatz for the real case which states the following [207]

Theorem 2.4.2 *The Real Nullstellensatz.*
Let $I \subset \mathbb{R}[x_1, \ldots, x_n]$ be the ideal defined by $\langle h_1, \ldots, h_m \rangle$. Then

$$-1 \in \Sigma + \langle h_1, \ldots, h_m \rangle \iff \mathcal{V}_{\mathbb{R}}(I) = \{\emptyset\}. \tag{2.93}$$

□

This can be interpreted as saying that if there exists $t_1, \ldots, t_m \in \mathbb{R}[x_1, \ldots, x_n]$ and $s \in \Sigma$ such that

$$-1 = s + t_1 h_1 + \cdots + t_m h_m, \tag{2.94}$$

then there exists no $\mathbf{x} \in \mathbb{R}^n$ such that $h_i(\mathbf{x}) = 0$ for all $i = 1, \ldots, m$. Therefore the real Nullstellensatz provides a certificate of emptiness of the variety of the ideal constructed from $h_1, \ldots, h_m$.

Example Taken from [204]. Let $h(x) = x^2 + 1$. It is immediatly clear that $\mathcal{V}_{\mathbb{R}}\{h\} = 0$. A suitable real Nullstellensatz construction which proves the emptiness of $\mathcal{V}_{\mathbb{R}}$ would be

$$-1 = s + th \tag{2.95}$$

where $s(x) = x^2$ and $t(x) = -1$.

□

The Positivstellensatz

In the case where polynomial inequalities are used to define the feasible set, it is necessary to move from speaking about varieties (which are defined by polynomial equalities) to semialgebraic sets.

Basic semialgebraic sets: A basic semialgebraic set defined by the polynomials $f_1, \ldots, f_s$ is the set

$$S = \{\mathbf{x} \in \mathbb{R}^n \mid f_i(\mathbf{x}) \geq 0, \ \forall i = 1, \ldots, s\}. \tag{2.96}$$

Examples include the nonnegative orthant in $\mathbb{R}^n$, the cone of PSD matrices $\mathbb{S}^p_+$ and their intersections. Given any basic semialgebraic sets, semialgebraic sets may be generated by performing union, intersection and complement operations on them. The feasible set of a polynomial optimisation problem is a semialgebraic set since it is defined by polynomial equalities and inequalities. Every real variety is also a semialgebraic set. Real algebraic geometry provides a way to produce certificates of emptiness (refutations) for semialgebraic sets in general through a key theorem in real algebraic geometry known as Stengle's Positivstellensatz [208]:

Theorem 2.4.3 *Stengle's Positivstellensatz*
Given a semialgebraic set S defined by inequalities $f_1, \ldots, f_s$, and equalities $h_1, \ldots, h_m$,

$$-1 \in \mathbf{cone}\{f_1, \ldots, f_s\} + \mathbf{ideal}\{h_1, \ldots, h_m\} \Longleftrightarrow S = \{\emptyset\}, \tag{2.97}$$

where $\mathbf{cone}\{f_1, \ldots, f_s\}$ *denotes the cone generated by* $\{f_1, \ldots, f_s\}$ *and* $\mathbf{ideal}\{h_1, \ldots, h_m\}$ *denotes the ideal generated by* $\{h_1, \ldots, h_m\}$.

□

An alternative and more concrete statement of the theorem is the following [207]:

Theorem 2.4.4 *Stengle's Positivstellensatz*
Given the sets of polynomials $\{f_1, \ldots, f_s\}$,$\{g_1, \ldots, g_t\}$ and $\{h_1, \ldots h_m\}$. Let P be the cone generated by $\{f_1, \ldots, f_s\}$, M the monoid generated by $\{g_1, \ldots, g_t\}$ and I the ideal generated by $\{h_1, \ldots h_m\}$. Then the following statements are equivalent:

1. *The set*

$$\left\{ x \in \mathbb{R}^n \;\middle|\; \begin{array}{ll} f_j(x) \geq 0, & j = 1, \ldots, s \\ g_k(x) \neq 0, & k = 1, \ldots, t \\ h_l(x) = 0, & l = 1, \ldots, m \end{array} \right\} \text{ is empty.} \tag{2.98}$$

2. *There exists $f \in P$,$g \in M$, and $h \in I$, such that $f + g^2 + h = 0$.*

□

The Positivstellensatz is a generalisation of convex optimisation duality which was discusssed in section 2.2.3 and Hilbert's Nullstellensatz. In words it states that for a system of polynomial equalities and inequalities either one of two things must exist:

- A solution in $\mathbb{R}^n$.
- A polynomial identity (refutation) that bears witness to the fact that no solution exists.

The Positivstellensatz guarantees the existence of these polynomials but doesn't provide any information on how to obtain them. In fact the best complexity bounds on finding the necessary polynomials are triply exponential [199]. However if the degree of the refutation is bounded from above then the search for a suitable refutation may be carried out using a semidefinite programme. Therefore the search for a refutation within a given bound on the degree is achievable in polynomial time. If the degree of the refutation is high enough, the SDP will be feasible and will produce a certificate of emptiness via its solution [199].

Practical interpretation of the positivstellensatz

To construct a Positivstellensatz refutation for a polynomial optimisation problem

$$\begin{aligned} \text{minimise:} \quad & f_0(\mathbf{x}), \\ \text{subject to:} \quad & \mathbf{x} \in S, \end{aligned} \tag{2.99}$$

where S is defined by $f_i(\mathbf{x}) \geq 0,\ \forall i = 1, \ldots, s$, and $h_i(\mathbf{x}) = 0,\ \forall i = 1, \ldots, m$. λ is a lower bound on $f_0(\mathbf{x})$ over S if and only if the semialgebaic set

$$\{\mathbf{x} \in S, f_0(\mathbf{x}) - \lambda < 0\} = \{\emptyset\}. \tag{2.100}$$

This can be interpreted as finding the coefficients of t_i, s_i and r_{ij} such that

$$f_0 - \lambda - \Sigma_i h_i t_i - \Sigma_i s_i f_i - \Sigma_{i \neq j} r_{ij} f_i f_j - \cdots \text{ is SOS}, \tag{2.101}$$

where t_i are polynomials and s_i and r_{ij} are SOS polynomials and $\cdots$ refers to $\mathbf{monoid}\{f_1, \ldots, f_s\}$.

The positivstellensatz can be regarded as a general tool for incorporating both linear/nonlinear equality and inequality contraints into a SOS optimisation problem. In this respect it can be viewed as a generalisation of the Lagrangian with the SOS polynomials taking the place of nonnegative scalar Lagrange multipliers.

2.4.4 Alternative sum of squares representations

There are other ways of generating SOS representations which although not as general as the positivstellensatz are useful in many cases and sometimes less computationally expensive to compute. These can all be regarded as specialisations of the positivstellensatz [199].

Schmüdgen's Positivstellensatz

This is a stronger form of Stengle's Positivstellensatz. Given a semialgebraic set defined as in equation 2.96, every polynomial in $\mathbf{cone}\{f_1, \ldots, f_m\}$ is nonnegative on S. If S is a compact set Schmüdgen showed that

$$f \in \mathbf{cone}\{f_1, \ldots, f_m\} \quad \Rightarrow \quad f(\mathbf{x}) > 0 \quad \text{for all } \mathbf{x} \in S. \tag{2.102}$$

(In Euclidean space a set is compact if it is closed and bounded). In other words finding a representation of f in the following form

$$f = s_0 + \sum_i s_i f_i + \sum_{i,j} r_{ij} f_i f_j + \cdots, \tag{2.103}$$

where s_i and r_{ij} are SOS polynomials and $\cdots$ implies $\mathbf{monoid}\{f_1, \ldots, f_m\}$, provides a certificate of non-negativity for f.

Putinar's representation

If there is a single polynomial f_k in the constraint set such that

$$\{\mathbf{x} \in \mathbb{R}^n \mid f_k(\mathbf{x}) \geq 0\} \quad \text{is compact,} \tag{2.104}$$

then finding

$$f = s_0 + \sum_i s_i f_i, \tag{2.105}$$

for some SOS polynomials s_i, provides a certificate of nonnegativity for f over the compact set S.

2.5 Putting it all together

At this point the most important topics have been dealt with and it is now possible to describe the sum of squares optimisation algorithm.

Algorithm 2.5.1 *The sum of squares optimisation algorithm*

1. *Write the potential energy minimisation problem as a standard polynomial optimisation problem.*
2. *Transform the optimisation problem into epigraph form.*
3. *Relax the optimisation problem to a SOS optimisation problem using if necessary an appropriate nonnegative representation such as Stengle's Positivstellensatz.*
4. *Express the optimisation problem in Gram matrix form.*
5. *Map the SOS optimistion problem into a primal-dual semidefinite programme.*
6. *Solve the semidefinite programme to obtain a lower bound on the global minimum from the primal SDP.*
7. *If the rank of the dual matrix is 1, then obtain the minimiser coordinates from the dual solution matrix (see figure 3.8 and section 3.4 for details).*
8. *If the rank of the dual matrix $>$ 1, check for symmetry or repeat from step 3 using a higher dimensional lifting (see figure 3.8 and section 3.4 for details).*

The steps of the algorithm are described in flowchart form in figure 2.5.

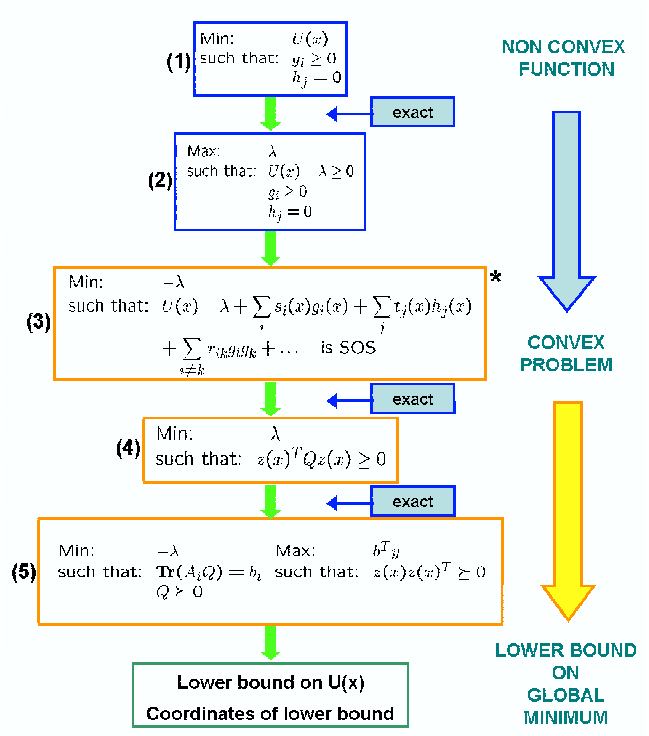


Figure 2.5: Schematic flowchart of a SOS optimisation. The original global energy minimisation problem is reformulated in a series of steps: (1) a polynomial optimisation including any constraints is transformed to (2) a maximisation of a real number λ. A sufficient but not necessary condition for this to be true is that the function is a SOS, so the problem is rewritten as (3). The constraints can be incorporated as well using the Positivstellensatz. The additional terms provide a series of hierarchical approximations to the true solution. (4) A function being a SOS is equivalent to finding a positive semidefinite matrix $\mathbf{Q}$ where $\mathbf{z}$ is a vector of monomials of U. Note that linear constraints on the entries of $\mathbf{Q}$ relates the q_{ij} to the coefficients of U. Finding $\mathbf{Q}$ is equivalent to the semidefinite programme given in (5), with the primal (left) and dual (right) problems. Each step is an exact reformulation of the previous (apart from (3)) and for a sufficiently high degree is guaranteed to give an optimal bound to the original problem. When expressed in the appropriate monomial basis the problem of locating the coefficients of $\mathbf{Q}$ becomes convex and hence efficiently solvable. However, the SOS function (3) is still a non-convex function in configuration space and is a good representation of the original function.

2.6 Sum of squares representations of rational polynomials

So far sum of squares relaxations have been applied to optimisation of polynomial functions. An important category of functions which have not been discussed yet are rational functions. These play an important role in theoretical and computational chemistry as long range interactions such as Van der Waals and electrostatic are commonly represented using rational functions. The most obvious example is the Lennard-Jones potential

$$U_{LJ} = 4\epsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^{6} \right], \tag{2.106}$$

where ϵ is the well depth and σ is the distance at which the potential is zero.

2.6.1 Jibetean's approach

A general approach to providing relaxations to the minimisation of multivariate rational polynomials has been published by Jibetean [196] using a modification of the approach of Parrilo and Lasserre. The problem of minimising a rational polynomial can be expressed as

$$p^* := \inf_{\mathbf{x}\in S,\ q(\mathbf{x})\neq 0} f, \quad \text{where } f := \frac{p(\mathbf{x})}{q(\mathbf{x})}, \tag{2.107}$$

where $p(\mathbf{x})$ and $q(\mathbf{x})$ are polynomials with no common factors and real coefficients. S is either defined to be $\mathbb{R}^n$, (i.e. unconstrained optimisation) or a semialgebgraic set. An important qualification is that in order for f to be bounded from below q cannot change sign on S. Using this assumption the minimum can be expressed as

$$p^* = \sup\{\alpha \mid p(\mathbf{x}) - \alpha q(\mathbf{x}) \geq 0, \quad \forall \mathbf{x} \in S\}. \tag{2.108}$$

Relaxing the nonnegativity condition to a sum of squares condition Jibetean shows results can be obtained which differ according to the following categories, unconstrained univariate, bivariate and multivariate and constrained univariate and multivariate. If the problem is unconstrained the following results apply:

Univarite case: Since the nonnegative and SOS cones are identical when $n = 1$, equation 2.108 can be replaced with

$$p^* = \sup\{\alpha \mid p(\mathbf{x}) - \alpha q(\mathbf{x}) \text{ is SOS}, \quad \forall \mathbf{x} \in S\}. \tag{2.109}$$

Multivariate case: In the multivariate case Jibetean noted that lower bounds $\bar{p}$ can be computed on p^* using the approach of Parrilo outlined earlier to produce a hierarchy of lower bounds

$$\bar{p} = \sup\left\{\alpha \mid (p(\mathbf{x}) - \alpha q(\mathbf{x}))\left(1 + \sum_{i-1}^{n} \mathbf{x}_i^2\right)^r \in \Sigma_{n,d+2r}\right\}, \tag{2.110}$$

where $\Sigma_{n,d}$ is the SOS cone of degree d in n variables. Therefore in the limit $r \to \infty,\ \bar{p} \to p^*$.

In the constrained case the following results apply:

Univariate case: The constraints are defined to restrict the feasible set to the interval $S = [\mathbf{a}, \mathbf{b}]$. Therefore assuming that $q(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in S$

$$p^* = \sup\left\{\alpha \mid p(\mathbf{x}) - \alpha q(\mathbf{x}) = \Sigma_{1,2d} + (\mathbf{x} - \mathbf{a})(\mathbf{b} - \mathbf{x})\Sigma_{1,2d-2}\right\}. \tag{2.111}$$

Multivariate case: Here the problem considered is the same as equation 2.107 but where S is defined to be

$$S := \{\mathbf{x} \in \mathbb{R}^n \mid f_i(\mathbf{x}) \geq 0,\ i = 1, \ldots, s\}. \tag{2.112}$$

A hierarchy of relaxations due to Lasserre [192] are used

$$p^{(r)} = \sup \left\{ \alpha \mid p(\mathbf{x}) - \alpha q(\mathbf{x}) \in \Sigma_{n,2r} + \sum_{i=1}^{s} f_i(\mathbf{x}) \Sigma_{n,2r} \right\}, \tag{2.113}$$

which is polynomial in m, n and in the degrees of p and q for any fixed r. Again in the limit $r \to \infty,\ p^{(r)} \to p^*$.

2.6.2 Change of variables

An alternative approach to Jibetean's which works very well when the rational polynomial has constants in the numerator and variables only in the denominator is to use a change of variables. Therefore it works well for functionals such as the Lennard-Jones and Coulombic potentials.

Given a rational polynomial such as the Lennard-Jones functional (equation 2.106) define a new variable $a = 1/r_{ij}$. Equation 2.106 can then be written as a polynomial in a

$$4\epsilon \left[(\sigma a)^{12} - (\sigma a)^6 \right], \tag{2.114}$$

with the additional constraint $ar_{ij} = 1$. The relative merits of both Jibeteans and this approach for modelling Van der Waal's and Coulombic interactions are discussed in detail in section 3.5.3.

2.7 Conclusions

The approach to global optimisation described in this chapter has emerged over the last five years and has generated significant interest in the optimisation community. Evidence of this interest is a summer school devoted exclusively to this particular topic organised last year which was very well attended and involved many of the leading researchers in the area of convex optimisation and sum of squares theory [209]. The SOS decomposition approach to global optimisation is both mathematically rigorous and beautiful in its deep algebraic foundations, both of which make it very appealing. Past experience has also demonstrated that algorithms with proven polynomial computational complexity bounds tend to become favoured over heuristic approaches. An obvious example is the rise in popularity of interior

point methods when compared to the simplex method which once dominated linear programming.

The realisation that the SOS decomposition problem can be formulated as a convex optimisation problem is therefore very significant indeed and has opened up the possibility of the method being developed into a powerful global optimisation approach. The long term success of this avenue of research in the author's opinion depends on three factors:

- The implementation of specialised mathematical theories such as invariant theory to reduce the dimensionality of both the initial polynomial optimisation problem and the resulting SDP.
- Future progress in the scaling up of interior point methods or alternative approaches to solving semidefinite programmes.
- The willingness of practitioners from the applied optimsiation community and scientists in general who use optimisation techniques to apply the approach to their problems.

The contribution of this thesis clearly fits into the third category. In this work attention is focused on demonstrating the applicability of the approach to modest energy landscape problems which are within the tractability limit of current SDP solvers. This is the subject of the next chapter.

Chapter 3

Sum of squares for energy landscape problems

3.1 Introduction

At the end of chapter 1 the motivation for obtaining rigorous lower bounds to the PES was highlighted and a method (sum of squares decompositions) for obtaining such bounds on polynomial functions was identified and described in detail in chapter 2. This chapter describes the application of the sum of squares approach to potential energy landscapes. The chapter has four primary goals:

- To demonstrate using model potential energy surfaces how the issues with finding the global minimum using current energy landscape methods may be addressed using a sum of squares approach.
- To establish that standard classical potential energy functionals (force fields) can be represented using polynomial functionals and that these polynomials can be used to build a statement of a polynomial optimisation problem which represents the global optimisation problem on potential energy surfaces.
- To show how a sum of squares relaxation may be applied to this polynomial optimisation problem thereby producing provable bounds on the global minimum of the potential energy surface.
- To demonstrate the application of the sum of squares approach to a real chemical system, namely the ethane molecule.

First a brief description is made of the software tools used in obtaining a SOS decomposition and performing global optimisation (section 3.2). Then the topological features of the PES which makes identifying the global minimum difficult are described in section 3.3.1. To evaluate the feasibility of applying the sum of

squares approach to the study of PES problems, three model problems were identified which exhibited these topological characteristics. These model landscapes and their solution using the SOS approach are described in section 3.3.

The remainder of the chapter describes a polynomial representation of the PES and the process of obtaining the global minimum using a SOS decomposition. An early decision to be made was which coordinate system was most suitable for a polynomial representation of the PES. An internal coordinate system was chosen for reasons described in section 3.5.1. A parameterisation for components of the Amber94 force field is described (section 3.5). This is followed by a detailed case study on the application of the SOS decomposition method to the potential energy surface of the ethane molecule as defined by the Amber94 forcefield (section 3.5.2). Finally a discussion of the merits and limitations of applying SOS decompositions to a PES are discussed (section 3.8).

3.2 Software for sum of squares decompositions

All of the SOS decompositions described in this chapter were carried out using a Matlab toolbox SOSTools [210]. The toolbox allows representation of a polynomial optimisation problem in symbolic form using the Matlab symbolic toolbox or the Multivariate Polynomial Toolbox (part of SOSTools). For the majority of examples described here the symbolic toolbox was used, however this was primarily for legacy reasons since SOSTools was first developed for use with the symbolic toolbox. SOSTools has its own commands and syntax for building a sum of squares program. It accepts polynomial objective and constraint symbolic functions and outputs a semidefinite programme. The semidefinite programme is then solved using a semidefinite programme solver such as SeDuMi [211] and the solution is returned to SOSTools which extracts the lower bound. There are currently a number of SDP solvers under development and all are freely available for use under the GNU general public license or similar schemes. Each has advantages and some are more suited to large scale problems than others, however SeDuMi is considered to be the most robust for small to medium scale problems. Because software for solving semidefinite programmes is relatively new there are currently a number of different formats for presenting a semidefinite programme to a solver and the SeDuMi format is one of the most popular. In addition to interfacing with SeDuMi, SOSTools is also configured to interface directly with SDPT3. As part of this project interfaces were written for interfacing SOSTools to SDPARA (a parallel version of the SDPA solver [212]) and to CSDP [213]. This involved reformulating the semidefinite programme without free variables (variables that can take on a negative value) as λ in a standard SOS programme is a free variable. The extension to SDPARA allowed the author

to experiment with running sum of squares optimisation problems on a 12 processor linux cluster with 12 Gb RAM. This was particularly useful for overcoming the memory problems described in appendix B associated with inverting the Hessian in primal-dual interior point algorithms.

SOSTools has routines for producing a minimal monomial vector through the construction of a Newton polytope of monomials as described in section 2.3.4. For small and medium sized problems SOSTools uses the Matlab command `convhulln`, but for large problems it interfaces with an external program CDD [214] to generate the vertices of the convex hull of the monomial points. CDD is a C++ implementation of the Double Description Method of Motzkin et al. for generating all vertices (i.e. extreme points) and extreme rays of a general convex polyhedron in $\mathbb{R}^n$ given by a system of linear inequalities. The program can be used for the reverse operation (i.e. convex hull computation). This means that one can move back and forth between an inequality representation and a generator (i.e. vertex and ray) representation of a polyhedron.

3.3 Model energy landscape problems

3.3.1 Challenging topological features of the PES

The topological features that make finding the global minimum difficult on a PES when using a sampling strategy have been discussed in chapter 1. Two landscape topologies which are particularly problematic for the sampling techniques described are:

- Landscapes with minima which are low in energy yet far apart in configurational space. This can be understood in terms of a disconnectivity graph with many minima on a pathway connecting the two low lying minima. The case of the LJ_{38} cluster is a good example of a system with such a PES topology.

- Landscapes that are mainly flat except for a localised region where the there is a deep tree leading to the global minimum. This type of topology is particularly problematic for sampling algorithms which must sample very close to the global minimum to identify the presence of the well in the first place.

This section describes the application of the SOS approach to model landscapes with these topologies and demonstrates its advantages over the conventional approaches. The first three examples (section 3.3.2) address item 1 while item 2 is addressed in the fourth example (section 3.3.3).

A fifth model landscape demonstrates how the SOS approach can be applied to non-polynomial landscapes through a simple fitting procedure. This landscape

is also an important model which is widely used to study dynamics, for example see [215,216]. The application of the SOS approach to these model energy landscapes has been recently published by the author in [217].

3.3.2 Multi-well potentials

Here are three potentials with many minima whose values are very close to the global minimum but are separated by a number of high barriers (relative to the energy difference between the minima).

1-dimensional 16^{th} degree polynomial landscape

The first problem is a simple one dimensional problem (see figure 3.1) in which the typical ratio of barrier height/local minima energy difference is 80. In addition, the minimum closest in energy to the global minimum is furthest apart in domain space. This type of poor connectivity is not uncommon in potential energy surface problems where a number of nodes on the disconnectivity graph have energy levels that are almost identical (banyan tree disconnectivity graph - see section 1.4.2). The function is defined as

$$\begin{aligned} U = {} & 0.01x - 128.1x^2 + 2688x^4 - 21504x^6 + 84480x^8 \\ & -180224x^{10} + 212992x^{12} - 131072x^{14} + 32768x^{16}. \end{aligned} \tag{3.1}$$

The SOS approach provides a lower bound and the coordinates of the minimiser are obtained from the dual solution (table 3.1). The lower bound is verified as the global minimum by comparing the upper bound (produced by evaluating the function at the minimiser) with the lower bound produced in the primal problem. As can be seen from figure 3.1, the SOS decomposition obtained is a faithful representation of the original function. The SOS approach was compared to two other common minimisation approaches: the downhill simplex method [218] and the BFGS quasi-newton scheme [55–60], (used in the basin-hopping algorithm). Both methods were applied iteratively from a grid of starting points spanning the domain of the problem with the grid density being increased until the global minimum was obtained. To find the global minimum, at least 10 separate minimisations from starting points on an evenly spaced grid of points spanning the domain were required. This is in fact more attempts than the number of minima in the domain. The point however remains that these algorithms have no way of establishing that the minimum that they locate is in fact the global minimum.

It should be noted that there is nothing special about the form of the potential presented. Several multi-well examples with different polynomial representations, more asymmetric wells, and, separations have been addressed with identical success.

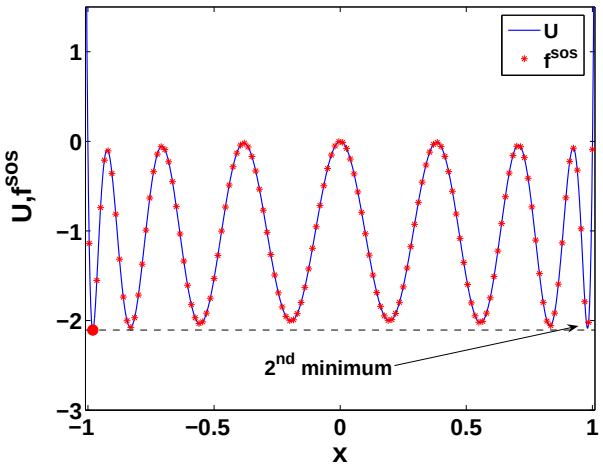


Figure 3.1: Sum of squares optimisation on a multi-well potential U(-) with 8 minima, the leftmost being the global minimum. The second minimum, farthest away in domain space, is very close in value to the global minimum. The global minimum is correctly identified. The SOS function $f^{SOS}(*)$ is in excellent agreement with U.

Interestingly, the SOS approach can identify the global minimum even if the wells are nearly-degenerate with arbitrarily close energy.

4-dimensional, 16^{th} degree polynomial landscape

Another example is a polynomial function in 4 variables and of degree 16 given by

$$U(x_1,\ldots,x_4) = \sum_{i=1}^{4} a_0 + a_1 x_i + a_2 x_i^2 + a_3 x_i^4 + a_4 x_i^6 + a_5 x_i^8 + a_6 x_i^{10} + a_7 x_i^{12} + a_8 x_i^{14} + a_9 x_i^{16}, \tag{3.2}$$

where $a_0 = 1$, $a_1 = 0.5$, $a_2 = -130.1$, $a_3 = 2688$, $a_4 = -21504$, $a_5 = 84480$, $a_6 = -180224$, $a_7 = 212992$, $a_8 = -131072$, $a_9 = 32768$. The lower bound, $U = -14.036$, was identified as the global minimum and the location of the minimiser, $[-0.9798, -0.9798, -0.9798, -0.9798]$ was correctly located among on the 4096 densely spaced minima (figure 3.2).

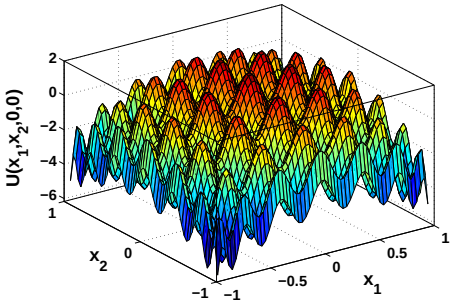


Figure 3.2: A 2D projection of a 4-dimensional 16^{th} degree multi-well potential (equation 3.2).

10-dimensional, 6^{th} degree polynomial landscape

Finally the SOS approach was applied to a ten dimensional potential with cross terms given by

$$U(x_1,\ldots,x_{10}) = \sum_{i=1}^{5} a_1 x_i + a_2 x_i^2 + a_3 x_i^3 + a_4 x_i^4 + a_5 x_5 + a_6 x^6 + \sum_{i=6}^{10} b_1 x_i + a_2 x_i^2 + a_3 x_i^3 + a_4 x_i^4 + a_5 x_5 + a_6 x^6 + \sum_{\substack{i=1,\ldots,9 \\ j=i+1,\ldots,10}} r_{(1)ij} x_i x_j - r_{(2)ij} x_i^2 x_j^4, \tag{3.3}$$

where $r_{(1)ij}$ and $r_{(2)ij}$ are real numbers chosen randomly between 0 and 10. The lower bound on the global minimum of $U = -193.685$ was exact, and the minimiser $[-0.7237, 0.7468, 0.7431, 0.7487, -0.7226, 0.7477, 0.7463, 0.7491, 0.7531, -0.7319]$ was identified. The duality gap was -2.316×10^{-3}.

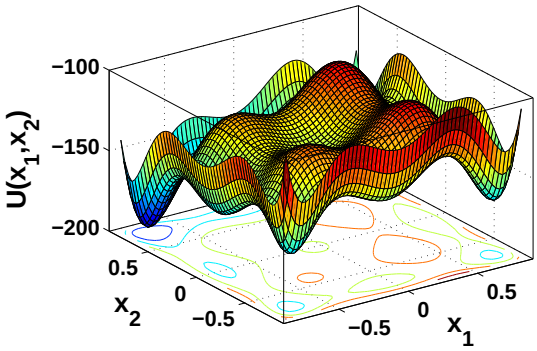


Figure 3.3: A 2D projection of a 10-dimensional multi-well potential.

3.3.3 Golf-course-like potential

Another category of energy landscapes exhibit flat regions with deep and narrow wells which become especially challenging and rapidly intractable for searching algorithms. The reason is that the flat regions provide no energy differences that can guide any search. Exhaustive searching is the only option, which becomes intractable very rapidly as the domain increases. Using a SOS decomposition it was possible to locate the global minimum of a delta function in $\mathbb{R}^2$. The potential is defined by

$$\frac{-1}{\pi}\left(\frac{e_1}{(x-5)^2+(y-4)^2+e_1^2}+\frac{e_2}{(x+7)^2+(y+6)^2+e_2^2}\right), \tag{3.4}$$

where $e_1 = 0.09, e_2 = 0.1$. The potential is shown in figure 3.4. Note that the function is almost flat everywhere except in two locations where almost identical delta function wells are present and that it is defined over the entire $\mathbb{R}^2$ domain. The SOS approach correctly identifies the minimum and its location (table 3.1). It also succeeds when the higher energy well is widened.

3.3.4 Müller potential

The Müller potential [219] is commonly used [215, 216] as a non-trivial test for reaction path methods due to the complexity of its minima structure that leads to a highly contorted reaction coordinate which requires a sharp change of direction at the saddle. The potential surface (figure 3.5a) contains three minima, with the global minimum of -146.7 at $(-0.558, 1.442)$:

$$U(x,y) = \sum_{i=1,\ldots,4} A_i \exp[a_i(x-\hat{x}_i)^2 + b_i(x-\hat{x}_i)(y-\hat{y}_i) + c_i(y-\hat{y}_i)^2], \tag{3.5}$$

where $A = (-200, -100, -170, 15)$, $a = (-1, -1, -6.5, 0.7)$, $b = (0, 0, 11, 0.6)$, $c = (-10, -10, -6.5, 0.7)$, $\hat{x} = (1, 0, -0.5, -1)$, $\hat{y} = (0, 0.5, 1.5, 1)$. Its domain is constrained in order to exclude areas where the function is not well behaved. The definition of this function includes an exponential so in order to apply the method the function was fitted in a least squares sense to a polynomial of degree 16, U_f, which contains 153 monomials, 67 of which are odd. This leads to the minimisation of U_f subject to a series of inequality constraints, g_i, using a Positivstellensatz

representation:

$$\begin{aligned} \text{max:} \quad & \lambda, \\ \text{subject to:} \quad & U_f - \sum_{i=1}^{4} s_i g_i - \lambda \succeq 0 \\ & g1 : x \geq -1.5, \quad g2 : x \leq 0.7, \\ & g3 : y \geq -0.5, \quad g4 : y \leq 2.0, \end{aligned}$$

where s_i are additional SOS polynomials with unknown coefficients.

The SOS approach succeeded in locating the minimum which was proven to be exact by comparing the upper and lower bounds. It is interesting to point out that the final SOS decomposition of the potential is a good representation of the original function (see figure 3.5) which may allow for additional properties of the function to be explored in different representations. For example, it was possible to correctly locate the saddle through eigenvector following on the SOS function.

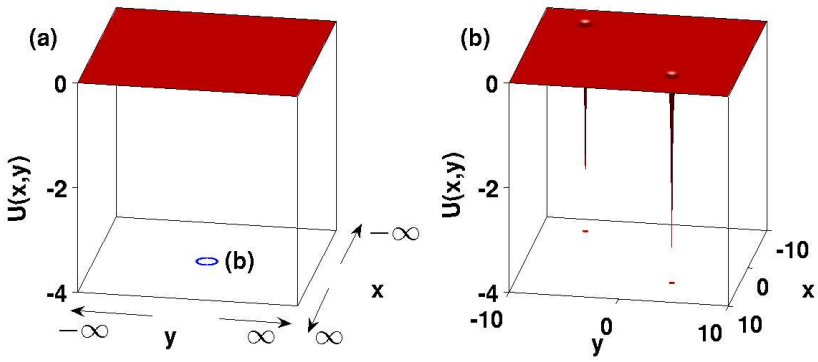


Figure 3.4: Contour plot of a Golf-like potential. The potential is defined over the entire $\mathbb{R}^2$ domain (left) .The area close to the origin (blue circle) is reproduced at greater magnification on the right. The two minima are $U(5,4) = -3.5369$ (global) and $U(-7,6) = -3.1832$. In this case the SOS optimisation produces a lower bound of $U = -3.539$. The configurational coordinates are obtained correctly to within numerical accuracy.

3.4 Symmetric polynomials

The study of the properties of symmetric polynomials forms a rich branch of mathematics called invariant theory which has been the subject of much interest in recent years. Invariant theory can on occasion, where the SOS optimisation problem being considered possesses inherent group symmetry, be used to reduce the size the problem. In the simplest case where the polynomial being minimised has a monomial

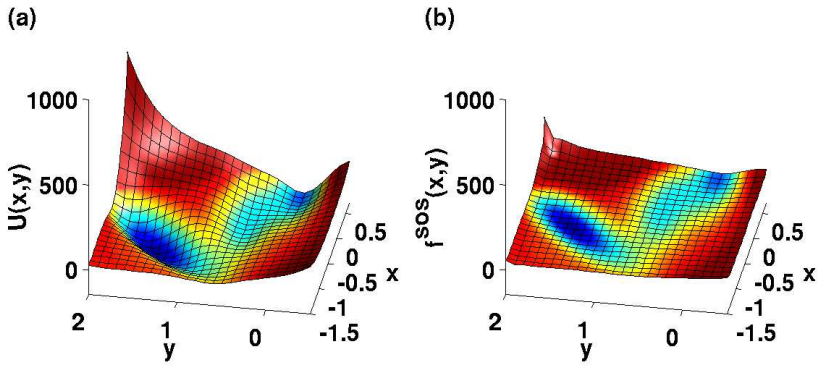


Figure 3.5: Contour plot of the original fitted Müller potential (a), and its SOS representation (b). The global minimum was identified correctly. f^{SOS} (b) is the projection to the original domain space from the lifted space of the SOS function including the inequality constraints. The difference between the two (a) and (b) is due to numerical errors during the SOS decomposition calculation, however much of the original function is preserved.

	1D Multi-well	10D Multi-well	Müller	Golf-like
CPU time (s)	0.65	6840.00	10.06	0.80
Lower bound	-2.106 [-2.106]	-193.685	-146.674 [-146.674]	-3.539 [-3.537]
Configuration	-0.980	see text	(-0.557,1.449)	(4.994,3.995)
Duality gap	3.857×10^{-6}	7.560×10^{-4}	1.366×10^{-4}	2.000×10^{-3}
Primal	2.105998	196.6848	146.6743	3.537
Dual	2.106002	193.6855	146.6745	3.539

Table 3.1: Summary of SOS optimisation results for the multi-well, Müller and golf-like potentials: the cpu time in a standard PC in seconds, the lower bound value with the exact global minimum in brackets, the configurational coordinates of the minimum, the duality gap, and the final values of the primal and dual problem.

of even degree as its monomial with the lowest non-zero coefficient, the polynomial has symmetry about the origin. Consider the one dimensional quartic polynomial,

$$U = x^4 - 40x^2, \tag{3.6}$$

shown in figure 3.6.

This 1-dimensional polynomial has no monomial of degree 1 and therefore has symmetry about the origin. It degenerate minimisers lie at $x = \pm 4.4722$ having identical values of -400.0. If this function is submitted to SOSTools for minimisation, the lower bound obtained is exact at -400.0. However the following dual matrix is obtained

$$\begin{bmatrix} 1.0000 & 0.0000 & 20.0002 \\ 0.0000 & 20.0002 & 0.0000 \\ 20.0002 & 0.0000 & 400.0091 \end{bmatrix} \tag{3.7}$$

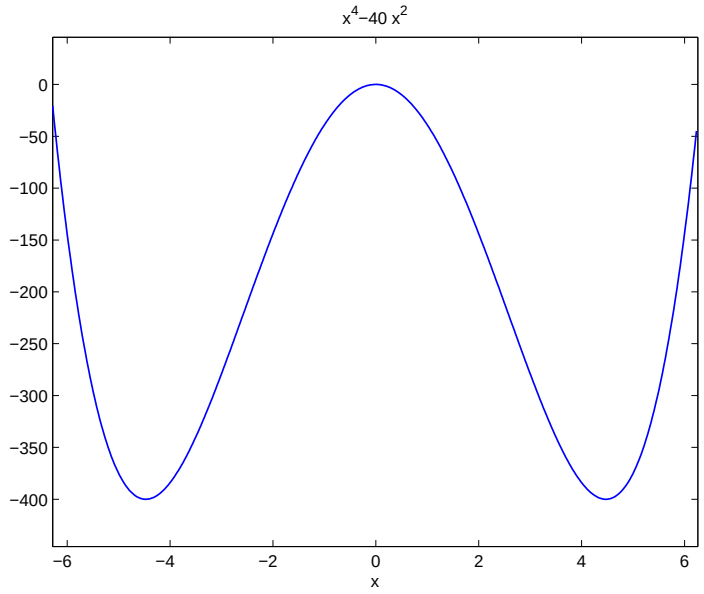


Figure 3.6: The polynomial function $x^4 - 40x^2$. It has two degenerate minima at $x = \pm 4.4722$ having identical values of -400.0.

An eigenvalue decompostion of this dual matrix gives

$$\left[401.0091, 20.00023, 5.259617 \times 10^{-11}\right] \tag{3.8}$$

showing that it has rank $= 2$. This implyies that the SOS relaxation is not exact.

The monomial basis vector for the SOS decompostion is as follows

$$\left[1, x, x^2\right]. \tag{3.9}$$

Evaluating the function at $x = 0.000$ from the second entry in the first row of the dual matrix gives an overestimate of the lower bound at $f(x = 0.000) = 0.000$. However mindful of the fact that the polynomial has an x^2 monomial as its monomial of lowest degree, the minimiser should instead be taken from the third entry in the first row of the dual matrix. This gives a minimiser at $x^2 = 20.0002$, i.e. $x = \pm 4.4722$ which are the exact minimisers of the function.

This method of extracting the minimisers from the dual solution matrix also works in higher dimensions. Consider the two dimensional equivalent of equation 3.6,

$$U = x^4 - 40x^2 + y^4 - 40y^2, \tag{3.10}$$

which is shown in figure 3.7 Again the lowest degree of any monomial is 2. The dual

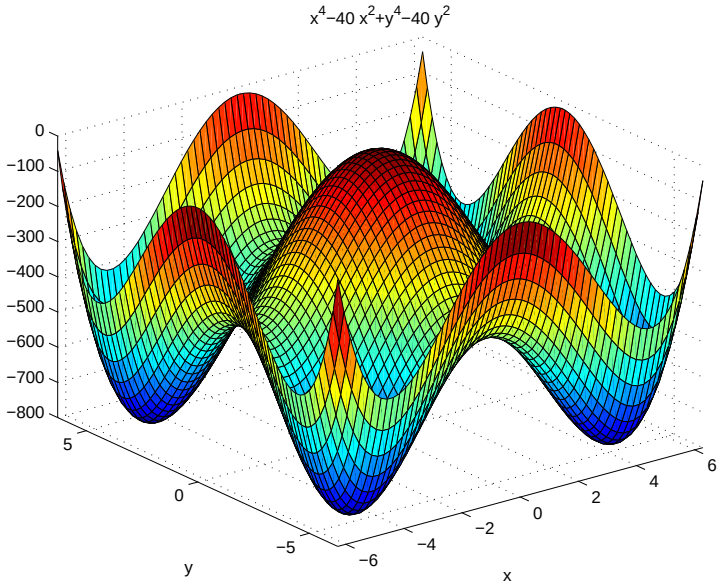


Figure 3.7: The polynomial function $x^4 - 40x^2$. It has four degenerate minima at $(x = \pm 4.4722, y = \pm 4.4722)$ having identical values of -800.0.

matrix obtained from SOSTools is

$$\begin{bmatrix} 1.0000 & -0.0000 & 0.0000 & 20.0003 & -0.0001 & 20.0003 \\ -0.0000 & 20.0003 & -0.0001 & -0.0000 & 0.0001 & -0.0000 \\ 0.0000 & -0.0001 & 20.0003 & 0.0001 & -0.0000 & 0.0001 \\ 20.0003 & -0.0000 & 0.0001 & 400.0116 & -0.0012 & 400.0116 \\ -0.0001 & 0.0001 & -0.0000 & -0.0012 & 400.0116 & -0.0012 \\ 20.0003 & -0.0000 & 0.0001 & 400.0116 & -0.0012 & 400.0116 \end{bmatrix}, \tag{3.11}$$

and the eigenvalue decomposition is

$$\left[801.0232,\ 400.0116,\ 20.00035,\ 20.00023,\ 9.161582 \times 10^{-08},\ 1.358453 \times 10^{-10}\right], \tag{3.12}$$

implying that the rank of the dual is 4. There appears to be a direct link between the rank of the dual and the number of degenerate global minima. The monomial vector is

$$\left[1, x, y, x^2, xy, y^2\right], \tag{3.13}$$

and the correct minimisers can be obtained from the 4^{th} and 6^{th} entries of the first row or column of the dual matrix. Therefore the minimisers are identified as being at $(x^2 = 20.003, y^2 = 20.003)$ and therefore $(x = \pm 4.4722, y = \pm 4.4722)$. Figure 3.8

shows a decision tree flowchart for postprocessing of SOS optimisation problems.

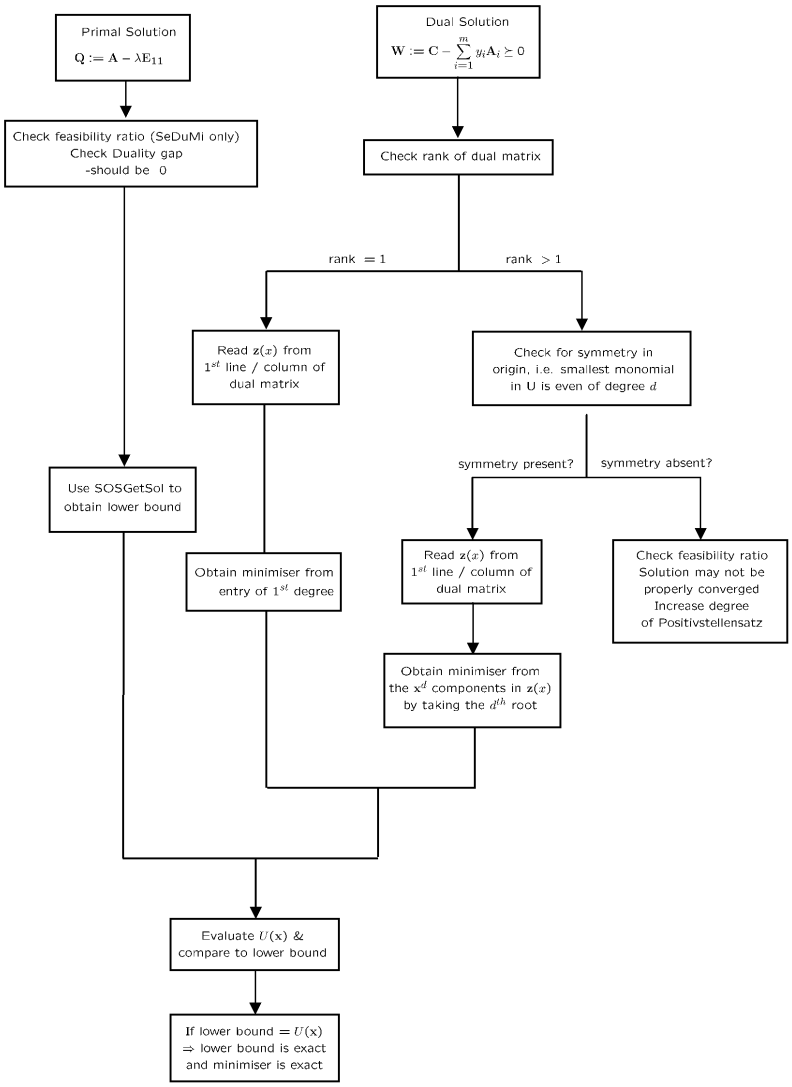


Figure 3.8: SOS optimisation postprocessing flowchart. The primal solution produces a lower bound which can be verified as tight by evaluating U at the minimiser obtained from the dual solution matrix. If the dual solution has rank $= 1$, the minimiser can be read directly from the appropriate entries in the first line or column. If the rank > 1, and if the lowest degree monomial in U has even degree d the minimiser coordinates should be read from the $\mathbf{x}^d$ entries of the first row or column by taking the d^{th} root. Alternatively if the lowest degree of U is $d = 1$, then the solution may not have fully converged and it is necessary to increase the degree of the Positivstellensatz.

3.5 Polynomial representation of a classical force field

The next topic to be discussed is how to represent a classical potential energy surface in polynomial form which of course is necessary for treatment using the SOS approach. In section 1.1.2 some widely used mathematical descriptions of potential energy surfaces (force fields) were discussed. In this section one of these, the Amber94 potential [19] is examined and a possible approach for mapping it into a form

suitable for treatment using the SOS approach is described. Amber94 is a widely used force field for biomolecular simulation being one of the most highly cited in the literature. One of the reasons for the importance of the Amber94 force field is that it was the first in which the parameterisation was performed on an algorithmic basis thereby ensuring extensibility to other molecules and functional groups not considered in the initial development. In fact a recent update (the Generalised Amber Force Field [220]) has extended it to cover organic molecules in general making applications such as drug-receptor binding free energy calculations accessible. However since most of the common functions have similar form the approach described here is more or less general to the other popular force fields.

3.5.1 The Amber94 force field

In common with other force fields the Amber94 force field is described in terms of internal coordinates as described in section 1.1.2 and the functional form is a sum of terms which can be divided into bonded and non-bonded terms:

$$E_{\text{total}} = \sum_{\text{bonds}} K_r(r - r_{eq})^2 + \sum_{\text{angles}} K_\theta(\theta - \theta_{eq})^2 + \sum_{\text{dihedrals}} \frac{V_n}{2}[1 + \cos(n\phi - \gamma)] + \sum_{i<j} \frac{A_{ij}}{r_{ij}^{12}} - \frac{B_{ij}}{r_{ij}^{6}} + \sum_{i<j} \frac{q_i q_j}{\mathcal{C} r_{ij}} \quad (3.14)$$

where the first three terms are the bonded terms representing the stretches, bends and torsions. The stretches are simple harmonic functions in the distance between two covalently bonded atoms and the angle terms are harmonic functions in the angle between a triple of atoms covalently bonded in a linear sequence. The torsional term is a Fourier expansion which represents the four body term between a linear sequence of covalently bonded atoms. The fourth term is the Lennard-Jones representation of the van der Waals interaction. Van der Waals interactions are a combination of the dispersion attractions which arise due to dipole-dipole interactions and the hard core repulsion between the nuclei. The most popular representation of these interactions is the Lennard-Jones functional which is a rational polynomial of degree 12 in the interatomic distance r_{ij} and the A_{ij} and B_{ij} are atom dependent parameters. The Lennard-Jones representation is particularly popular in computational evaluations as the 12^{th} degree term is just the square of the 6^{th} degree term making it more efficient to evaluate. The fifth term represents the interatomic Coulombic interactions where the q_i are the atomic charges and $\mathcal{C}$ is a constant. Each of the five terms is illustrated schematically in figure 1.2. It should be noted that there is no reason why this particular form of subdivision of the potential energy functional should be any more effective than another, however this form has become more or less standard, although some force fields include a separate term to capture hydrogen bonding.

The parameterisation is based upon the atom types involved in the interaction and the particular bonding configuration of the atom in question, for example an sp^3 carbon is given a different parameterisation to a double bonded sp^2 carbon which is also parameterised differently to an sp^2 carbon in an aromatic 5-membered ring. Further details are given in table 1 in [19].

Both internal coordinate and Cartesian coordinate bases were examined in this work, however it was found that internal coordinate representations were more suitable as it kept the number of variables to a minimum thereby reducing the size of the monomial basis and therefore the SDP which needed to be solved.

3.5.2 Ethane as a case study

In order to describe the application of the SOS method to a real system which contains the five components of the classical potential energy functional, the remainder of this section will describe the application of the method to finding the global minimum of the ethane molecule in vacuo (figure 3.9). Ethane, a simple hydrocarbon, is one of the smallest molecules for which all five terms in the potential energy functional are applicable since non-bonded terms are typically not included between atoms separated by less than three covalent bonds. In atoms connected by three or less covalent bonds the dispersion and Coulombic interactions are included in the stretch and bend parameters. Therefore in the case of ethane there are 7 stretches, 12 bends, 9 torsions and 9 van der Waals and Coulombic interactions making a total of 37 interaction terms for a molecule of just 8 atoms. The two carbons in ethane

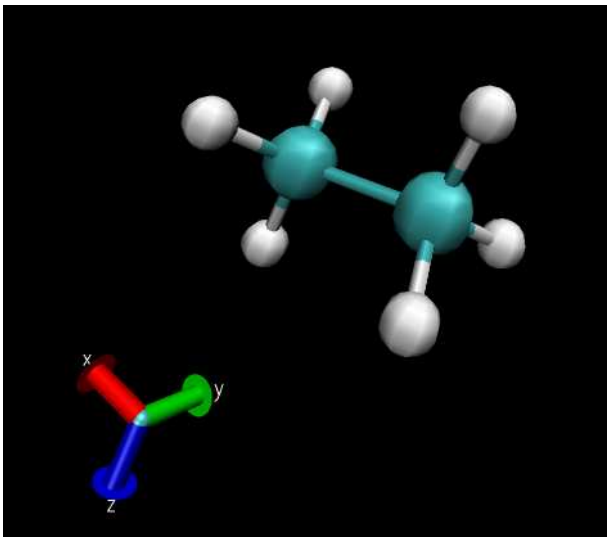


Figure 3.9: CPK representation of Ethane (C_2H_6) in the trans configuration. Carbons are rendered in cyan and hydrogens are rendered in white. Covalent bonds are shown as rods. Rendered using Visual Molecular Dynamics (VMD) [221].

are sp^3 hybridised and a single bond exists between them and between each carbon

and its three hydrogens. In general the stretch and bend internal coordinates are much stiffer than the torsional degrees of freedom as will be discussed in the coming paragraphs. As a result the torsional rotations are strongly coupled resulting in rotation of the two methyl groups as independent units relative to each other about the axis defined by the carbon-carbon bond. A numbering system will be useful for describing the hydrogens in ethane. The hydrogens are labelled according to their methyl group and are numbered 1 to 3 according to the following:

1. One of the methyl groups is labelled A and the other is labelled B.
2. The molecule is represented in its trans position.
3. One hydrogen on methyl group A is chosen as hydrogen 1 on methyl group A. This is therefore HA1.
4. Looking at the carbon-carbon bond end on, with methyl group A closest to the viewer, the remaining hydrogens in methyl group A are numbered in a clockwise increasing order giving HA2 and HA3.
5. This procedure is repeated for methyl group B by rotating the molecule to look down the carbon-carbon bond with methyl group B nearest the viewer. The hydrogen furthest from HA1 is labelled HB1 and the remaining hydrogens are labelled in a clockwise increasing order to give HB2 and HB3.

This labelling scheme is illustrated in figure 3.10.

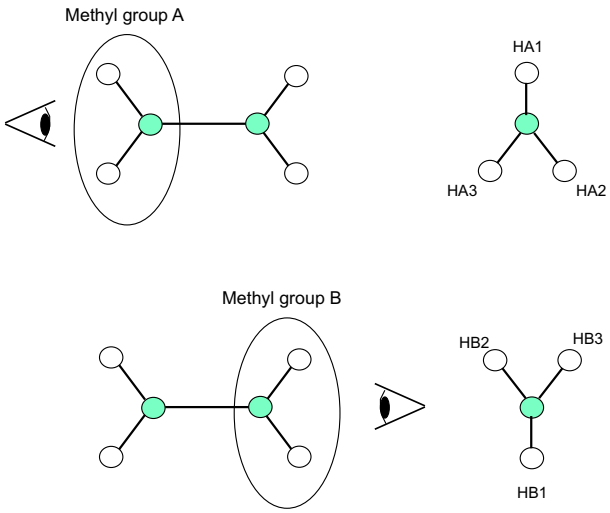


Figure 3.10: Numbering scheme for the hydrogen molecules in ethane.

3.5.3 Mapping the terms of the potential energy functional into polynomial form.

Stretch and bend terms

The representation of the stretch and bend terms in the SOS format is identical to their representation in the Amber94 functional. In fact in many calculations both stretch and bend internal coordinates are fixed at their equilibrium point, due to the high force constants involved, which result in only modest fluctuations from these equilibrium points at room temperature. For example consider the stretch term in the Amber94 force field between two sp^3 carbons (coded CT-CT in [19]):

$$E_{\text{stretch}} = 310.0\text{kcal}/(\text{mol Å}^2)(r_{ij} - 1.526\text{Å})^2. \tag{3.15}$$

This functional is plotted in figure 3.11. Notice that a ± 1 Angstrom displacement from the equilibrium distance raises the potential energy to over one hundred times that required for a full torsional rotation about a CT-CT bond (see figure 3.12). A similar case can be made for fixing the bending internal coordinates, with the

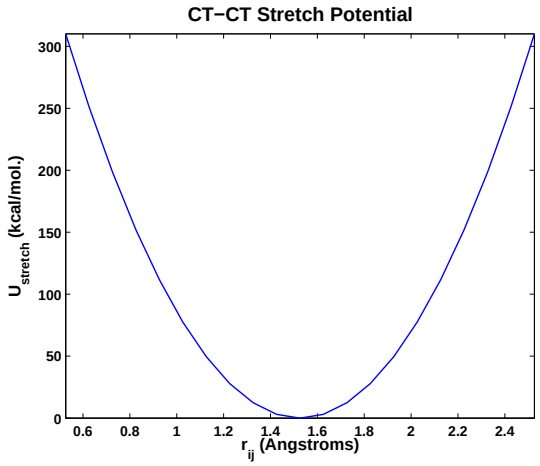


Figure 3.11: Potential energy function for bond stretching between two sp^3 carbons.

potential energy functional given by:

$$E_{\text{bend}} = 40.0\text{kcal}/(\text{mol radian}^2)(\theta_{ijk} - 109.50 \text{ degrees})^2. \tag{3.16}$$

For the ethane molecule the approach of fixing the bonds and stretches is adopted for these reasons. Therefore it is possible to consider the potential energy functional to be a function of a single parameter, namely the torsional angle about the carbon-carbon bond.

Torsional terms.

In the ethane example all the torsional terms involve rotation of two end hydrogens about a carbon-carbon axis where both carbons are sp^3 hybridised. Therefore in the Amber94 force field the appropriate torsional term is represented as X-CT-CT-X where X can be any atom and the CT's represent sp^3 hybridised carbons. The relevant parameters for the Fourier series description of the torsional interaction, as shown in equation 3.14 are: i) the magnitude of the torsion in kcal/mol. : $V_n = 1.40$, ii) the phase offset in degrees : $\gamma = 0.0$, and the periodicity of the torsion : $n = 3.0$. The final parameter is the number of paths for the torsion which in this case is 9 so that each torsion has $V_n/18$ assigned to it. For the torsional terms a one dimensional polynomial fitting was used to represent the Fourier series. The fitted polynomial is (to five significant places) given by (see figure 3.12)

$$\begin{aligned} E_{\text{torsion}} \;=\; & 77.153\phi^{14} - 0.48792\phi^{13} - 383.76\phi^{12} + 2.1621\phi^{11} \\ & +826.09\phi^{10} - 3.9985\phi^{9} - 978.69\phi^{8} + 3.8710\phi^{7} \\ & +656.72\phi^{6} - 1.9735\phi^{5} - 227.07\phi^{4} + 0.45804\phi^{3} \\ & +30.943\phi^{2} - 0.031312\phi + 0.00082950. \end{aligned} \tag{3.17}$$

All polynomial fitting reported is least squares fitting and in this case the norm of the residual was 0.016816.

Lennard-Jones

The two non-bonded terms namely the Lennard-Jones and Coulombic interactions are rational polynomials and as such are more challenging to represent in polynomial form. There are two options for representing these terms:

1. Use the approach of Jibetean described in section 2.6.1.
2. Express the non-bonded terms as polynomials in $1/r_{ij}$ as described in section 2.6.2.

Both methods were tested and while Jibetean's approach works well for one dimensional problems it has an obvious scaling problem for multidimensional problems and in particular for rational polynomials of 12^{th} degree, which are required to consider Lennard-Jones interactions. Consider the problem of three atoms a, b, c interacting through a Lennard-Jones potential. The energy functional for this interaction is:

$$U_{LJ} = \left(\frac{A_{ab}}{r_{ab}^{12}} - \frac{B_{ab}}{r_{ab}^{6}}\right) + \left(\frac{A_{bc}}{r_{bc}^{12}} - \frac{B_{bc}}{r_{bc}^{6}}\right) + \left(\frac{A_{ac}}{r_{ac}^{12}} - \frac{B_{ac}}{r_{ac}^{6}}\right), \tag{3.18}$$

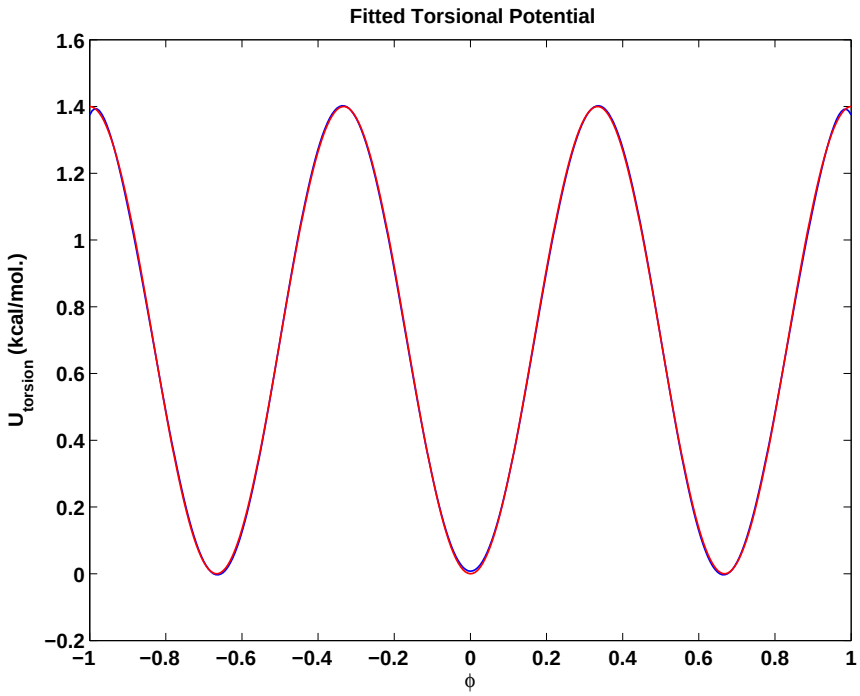


Figure 3.12: A polynomial fitting to a torsional potential energy functional for an X-CT-CT-X torsional bond where CT represents an sp^3 carbon and X represents any atom bonded to the sp^3 carbon. The original functional is shown in red and the fitted polynomial is shown in blue. ϕ is the torsional variable which is rescaled to vary between -1.0 and 1.0.

where A_{ij} and B_{ij} are the appropriate constants. To use Jibetean's method requires that the rational polynomial be written in the form p/q as follows:

$$\frac{(A_{ab}r_{bc}^{12}r_{ac}^{12} - B_{ab}r_{ab}^{6}r_{bc}^{12}r_{ac}^{12}) + (A_{bc}r_{ab}^{12}r_{ac}^{12} - B_{bc}r_{ab}^{12}r_{bc}^{6}r_{ac}^{12}) + (A_{ac}r_{ab}^{12}r_{bc}^{12} - B_{ac}r_{ab}^{12}r_{bc}^{12}r_{ac}^{6})}{r_{ab}^{12}r_{bc}^{12}r_{ac}^{12}}.$$

When written in the form $p(x) - \alpha q(x)$ (see section 2.6.1) this gives a polynomial of degree 36 for interactions between 3 particles. For ethane in which there are 9 pairwise interactions, this rises to a polynomial of degree 108.

Instead it was found that expressing the rational polynomial as a polynomial in $1/r_{ij}$ (option 2) proved to be the most effective approach for multi-dimensional problems. In ethane there are 9 van der Waals and 9 Coulombic interactions to consider between the 6 hydrogens. The first step is to obtain $1/r_{ij}$ as a function of ϕ for each pair of hydrogens in different methyl groups. Figure 3.13 shows the distance (r_{ij}) between the hydrogen atoms as a function of the torsional angle ϕ. This highlights the symmetry in the ethane molecule (at least when stretches and angles are modelled as fixed). The 9 inter-hydrogen distances can be divided into 3 sets with all distances in each set being identical. The sets are {HA1-HB1, HA2-HB3, HA3-HB2}, {HA1-HB2, HA2-HB1, HA3-HB3}, {HA1-HB3, HA2-HB2, HA3-HB1}.

This means that only 3 distance functions are required to describe the distance - torsional angle dependence.

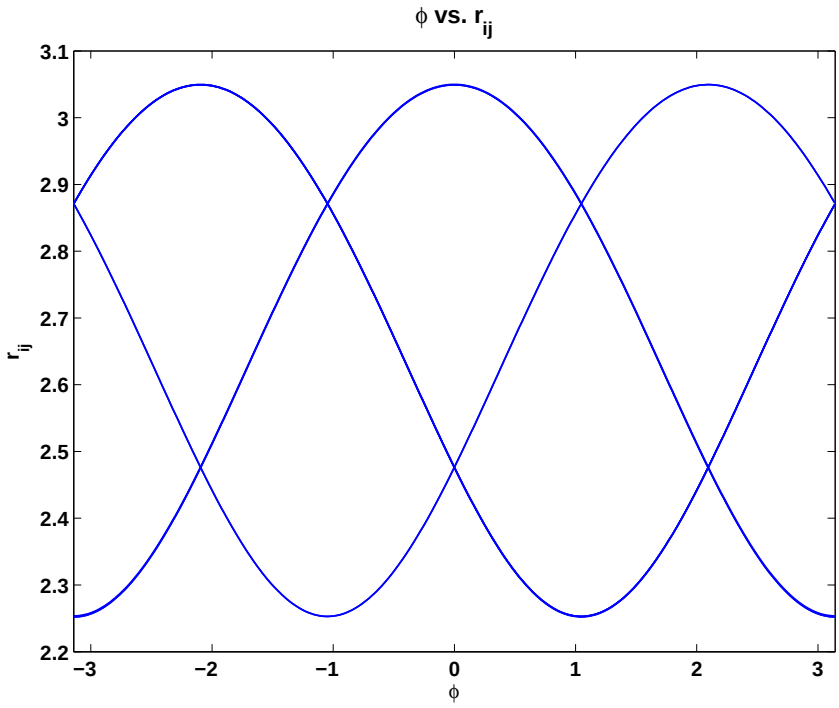


Figure 3.13: Plots of ϕ vs. r_{ij} for the hydrogens in ethane. Although there are 9 pairwise interactions, 3 sets of 3 distances are identical meaning that only 3 distances are required to describe the non-bonded energy functionals.

This $\phi - r_{ij}$ relationship can be converted into a relationship between ϕ and $1/r_{ij}$. These curves are still smooth and continuous and can be approximated reliably using polynomial functions. Figure 3.14 shows a plot of ϕ vs. $1/r_{ij}$ and the fitted polynomials.

The polynomials that were chosen were the following:

$$\begin{aligned} P_1 := 0.38740 \times 10^{-1}x^8 - 0.19035 \times 10^{-3}x^7 - 0.13182x^6 + 0.44786 \times 10^{-3}x^5 \\ +0.28169 \times 10^{-1}x^4 - 0.75224 \times 10^{-4}x^3 \\ +0.18055x^2 - 0.18150 \times 10^{-3}x + 0.32793, \end{aligned} \tag{3.19}$$

$$\begin{aligned} P_2 := -0.14126x^8 - 0.14561x^7 + 0.31544x^6 + 0.40187x^5 - 0.17084x^4 \\ -0.44495x^3 - 0.59825 \times 10^{-1}x^2 + 0.18800x + 0.40410, \end{aligned} \tag{3.20}$$

$$P_3 := -0.14053x^8 + 0.14640x^7 + 0.31319x^6 - 0.40329x^5 - 0.16819x^4 + 0.44545x^3 - 0.61274 \times 10^{-1}x^2 - 0.18786x + 0.40429. \quad (3.21)$$

The norm of the residuals from the least squares fit for these polynomials were 0.001083, 0.008080 and 0.016816 respectively.

Although the van der Waals contributions to the overall potential energy will be small, the contribution from the electrostatics will be more significant and since this is also pairwise distance dependent, accurate calculation of the electrostatics will also require these polynomials to be accurate and favourable to numerical solution.

The van der Waals energy also requires calculation of the atom dependent constants A_{ij} and B_{ij}. These are defined as follows:

$$\begin{aligned} A_{ij} &= \epsilon_{ij} r_{0ij}^{12}, \\ B_{ij} &= 2\epsilon_{ij} r_{0ij}^{6}, \end{aligned} \quad (3.22)$$

where r_{0ij} is the van der Waals equilibrium distance in Angstroms and ϵ_{ij} is the van der Waals well depth in kcal/mol. The parameters from the Amber94 force field for a H-H interaction are $r_{0ij} = 1.4870$ Angstroms and $\epsilon_{ij} = 0.0157$ kcal/mol. A final important subtlety within the Amber94 force field is that for atoms separated by 3 covalent bonds i.e. 1-4 interactions, all van der Waals interaction energies are scaled by 0.5. The reader is referred to the Amber94 paper [19] for a discussion of this scaling. The van der Waals interactions are plotted as a function of ϕ in figure 3.15. Notice that there are three distinct peaks corresponding to the three possible eclipsed conformations of the ethane molecule.

Coulombic interactions

In the Amber94 forcefield like in many of the other widely used force fields, atomic charges and dipoles are represented by atom centred partial charges. These are usually derived by fitting the charges to a quantum mechanically derived electrostatic potential (ESP). It has been found that an ESP fit from a quantum mechanical calculation performed at the 6-31G* level of Gaussian type orbital approximation gives excellent reproduction of inter-molecular properties such as liquid enthalpies, densities and free energies of solvation [19]. The Amber94 force field uses an improved version of the ESP fit known as a restrained ESP (RESP) fit which overcomes some of the issues with the standard ESP approach (see [19] for a more in depth discussion). These partial charges can be used in the Coulombic part of the potential energy functional (equation 3.14) to calculate electrostatic interactions in vacuo. In a condensed phase the electrostatic shielding by the solvent must be accounted for but this is not dealt with in this current work.

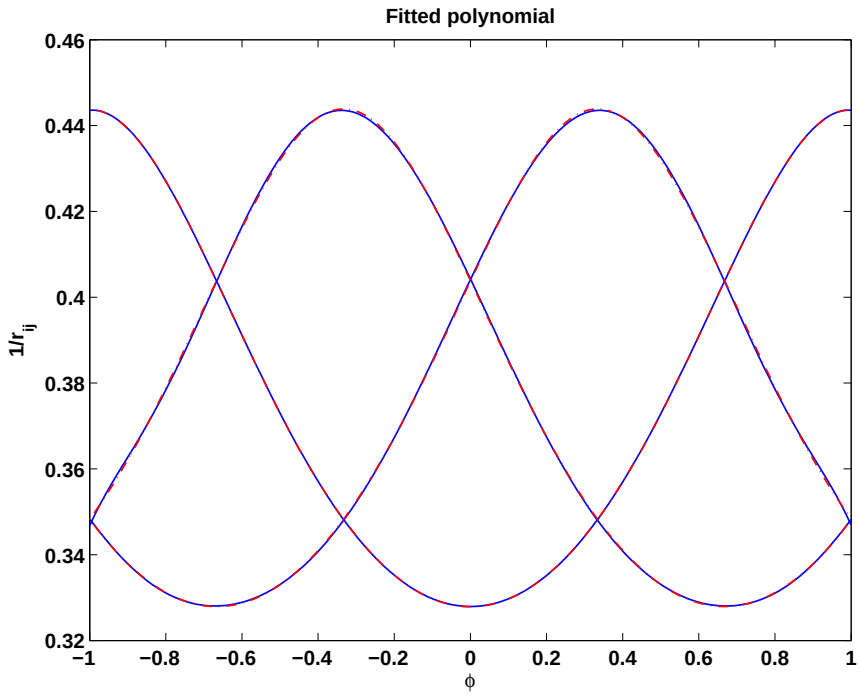


Figure 3.14: Plots of ϕ vs. $1/r_{ij}$ for the hydrogens in ethane (blue lines) and fitted 8^{th} degree polynomials (dashed red lines). Notice that the abscissa has been rescaled from $[-\pi\ \pi]$ to $[-1.0\ 1.0]$.

The only parameters required are the charges themselves and the constant of proportionality $\mathcal{C}$ in equation 3.14. In the Amber94 force field paper the charges are presented as they appear on the various amino acid residues. For the charges on the hydrogens in ethane it was decided to use the charges on the methyl hydrogens on the side-chain of the alanine residue since this residue is overall charge neutral and is also non-polar. Therefore the charge used for the hydrogens was $q_i = 0.0603 \times e$ C where $e = 1.6 \times 10^{-19}$ C is the charge on a proton. The constant of proportionality $\mathcal{C}$ is defined as $\mathcal{C} = 1/4\pi\epsilon_0$ where ϵ_0 is the permittivity of a vacuum (8.85×10^{-12} farad m^{-1}). In analogy to the van der Waals terms, in all Amber94 calculations, a scaling factor of 1/1.2 must be applied to all 1-4 electrostatic interactions. The Coulombic contributions to the nine hydrogen interactions are plotted as a function of the torsional angle ϕ in figure 3.16.

Abscissa rescaling

A key to obtaining a numerically successful sum of squares decomposition is that the coefficients of the polynomial be well suited to producing an SDP which is solvable in a numerically stable way by interior point methods. For this to be the case it has been observed that it is preferable if the ratio of the magnitudes of the largest and smallest coefficients in the potential energy functional is not too great. Additionally it is desirable to have the coefficients as close to unity as possible.

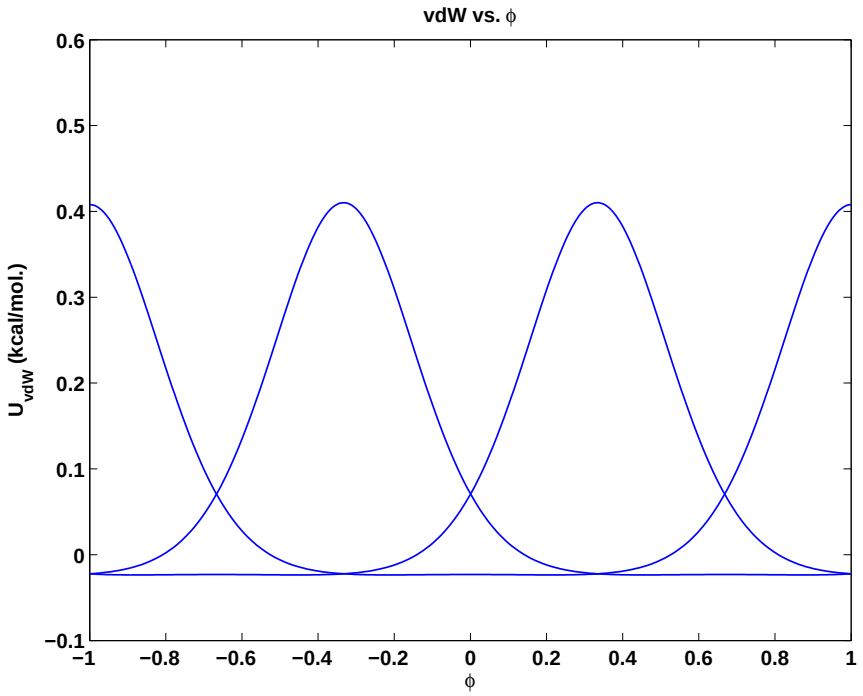


Figure 3.15: van der Waals energies from the nine hydrogen interactions as a function of the torsional angle ϕ. These curves are plotted using the fitted $1/r_{ij}$ vs. ϕ functionals. Note that as a result of the symmetry in the interatomic distances only 3 energy curves are shown. Each curve contains the sum of the contributions of the 3 symmetrical interactions. These energies also include the rescaling of the 1-4 interactions.

In this work this is referred to as coefficient conditioning. Rescaling the abscissa can result in much improved coefficient conditioning and this approach is employed where possible. This has no effect on the final result and the solution can be rescaled back to the original units. Examples are shown in figures 3.12 and 3.14. Table 3.2 gives the coefficients of polynomial fits to the torsional potential energy functional for the ranges $[0, 2\pi]$, $[-\pi, \pi]$, $[-1.0, 1.0]$ and $[0.0, 1.0]$. Similarly table 3.3 gives the coefficients for the polynomial fits to the torsion angle - inverse distance $(\phi - 1/r_{ij})$ functionals for the same ranges.

In the case of the non-bonded interactions the polynomial expressions given in equations 3.19, 3.20 and 3.21 must be substituted into the Lennard-Jones and Coulombic functionals. Since the Lennard-Jones functional is of 12^{th} degree in $1/r_{ij}$ care must be taken to ensure that the coefficients in these equations are as close in magnitude to unity as possible. The distribution of the $\log_{10}$ of the absolute of the coefficients was examined for the fitted polynomials for the four ranges $[0, 2\pi]$, $[-\pi, \pi]$, $[-1.0, 1.0]$, $[0.0, 1.0]$. The means and standard deviations for the $\log_{10}$ of the coefficients in the torsional energy functional and torsion angle - inverse distance functionals are listed in table 3.4 for the four ranges. The range $[-1.0, 1.0]$ has the closest distribution to zero so this range was chosen for the SOS model.

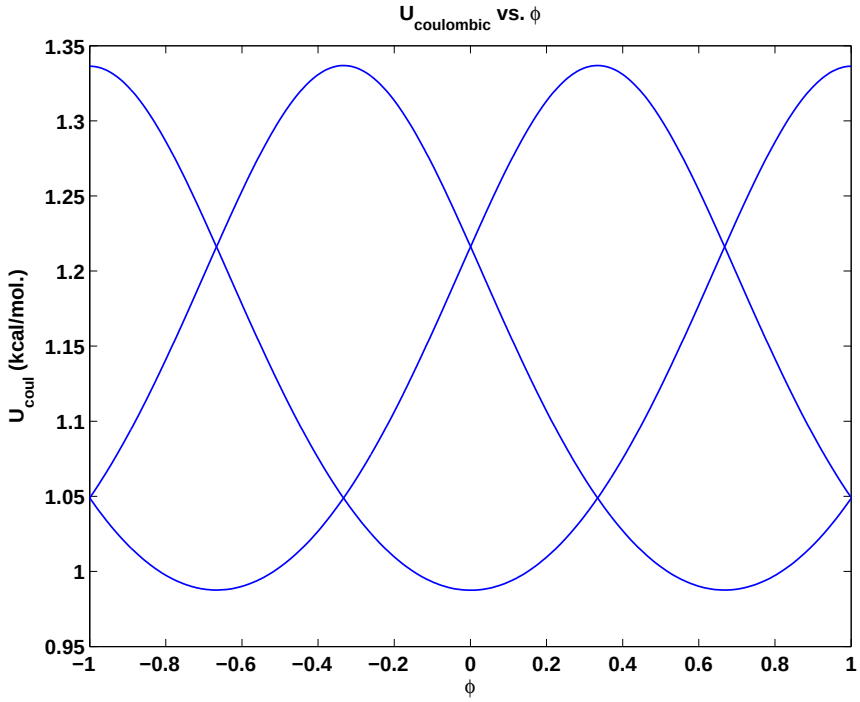


Figure 3.16: Coulombic energies from the nine hydrogen interactions as a function of the torsional angle ϕ. These curves are plotted using the fitted $1/r_{ij}$ vs. ϕ functionals. Note that as a result of the symmetry in the interatomic distances only 3 energy curves are shown. Each curve contains the sum of the contributions of the 3 symmetrical interactions. These energies also include the rescaling of the 1-4 interactions.

In spite of the rescaling the coefficients remain relatively small and this results in the van der Waals terms making a very small contribution to the overall potential energy. This is to be expected since the r_{ij} are never less than 2.25 Angstroms.

3.6 SOS programme for ethane model

The van der Waals and Coulombic components of the potential energy functional (in polynomial form) can be summed to form a polynomial function expressing the total non-bonded potential energy in terms of a variable $1/r_{ij}$ and the $\phi - 1/r_{ij}$ functions can then be used to express the potential energy as a function of the torsion angle. This can be done in two ways, i) by direct substitution of the $\phi - 1/r_{ij}$ functions into the potential energy functionals or ii) by defining 3 new variables r_1, r_2, r_3 and setting them equal to the $\phi - 1/r_{ij}$ polynomials using equality constraints. Both approaches have advantages and disadvantages. The first approach keeps the number of constraints which must be incorporated into the sum of squares representation to a minimum and also results in a relatively modest number of components in the monomial basis (49), but results in high degrees in the monomial basis (48). The resulting SDP is therefore relatively small. The second approach results in a

range	$[0,2\pi]$	$[-\pi,\pi]$	$[-1.0,1.0]$	$[0.0,1.0]$
x^{14}	8.5180×10^{-06}	8.5180×10^{-06}	$7.7153\times10^{+01}$	$1.2641\times10^{+06}$
x^{13}	-3.7462×10^{-04}	2.0780×10^{-08}	-4.8792×10^{-01}	$-8.8525\times10^{+06}$
x^{12}	7.2317×10^{-03}	-4.1774×10^{-04}	$-3.8376\times10^{+02}$	$2.7212\times10^{+07}$
x^{11}	-8.0373×10^{-02}	-5.9366×10^{-07}	$2.1621\times10^{+00}$	$-4.8158\times10^{+07}$
x^{10}	5.6715×10^{-01}	8.8662×10^{-03}	$8.2609\times10^{+02}$	$5.4112\times10^{+07}$
x^{9}	-2.6464×10^{-00}	6.4578×10^{-06}	$-3.9985\times10^{+00}$	$-4.0206\times10^{+07}$
x^{8}	8.2785×10^{-00}	-1.0356×10^{-01}	$-9.7869\times10^{+02}$	$2.0028\times10^{+07}$
x^{7}	-1.7417×10^{-01}	-3.3319×10^{-05}	$3.8710\times10^{+00}$	$-6.7095\times10^{+06}$
x^{6}	2.5018×10^{-01}	6.8519×10^{-01}	$6.5672\times10^{+02}$	$1.5346\times10^{+06}$
x^{5}	-2.5839×10^{-01}	8.2372×10^{-05}	$-1.9735\times10^{+00}$	$-2.5239\times10^{+05}$
x^{4}	1.9330×10^{-01}	-2.3358×10^{-00}	$-2.2707\times10^{+02}$	$3.0066\times10^{+04}$
x^{3}	-6.7994×10^{-00}	-8.5744×10^{-05}	4.5804×10^{-01}	$-1.6840\times10^{+03}$
x^{2}	-1.6764×10^{-00}	3.1384×10^{-00}	$3.0943\times10^{+01}$	$-6.6114\times10^{+01}$
x	-1.3641×10^{-01}	2.4939×10^{-05}	-3.1312×10^{-02}	-8.5668×10^{-01}
1	1.4029×10^{-00}	8.2157×10^{-04}	8.2950×10^{-04}	$1.4029\times10^{+00}$

Table 3.2: Coefficients of polynomials which were fit to the X-CT-CT-X torsional functional when graphed on 4 different abscissas, namely $[0, 2\pi]$, $[-\pi, \pi]$, $[-1.0, 1.0]$, $[0.0, 1.0]$. The fit was a least squares fit and the norm of the residual was 1.6816×10^{-02} in each case.

longer monomial basis (495 components when unoptimised, 310 when optimised using CDD) but lower monomial degrees (16) in the basis. The first approach produces a relatively small SDP to solve but this does not converge very well due to the poor coefficient conditioning which results from the multiplicative effect of incorporating the P_1, P_2 and P_3 polynomials into the Lennard-Jones functionals. On the other hand the second approach leads to a much larger SDP but gives better convergence. This second approach will be described here.

3.6.1 Definition of the potential energy functional

The total potential energy function is then defined as a sum of the three components

$$U_{\text{total}} = U_{\text{Coulombic}} + U_{\text{van der Waals}} + U_{\text{torsion}}, \tag{3.23}$$

where $U_{\text{Coulombic}}$ and $U_{\text{van der Waals}}$ are functions of the variables r_1, r_2, r_3 and U_{torsion} is a function of ϕ. This is a polynomial function in 4 variables of degree 16. Particular care must be taken with the representation of polynomials functions in Matlab. When the symbolic toolbox is used to represent polynomials, Matlab stores the coefficients as rational numbers which can specify the coefficients to very high accuracy. It can considerably speed up the preprocessing phase of a sum of squares programme if these rationals are converted to double floating point numbers using the `vpa` (variable precision arithmetic) function in the symbolic toolbox. However this approximation can lead to significant errors in the representation of the energy surface and must be done with care. It is useful to always re-plot functions before

range	x^8	x^7	x^6	x^5	x^4	x^3	x^2	x	1
[0,2π]									
P_1	4.0994×10^{-06}	-1.0304×10^{-04}	9.9558×10^{-04}	-4.5278×10^{-03}	7.8911×10^{-03}	1.1370×10^{-02}	-5.5098×10^{-02}	2.4552×10^{-03}	4.4357×10^{-01}
[0,2π]									
P_2	-1.4948×10^{-05}	3.2712×10^{-04}	-2.7342×10^{-03}	1.1015×10^{-02}	-2.3051×10^{-02}	2.4643×10^{-02}	6.6083×10^{-03}	-3.7015×10^{-02}	3.4829×10^{-01}
[0,2π]									
$P3$	-1.4871×10^{-05}	4.2221×10^{-04}	-4.8479×10^{-03}	2.8377×10^{-02}	-8.6566×10^{-02}	1.2324×10^{-01}	-6.9507×10^{-02}	6.2822×10^{-02}	3.4679×10^{-01}
[$-\pi$,π]									
P_1	4.0994×10^{-06}	-1.1017×10^{-08}	-1.3753×10^{-04}	1.5301×10^{-07}	2.8977×10^{-04}	-5.8375×10^{-07}	1.8312×10^{-02}	5.2641×10^{-07}	3.2793×10^{-01}
[$-\pi$,π]									
P_2	-1.4948×10^{-05}	-4.8571×10^{-05}	3.2857×10^{-04}	1.3197×10^{-03}	-1.7469×10^{-03}	-1.4383×10^{-02}	-6.1364×10^{-03}	5.9853×10^{-02}	4.0420×10^{-01}
[$-\pi$,π]									
P_3	-1.4871×10^{-05}	4.8455×10^{-05}	3.2730×10^{-04}	-1.3181×10^{-03}	-1.7407×10^{-03}	1.4377×10^{-02}	-6.1460×10^{-03}	-5.9847×10^{-02}	4.0420×10^{-01}
[-1,1]									
P_1	3.8740×10^{-02}	-1.9035×10^{-04}	-1.3182×10^{-01}	4.4786×10^{-04}	2.8169×10^{-02}	-7.5224×10^{-05}	1.8055×10^{-01}	-1.8150×10^{-04}	3.2793×10^{-01}
[-1 1]									
P_2	-1.4126×10^{-01}	-1.4561×10^{-01}	3.1544×10^{-01}	4.0187×10^{-01}	-1.7084×10^{-01}	-4.4495×10^{-01}	-5.9825×10^{-02}	1.8800×10^{-01}	4.0410×10^{-01}
[-1,1]									
P_3	-1.4053×10^{-01}	1.4640×10^{-01}	3.1319×10^{-01}	-4.0329×10^{-01}	-1.6819×10^{-01}	4.4545×10^{-01}	-6.1274×10^{-02}	-1.8786×10^{-01}	4.0429×10^{-01}
[0 1]									
P_1	$9.9174\times10^{+00}$	$-3.9694\times10^{+01}$	$6.1071\times10^{+01}$	$-4.4226\times10^{+01}$	$1.2274\times10^{+01}$	$2.8159\times10^{+00}$	$-2.1730\times10^{+00}$	1.5419×10^{-02}	4.4357×10^{-01}
[0,1]									
P_2	$-3.6163\times10^{+01}$	$1.2602\times10^{+02}$	$-1.6772\times10^{+02}$	$1.0759\times10^{+02}$	$-3.5854\times10^{+01}$	$6.1035\times10^{+00}$	2.6062×10^{-01}	-2.3245×10^{-01}	3.4829×10^{-01}
[0,1]									
P_3	$-3.5976\times10^{+01}$	$1.6265\times10^{+02}$	$-2.9738\times10^{+02}$	$2.7718\times10^{+02}$	$-1.3464\times10^{+02}$	$3.0523\times10^{+01}$	$-2.7412\times10^{+00}$	3.9452×10^{-01}	3.4679×10^{-01}

Table 3.3: Coefficients of fitted polynomials to the ϕ vs. $1/r_{ij}$ curves when graphed on 4 different abscissas namely [0 2π], [$-\pi$ π], [-1.0 1.0], [0.0 1.0]. The labels P_1,P_2 and P_3 refer to the curves representing the identical distances {HA1-HB1, HA2-HB3, HA3-HB2}, {HA1-HB2, HA2-HB1, HA3-HB3}, {HA1-HB3, HA2-HB2, HA3-HB1}. The fit was a least squares fit and the norm of the residuals were 1.0831×10^{-03}, 8.0801×10^{-03} and 8.1387×10^{-03} for P_1, P_2 and P_3 respectively in each of the four abscissa ranges.

Range	$[0,2\pi]$	$[-\pi,\pi]$	$[-1.0,1.0]$	$[0.0,1.0]$
Mean	$-1.5456\times10^{+00}$	$-3.3416\times10^{+00}$	-4.3943×10^{-01}	$2.5012\times10^{+00}$
Std. dev.	$1.8089\times10^{+00}$	$2.1518\times10^{+00}$	$1.7082\times10^{+00}$	$2.7803\times10^{+00}$

Table 3.4: Mean and standard deviations of the distributions of $\log_{10}$ of the coefficients from the polynomials fitted to the torsion potential and torsion angle - inverse distance functionals.

and after such an operation. It has been observed that rounding errors due to the conversion affects computations in models such as this one.

3.6.2 Adding constraints

Two types of constraints are required in this model. The first are box constraints on the values that ϕ can take limiting it to the range $[-1.0, 1.0]$ in the rescaled case. These are inequality constraints and therefore must be multiplied by sum of squares variables. A sum of squares representation in the form of a Positivstellensatz representation (see section 2.4.4) is suitable:

$$U_{\text{total}} - g_1 s_1 - g_2 s_2 \quad \text{is SOS}, \tag{3.24}$$

where $g_1 := \phi \geq -1.0$ and $g_2 := \phi \leq 1.0$ are the inequality constraints. s_1 and s_2 are sum of squares polynomial multipliers of degree 6.

Another set of constraints act to constrain the variables r_1, r_2, r_3 to be equal to their respective polynomials P_1, P_2, P_3. The following equality constraints are defined:

$$h_1 := r_1 - P_1, \qquad h_2 := r_2 - P_2, \qquad h_3 := r_3 - P_3. \tag{3.25}$$

To incorporate these constraints the SOS representation is augmented with equality constraints multiplied by polynomial variables. These polynomial multipliers p_1, p_2, p_3, must be at least defined in the variables of their respective constraint and be at least of degree 8 (the degree of the constraints). The Positivstellensatz representation is therefore

$$U_{\text{total}} - g_1 s_1 - g_2 s_2 - p_1 h_1 - p_2 h_2 - p_3 h_3 \quad \text{is SOS}. \tag{3.26}$$

It is now possible to write the SOS programme as:

$$\begin{aligned} \text{maximise:} \quad & \lambda, \\ \text{subject to:} \quad & U_{\text{total}} - g_1 s_1 - g_2 s_2 - p_1 h_1 - p_2 h_2 - p_3 h_3 - \lambda \quad \succeq 0. \end{aligned} \tag{3.27}$$

3.6.3 Running the SDP

The SOS optimisation problem was generated using SOSTOOLS and SeDuMi was used to solve the SDP. The following is the output from the SeDuMi solver demonstrating the convergence of the solution. The interested reader is referred to the SeDuMi user manual for a detailed interpretation of the output.

```
SeDuMi 1.05 by Jos F. Sturm, 1998, 2001.
Alg = 2: xz-corrector, Step-Differentiation, theta = 0.250, beta = 0.500
eqs m = 2855, order n = 644, dim = 93324, blocks = 4
nnz(A) = 50673 + 0, nnz(ADA) = 8151025, nnz(L) = 4076940
 it :     b*y       gap    delta  rate   t/tP*  t/tD*   feas cg cg
  0 :            3.50E+01 0.000
  1 :  -3.23E+03 1.37E+01 0.000 0.3911 0.9000 0.9106   1.10  1  1
  2 :  -3.23E+03 3.48E+00 0.000 0.2542 0.9000 0.0000   2.48  1  1
  3 :   1.42E+01 6.29E-07 0.000 0.0000 0.5991 0.9000   1.84  1  1
  4 :  -7.18E+01 2.19E-07 0.000 0.3478 0.7875 0.9000   2.07  1  1
  5 :  -1.44E+01 6.36E-08 0.000 0.2908 0.8416 0.9000   1.66  1  1
  6 :  -4.30E+00 1.64E-08 0.000 0.2579 0.9000 0.9000   1.29  1  1
  7 :  -3.27E+00 4.32E-09 0.000 0.2635 0.9000 0.9000   1.11  1  1
  8 :  -3.56E+00 1.07E-09 0.000 0.2474 0.9000 0.9000   1.09  1  1
  9 :  -3.67E+00 2.95E-10 0.000 0.2763 0.9000 0.7831   1.12  1  1
 10 :  -3.56E+00 5.59E-11 0.000 0.1892 0.9000 0.8911   1.12  1  1
 11 :  -3.54E+00 1.04E-11 0.000 0.1865 0.9000 0.9000   1.04  2  2
 12 :  -3.54E+00 2.34E-12 0.000 0.2248 0.9000 0.9000   0.96  2  2
 13 :  -3.54E+00 2.05E-13 0.000 0.0877 0.9900 0.9900   0.96  3  3
 14 :  -3.54E+00 1.61E-14 0.000 0.0785 0.9900 0.9900   1.01  6  6
 15 :  -3.54E+00 3.14E-15 0.000 0.1952 0.9000 0.9000   1.01  9  9
iter seconds digits       c*x               b*y
 15    418.0   Inf -3.5378075683e+00 -3.5378072294e+00
|Ax-b| =   1.8e-05, [Ay-c]_+ =   2.3E-10, |x|=  3.9e+04, |y|=  1.7e+00
Max-norms: ||b||=1.127412e+04, ||c|| = 1,
Cholesky |add|=0, |skip| = 1, ||L.L|| = 3.5651.

Residual norm: 1.8431e-05

       cpusec: 4.180300000000000e+02
         iter: 15
    feasratio: 1.01235941469921
         pinf: 0
         dinf: 0
       numerr: 0
```

The initial monomial basis (without the SOS constraints) had 495 components implying that there would be $495^2 = 245025$ variables in the SOS programme (SeDuMi does not take the symmetry of the matrix into account at this point). However once advantage was taken of the Newton polytope method described in section 2.3.4 this was reduced to a monomial basis of 301 components. In addition each SOS polyno-

mial added an additional 70 variables to the basis. Therefore the number of positive semidefinite variables in the SDP was

$$70^2 + 70^2 + 301^2 = 100401, \tag{3.28}$$

since SeDuMi can handle sparse matrices (see figure 3.17). Each of the three polynomial variables had a monomial basis of length 45 adding a total of 135 variables. These are free variables meaning that they can take on negative values and are not included in the positive semidefinite constraints. λ is also a free variable bringing the total number of free variables to 136. Therefore the total number of variables in the SDP that was submitted to SeDuMi contained 100537 variables. Among the 100401 positive semidefinite variables there were 2855 constraints. This is not a particularly small SDP and the solution time was 419.5 cpu seconds. Some important points from the output to note are the fact that the SDP is both primal and dual feasible and that the feasibility ratio is almost 1.00. (The feasibility ratio is an indicator of the quality of the self-dual embedding which was achieved by the solver).

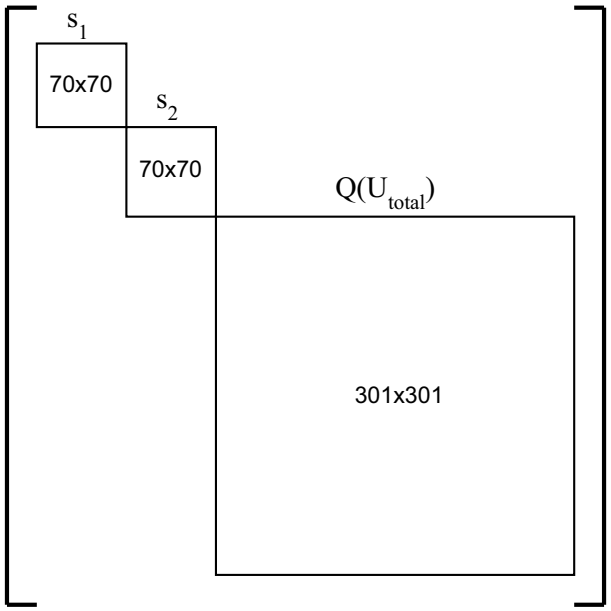


Figure 3.17: Block diagonal positive semidefinite SOS matrix showing the two blocks for the SOS constraints and the main block for the potential.

3.6.4 Post-processing and results

The lower bound reported by SOSTOOLS was 3.5378 kcal/mol. and the duality gap was -3.3891×10^{-7} kcal/mol. The duality gap is small enough to indicate that the SDP is very well converged and this level of convergence should give a very good lower bound on the global minimum if a sufficient Positivstellensatz representation has been used.

minimum	minimiser	function value
1	-6.6729e-01	3.537836
2	4.9709e-04	3.544585
3	6.6830e-01	3.537804

Table 3.5: Local minima of total potential energy function obtained using a golden section method. In the Amber94 force field these minima would be identical and symmetric about the origin. However errors were introduced by the fitting process and in particular during the fitting of the torsional potential.

In this SOS programme there are 4 variables and the solution vector obtained from the 85^{th}, 29^{th}, 8^{th} and 2^{nd} entries of the 1^{st} row of the dual solution matrix were

$$[x,\ r_1,\ r_2,\ r_3] \quad = \quad [0.61141,\ 0.40389,\ 0.40030,\ 0.33127]. \tag{3.29}$$

Therefore the minimiser is predicted to be 0.61141 (in rescaled units). The dual matrix is not exactly rank 1 as indicated by an eigenvalue decomposition where the first seven entries were

$$\begin{bmatrix} 2.7814 \times 10^{+00} \\ 0.1041 \times 10^{-01} \\ 2.4555 \times 10^{-05} \\ 1.8450 \times 10^{-06} \\ 1.4802 \times 10^{-06} \\ 6.7454 \times 10^{-08} \\ 5.7728 \times 10^{-08} \end{bmatrix}. \tag{3.30}$$

This indicates that there is a small error in the minimiser location and casts some uncertainty over the lower bound of 3.5378 obtained from the primal SDP solution. Evaluating the total energy function at the minimiser vector (equation 3.29) gives a value of 3.6195. Therefore what can be claimed with certainty is that the global minimum lies in the range $[3.5378, 3.6195]$ and that the location of the global minimum is close to $\phi = 0.61141$.

To confirm these predictions local minimisations of the total energy functional were carried out in the three minima to obtain the exact values. These were carried out using the `fminbnd` function in the Matlab optimisation toolbox, which is based on a simple golden section method and quickly converged to the numerically exact minima. These are presented in table 3.5. Notice that in the true Amber94 potential these minima would be identical in magnitude and would be symmetrical about the origin. However the fitting process has introduced small errors which derive mainly from the fitting of the torsional potential. As a result the energies of the three

minima are not identical with the 3^{rd} minimum being the lowest. Notice that the lower bound produced by SOSTOOLS is in fact accurate to 4 decimal places. In addition it has identified the correct well but not the exact location of the minimum due to the issue with the dual solution matrix being of rank greater than 1. The absolute error in the minimiser is 0.05689 which when normalised by the range of the domain gives a 2.8% error.

3.6.5 Improving the SOS representation

Increasing the degree of the SOS representation should improve the rank of the dual matrix, and to a large extent finding the appropriate level of SOS representation can be trial and error. It involves a trade off between accuracy and computational expense. The dual matrix is quite close to rank one and therefore a small increase in the level of the SOS representation may give a better solution. The original Positivstellensatz representation has SOS multipliers of degree 6, so another attempt was made with SOS variables of degree 8, in common with the polynomial multipliers.

The SDP solution took 422.9 seconds and the feasibility ratio was improved to 1.00155. The duality gap was improved to -8.7629×10^{-08}. The lower bound was identical at 3.537800. The rank of the dual matrix had considerably improved with the first seven entries of the eigenvalue decomposition being

$$\begin{bmatrix} 2.8780 \times 10^{+00} \\ 7.6707 \times 10^{-03} \\ 6.4439 \times 10^{-06} \\ 5.0678 \times 10^{-07} \\ 3.9065 \times 10^{-07} \\ 1.3871 \times 10^{-08} \\ 1.0819 \times 10^{-08} \end{bmatrix}. \tag{3.31}$$

The minimiser point was now much closer to the global minimum

$$[x,\ r_1,\ r_2,\ r_3] \quad = \quad [0.66411,\ 0.40389,\ 0.40328,\ 0.32829]. \tag{3.32}$$

i.e. the error is 4.19×10^{-3} which when normalised by the range is 0.21%. Evaluating the total energy function at the minimiser yields 3.5379 which implies that the global minimum is identified to within 0.0001 kcal/mol.

3.7 A trigonometric basis

The approach described thus far centred on mapping the expression for the potential energy of ethane into a function of the torsion angle ϕ between the two methyl

groups. However the potential energy functional can also be written in terms of the cosine and sine components of the torsion angle. Trigonometric functions can be represented in polynomial form by using the cosine and sine functions as a basis in a similar fashion to a Fourier series. In a sum of squares decomposition the basis for the SOS function contains variables representing the cosine and sine functions, and an equality constraint is added such that the variables meet the constraint

$$(\cos\phi)^2 + (sin\phi)^2 = 1. \tag{3.33}$$

3.7.1 SOS decomposition for the torsional potential in a trigonometric basis

The torsional component of the potential energy which is itself written as a Fourier expansion

$$U_{\text{torsion}} = \frac{V_n}{2}(1 + cos(3\phi)), \tag{3.34}$$

can be expressed in terms of variables c_3 $(cos3\phi)$ and s_3 $(sin3\phi)$ where again an equality constraint $h_1 := c_3^2 + s_3^2 = 1$ is used. A sum of squares decomposition was obtained for the torsional potential alone using the Positivstellensatz representation

$$U_{\text{torsion}} - h_1 p_1 \tag{3.35}$$

where p_1 is a polynomial of degree 4. The solution was obtined in 0.93 seconds with a feasibility ratio of 0.9975. The lower bound produced was 0.00 which is exact and the minimiser was obtained at $c_3 = -1, s_3 = 0$. These trigonometric values correspond to torsional angles $\phi = \{\frac{\pi}{3}, \frac{2\pi}{3}, \pi, \ldots\}$ which are the exact minimisers of the torsional component of the ethane potential. Notice how the symmetry problems encountered with the torsional potential in the previous approach are eliminated with this approach. In addition the need for rescaling of the abscissa is eliminated.

3.7.2 SOS decomposition for the remainder of the potential using a trigonometric basis

Expressing the entire potential in a trigonometric basis necessitated expressing the inverse distance between the hydrogens as a funcion of $cos\phi$ and $sin\phi$ in order to model the non-bonded terms, i.e. the dispersion and Coulombic terms. Once again

this was carried out using a fitting procedure. The fitted expressions are as follows:

$$f_1 := .49587 \times 10^{-5}c^8 + .17588 \times 10^{-4}c^7 + .53636 \times 10^{-4}c^6 + .19977 \times 10^{-3}c^5 + .76276 \times 10^{-3}c^4 + .29642 \times 10^{-2}c^3 + .12095 \times 10^{-1}c^2 + .54848 \times 10^{-1}c + .37307. \tag{3.36}$$

$$f_2 := -.10817s^8 + .15632s^6 - .79413 \times 10^{-1}s^4 - .35235 \times 10^{-1}s^2 + .44378. \tag{3.37}$$

$$f_3 := .98107 \times 10^{-1}c^8 - .14701 \times 10^{-1}c^7 - .14167c^6 + .16201 \times 10^{-1}c^5 + .73959 \times 10^{-1}c^4 - .92989 \times 10^{-2}c^3 + .86558 \times 10^{-2}c^2 - .19000 \times 10^{-1}c + .33329. \tag{3.38}$$

$$f_4 := .62466 \times 10^{-1}s^8 - .16113 \times 10^{-1}s^7 - .89763 \times 10^{-1}s^6 + .17370 \times 10^{-1}s^5 + .47735 \times 10^{-1}s^4 - .11598 \times 10^{-1}s^3 + .11286 \times 10^{-1}s^2 - .38086 \times 10^{-1}s + .34847. \tag{3.39}$$

$$f_5 := -.98040 \times 10^{-1}c^8 + .14652 \times 10^{-1}c^7 + .14170c^6 - .16364 \times 10^{-1}c^5 - .74949 \times 10^{-1}c^4 + .16010 \times 10^{-1}c^3 - .20627 \times 10^{-1}c^2 - .43080 \times 10^{-1}c + .43178. \tag{3.40}$$

$$f_6 := .62355 \times 10^{-1}s^8 + .16078 \times 10^{-1}s^7 - .89603 \times 10^{-1}s^6 - .17332 \times 10^{-1}s^5 + .47650 \times 10^{-1}s^4 + .11576 \times 10^{-1}s^3 + .11279 \times 10^{-1}s^2 + .38074 \times 10^{-1}s + .34843. \tag{3.41}$$

where c and s are variables representing cos and sin respectively and monomials with coefficients less than 1×10^{-10} have been omitted. The following equality constraints were defined equating these polynomials in c and s to the inverse distances $\{r_{i1}, r_{i2}, r_{i3}\}$. The following equality constraints were defined equating these

polynomials in c and s to the inverse distances $\{r_{i1}, r_{i2}, r_{i3}\}$,

$$\begin{aligned} h_1 &:= f_1 - r_{i1} = 0, \\ h_2 &:= f_2 - r_{i2} = 0, \\ h_3 &:= f_3 - r_{i3} = 0, \\ h_4 &:= f_4 - r_{i1} = 0, \\ h_5 &:= f_5 - r_{i2} = 0, \\ h_6 &:= f_6 - r_{i3} = 0. \end{aligned} \tag{3.42}$$

Inequality constraints were defined such that both variables were constrained to the range $[-1, 1]$ as follows:

$$\begin{aligned} g_1 &:= c + 1 \geq 0, \\ g_2 &:= 1 - c \geq 0, \\ g_3 &:= s + 1 \geq 0, \\ g_4 &:= 1 - s \geq 0. \end{aligned}$$

A further equality contraint,

$$h_7 := c^2 + s^2 - 1 = 0, \tag{3.43}$$

was included to maintain the trigonometric relationship between c and s. The total energy U_{tot} is defined as

$$U_{\text{total}} = U_{\text{Coulombic}} + U_{\text{van der Waals}}, \tag{3.44}$$

where the torsional component of the energy functional has been neglected for the moment. Using this as the objective function and the above constraints a Positivstellensatz representation was defined as follows

$$\begin{aligned} U_{\text{tot}} - g_1 s_1 - g_2 s_2 - g_3 s_3 - g_4 s_4 \\ -h_1 p_1 - h_2 p_2 - h_3 p_3 - h_4 p_4 - h_5 p_5 - h_6 p_6 - h_7 p_7 \text{ is SOS}, \end{aligned} \tag{3.45}$$

where $s_1, \ldots, s_4$ are SOS polynomial variables and $p_1, \ldots, p_7$ are polynomial variables.

Once again SOSTOOLS was used to define the SOS optimisation programme and convert it to an SDP. The SDP was submitted to SeDuMi on the same PC, however the entire problem could not be held in memory and the solver failed to produce any solution. As a result no attempt was made to incorporate the torsional

component of the energy function.

3.8 Discussion

At this stage a case can be made that all four objectives stated at the start of this chapter have been achieved. In respect of objective 1, it has been demonstrated using the sum of squares decompositions of the model potential energy surfaces, that potential energy surfaces which exhibit pathological properties from a sampling perspective can be bounded from below. The method is particularly useful for the following cases i) where minima are close in energy but separated by other deep minima in configurational space, ii) where there are near degenerate global minima, iii) where there is some deep but very localised well on an otherwise seemingly flat energy landscape. Another important practical point which was demonstrated using the model potentials is that accurate polynomial representations of potential energy landscapes derived from complicated functionals containing non-polynomial terms including exponentials can be used to obtain an accurate SOS decomposition and therefore accurately identify the global minimum.

The ethane case study has demonstrated that the standard functional forms which are used in modern classical force fields can be represented using polynomial functionals (objective 2) provided sufficiently high degree polynomials are used. In particular the ability to represent potential energy functionals for organic compounds was demonstrated using the Amber94 force field as a case study. For the stretch and bend terms the equivalence is trivial. The torsional functionals require a polynomial fitting to be carried out. Some strategies for ensuring that the fitted polynomials are best suited to obtaining a high quality sum of squares decomposition were also discussed. These include rescaling the range from $[0\ 2\pi]$ to $[-1.0\ 1.0]$ and minimising the ratio of the largest and smallest coefficients. The van der Waals and Coulombic terms are rational polynomials but can be dealt with by expressing the interatomic distance relationships in terms of $1/r_{ij}$. The resulting expressions are polynomials in $1/r_{ij}$. The interatomic distance relationships may then be either directly incorporated into the energy functional or equality constraints may used as described in section 3.6.2. The SOS decomposition obtained using this approach yielded an excellent lower bound and minimiser for the global minimum of ethane as represented in the Amber94 forcefield (objective 3).

The approach of using trigonometric variables is a particularly elegant way to deal with angle dependent potential energy functionals. however the additional cost of requiring an extra variable in the basis for each angle made the approach too expensive even for this low dimensional case.

The fourth objective of demonstrating the application of the SOS decomposition

approach to a potential energy surface of a real chemical system was also met using the ethane in vacuo model. Although ethane is a small molecule by any standards and is amenable to optimisation by exhaustive search approaches even using ab initio methodologies to calculate the interatomic forces, the fact that its classical potential energy surface can be accurately represented using polynomial functionals is important for future research in this direction. Of course an exhaustive search method is more efficient than computing a SOS decomposition in such a small system, yet a fundamental principle of complexity theory is that a polynomial time approach such as SOS optimisation will always be more efficient than an exponential time approach such as exhaustive search for a large enough problem. It is the size of the problem at which this transition is made that is important (see figure 1.9). The practical issue that large enough problems cannot yet be dealt with using interior point methods however must also be acknowledged.

3.8.1 Lifting

Of course there are a number of other serious issues that remain. The most significant is the massive lifting that is required to obtain an accurate lower bound and minimiser coordinates using the SOS decomposition approach. Ethane contains 8 atoms each requiring 3 coordinates to completely describe its position in configurational space. Therefore it has $24 - 6 = 18$ independent degrees of freedom (where the 6 degrees of freedom account for translational and rotational invariance of the potential energy surface). Notwithstanding the fact that the bond angles and stretches are constrained the resulting SDP contains 100537 variables and 2885 constraints. In spite of the advantage of moving from a highly non-convex problem to a convex problem the number of variables in the problem (given the size of the physical problem) is truly enormous. The largest SDP which has been observed by the author in the literature to date contains 127600 constraints [222]. However in a great many cases, long before this limitation is reached, problems are experienced in the preprocessing stage when the Maple kernel crashes due to lack of memory. This problem was already experienced when applying the same methodology to propane (an 11 atom hydrocarbon shown in figure 3.18). A similar problem was encountered when the multivariate polynomial toolbox was used instead of the symbolic toolbox (`Maximum variable size allowed by the program is exceeded.`). The memory limitations while manipulating polynomials could be improved by writing a C polynomial manipulation package.

The primary reason that the interatomic distances were fitted to a functional in ϕ is that introducing Cartesian coordinates and specifying the interatomic distances as norm constraints results in an SDP which is too large to solve on a 3GHz Pentium 4 with 3Gb RAM using SeDuMi or SDPT3. The current approach however, would

not be ideal for larger systems, since the $\phi - 1/r_{ij}$ functional rapidly becomes high dimensional and the fitting process becomes onerous. Computational expense was also a consideration when deciding to fix the angles and bond lengths but little would be gained in the sophistication of the model by allowing them to vary, since these are quadratic functions.

3.8.2 Polynomial fitting

Another practical issue is the effect of fitting polynomials to predefined energy functionals. As was seen with the torsional functional in the ethane example, a compromise must be made between accuracy of the fit and degree of the polynomial (and therefore size of the monomial vector). In this case the true potential has 3 identical minima but the polynomial fit resulted in a functional with a single minimum albeit by a tiny difference to the other two. Nevertheless the SOS decomposition was able to detect this slight difference in minima and correctly identified a single minimum. This would not be a problem if the original data used to define the force field was fitted to polynomial functionals however SOS decomposition approaches would have to prove themselves for much larger systems for the large amount of work involved in this to be justified.

One possible approach to reducing the number of variables in the SDP may be to take a closer look at the monomial vector $\mathbf{z}$. At present when the monomial is constructed combinations of all the variables in the problem are included. Any methods for reducing the size of the basis such as the Newton polytope method are performed using this full basis. Because the potential energy functional is a sum of individual interactions it may be possible to neglect certain combinations of variables when constructing the monomial basis since interactions in the all to all sense are not always considered due to cutoffs.

3.8.3 Interior point methods

Although interior point methods for linear programmes as implemented in for example the CPLEX software package [223] can routinely solve linear programmes with hundreds of thousands of variables and constraints, software for solving SDP problems have not *yet* had the same level of success in solving large scale problems. As described in appendix B, the main differences between the different codes currently available are in the search directions employed and whether an infeasible start or self-dual embedding technique is used [224]. Choosing the correct search direction and method of dealing with infeasibility may have a significant effect on the speed and accuracy of the solution obtained but in general will have little effect on the storage requirements of the algorithm. Storage issues have a very significant bearing on the size of SDPs in particular, that can be solved using primal-dual interior point

methods and are often of greater concern than CPU time issues, in the sense that storage issues can prevent a solution being obtained, regardless of time available.

All the algorithms require the creation, storage and Cholesky factorisation of a large dense positive definite Schur Complement matrix at each iteration of the algorithm. The matrix is of size $p \times p$ where p is the number of linear equality constraints. In the HKM method (see appendix B) the Schur complement matrix is defined as

$$\mathbf{O} = [\mathbf{F}(\mathbf{Z}^{-1}\mathbf{F}_1\mathbf{X}), \mathbf{F}(\mathbf{Z}^{-1}\mathbf{F}_2\mathbf{X}), \ldots, \mathbf{F}(\mathbf{Z}^{-1}\mathbf{F}_p\mathbf{X})]. \tag{3.46}$$

Unfortunately the Schur complement matrix is fully dense for SDP problems and therefore the storage requirements are $\mathcal{O}(p^2)$. The $\mathbf{C}$ matrix which is also often dense requires $\mathcal{O}(n^2)$ worst case. The $\mathbf{X}$ and $\mathbf{Z}$ matrices often are block diagonal with i blocks of dimension $n \times n$. The storage requirements for these are $\mathcal{O}(n_1^2 + \cdots + n_i^2)$ each, while the storage requirements for the constraint matrices is $\mathcal{O}(pn^2)$worst case. However usually the constraint matrices are sparse requiring $\mathcal{O}(1)$ entries per constraint. Therefore the constraint matrices typically require $\mathcal{O}(p)$ storage. The approximate storage requirements for an SDP being solved by the CSDP programme, for example, are [224]

$$\text{Storage (bytes)} = 8(p^2 + 11(n_1^2 + \cdots + n_i^2)). \tag{3.47}$$

Given that a 32 bit pc can address a maximum of 4 Gb RAM, and a 22000×22000 matrix of double precision floating point numbers requires 3.9 Gb of RAM, this places obvious constraints on the ability of primal-dual interior point methods (PDIPM) to solve large problems.

Research is progressing in a number of directions to overcome these issues. Already two codes, CSDP and SDPA have been parallelised [224, 225] making use of the same algorithms deployed in a cluster environment. CSDP has also been ported to a 64-bit environment making larger ammounts of memory addressable. More fundamental research is looking at the way the Newton system in interior point methods is solved. Promising avenues include using iterative linear algebra routines to solve the Newton system. There is also research being carried out in methods which do not use the PDIPM formalism, but some of these will only work for a subset of SDP problems. Others can produce less accurate solutions which may be acceptable in many cases. Examples include bundle methods [226], the spectral bundle method [227] and boundary point methods [228].

3.8.4 Dealing with solvation

An important topic when considering treatment of potential energy surfaces for chemical systems and to which no reference has been made so far in the context of sum of squares decompositions is the topic of solvation. In dynamical sampling approaches such as molecular dynamics the solvent can be treated explicitly or implicitly. In an explicit solvent model the solvent molecules are typically modelled as a rigid unit of atoms and can include atom centred and off-atom charges depending on the complexity of the model. Water is the most common solvent in biological simulations and there are a number of models for water molecules ranging from the simple SPC type models [229] to polarisable models such as TIP5P [230]. An important effect of the solvent is its screening effect on the electrostatic interactions between the solute atoms. A clever and less computationally expensive alternative to an explicit model (at least for small systems) is the generalised Born approach of Still et al. [231]. In the generalised Born approach the free energy of solvation is given by

$$\Delta G = -\frac{1}{2}\left(1 - \frac{1}{\epsilon}\right)\sum_{i}^{n}\sum_{j}^{n}\frac{q_i q_j}{\sqrt{r_{ij}^2 + \alpha_i\alpha_j \exp(-r_{ij}^2/4\alpha_i\alpha_j)}}, \tag{3.48}$$

where ϵ is the relative permittivity and α_i and α_j are the Born Radii of the atoms i and j which would need to be known a priori. The exponential $\exp(-r_{ij}^2/4\alpha_i\alpha_j)$ could then be fitted to a polynomial in r_{ij} over the range of interest. The term in the sum could be squared above and below the line without altering its value resulting in a rational polynomial in r_{ij}. This could then be treated in a similar way to the van der Waals and Coulombic terms. However in order to calculate the Born radii a knowledge of the level of exposure of the atoms in question to the solvent is required, hence implying a priori knowledge of the systems structure.

3.8.5 Identifying other minima

Being able to identify the global minimum on the potential energy surface of a system does not provide thermodynamic information about the system. To obtain parameters such as free energy and heat capacities a knowledge of the accessible minima at a given temperature is required. However the global minimum is a very good place to start a sampling process such as molecular dynamics or Monte Carlo.

An alternative approach could be to take the first derivative functions of the entire potential energy functional and define a feasibility problem using these first derivative functions as equality constraints. The solutions to these equations would provide the coordinates of all minima on the surface. However this would only be feasible for the simplest of systems.

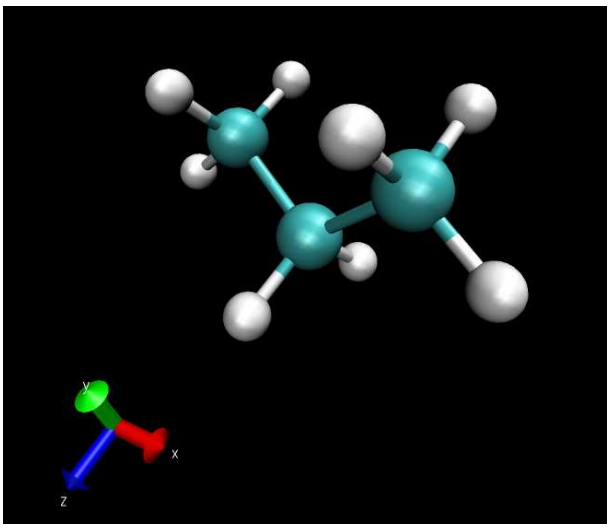


Figure 3.18: CPK representation of Propane (C_3H_8) in the trans configuration. Carbons are rendered in cyan and hydrogens are rendered in white. Covalent bonds are shown as rods. Rendered using Visual Molecular Dynamics (VMD) [221].

3.9 Conclusions

1. Model energy landscape problems have demonstrated the advantages a SOS decomposition approach brings to identifying the global minimum on a potential energy landscape.

2. A sum of squares decomposition for the potential energy surface of a simple molecular system using the Amber94 force field representation of the potential energy surface has been obtained. It bounds the potential energy exactly and obtains one of the three degenerate minimisers to within 0.21%.

3. The present limitations on the size of SDP which is solvable using interior point methods places restrictions on the practical applicability of the SOS decomposition approach to moderate to large scale global optimisation for molecular systems. It may be worthwhile to investigate alternative algorithms for solving SDPs. Polynomial complexity bounds will be lost in this case.

4. Polynomial representation of solvent contributions to a potential energy surface remains an outstanding challenge. If significantly more efficient SOS decompositions can be obtained, then an investigation into converting the Generalised Born model for solvation into a polynomial form will be worthwhile.

Chapter 4

Second order cone programming in coarse grained molecular dynamics

4.1 Introduction

One of the most difficult problems facing the molecular simulation community today involves the linking together of simulations at different length and time scales [232]. In chemistry important phenomena occur on vastly different timescales ranging from the femtosecond bond vibrations, through diffusion processes (nanosecond), to the folding of proteins which occurs on a microseconds to seconds timescale. Another interesting long timescale process is that of liquid crystalline ordering from disordered to nematic or smectic phases which has assumed huge technological importance and occurs on the order of tens of nanoseconds to microseconds. This problem is made even more difficult by the range of length-scales required to study these phenomena. Hydrogen bonding which is critical to stabilisation of protein structure during folding occurs on the Angstrom length-scale whereas a typical layer spacing in a liquid crystal display is of the order of $5\mu m$.

Although the quantum-atomistic bridge in length-scales is now pretty well defined the same cannot be said for the atomistic to mesoscopic regime. However some approaches have been found to be quite useful in many cases and these will be discussed later.

Molecular systems contain groups of covalently bonded atoms interacting via inter- and intra-molecular forces. There should in general be no difference in how these two sets of forces are treated but it can be assumed that intra-molecular forces are at least an order of magnitude greater than intermolecular forces. Problems arise in dynamical simulations of large molecular systems when a classical treatment of bond vibrations and bending is undertaken not least of which is the minuscule time-step required to integrate these motions. One solution is to take intra-molecular bonds to be of a fixed length, treating the molecule as a rigid body. This is of course an approximation but one which is quite good at ordinary temperatures

where the amplitudes of vibration are small relative to the size of the molecule itself. It is a poorer approximation for molecules with lower frequency degrees of freedom such as torsional angles. But in many cases the approximation works quite well and in this chapter systems of rigid molecules are treated using the sum of squares decomposition formalism. Here the sum of squares formalism is applied in a dynamical context rather than using it directly to minimise static potential energy surfaces. It is applied instead to a geometrical problem which arises when trying to propagate dynamical trajectories of rigid body representations of molecules.

Molecules which can be represented as coarse grained spheres could have their short range interactions represented by a specially parameterised Lennard-Jones functional. However in early computer simulations of rigid body aspherical molecules the short range attractive and repulsive intermolecular interactions were represented by Lennard-Jones interactions between specified atomic sites in the molecule. In large molecules this required the use of many sites (typically one per atom) to represent the short range interactions effectively. This was both computationally expensive and in many cases unnecessary for the level of detail required. These approaches were termed site-site potentials.

4.1.1 Single-site anisotropic potentials

One approach which avoids the computational expense of site-site potentials is the single site potential in which the entire intermolecular interaction is modelled by a single potential functional. A number of these have been reported in the literature [233–236].

The Gaussian overlap potential

One of the earliest mathematically tractable single site potentials was the Gaussian Overlap Potential (GOP) of Berne and Pechukas. This involved a Gaussian representation of the repulsive part of the intermolecular potential. Each molecule was regarded as an ellipsoid each of which was rotationally symmetric about a principal axis. A Gaussian function of a 3 dimensional vector $\mathbf{x}$ was defined as

$$\begin{aligned} G(\mathbf{x}) &= \exp(-\mathbf{x}\cdot\boldsymbol{\gamma}^{-1}\mathbf{x}), \\ \boldsymbol{\gamma} &= (\sigma_{\parallel}^2-\sigma_{\perp}^2)\mathbf{u}\mathbf{u}+\sigma_{\perp}^2\mathbf{I}. \end{aligned} \tag{4.1}$$

where $\mathbf{u}$ is a unit vector along the principal axis of the molecule and $\mathbf{I}$ is the unit matrix. G can be regarded as representing the distribution of matter with the spatial extent of the distribution being of order $2\sigma_{\parallel}$ along $\mathbf{u}$ and of order $2\sigma_{\perp}$ perpendicular to it. The interaction potential can be assumed to be proportional to the integral of the overlap of two Gaussians from separate molecules.

The overlap integral can be expressed in the form

$$S(\mathbf{u}_1, \mathbf{u}_2, \mathbf{r}) = \epsilon(\mathbf{u}_1, \mathbf{u}_2) \exp[-\mathbf{r}^2/\sigma^2(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}})], \tag{4.2}$$

where $\mathbf{r}$ is the inter-centre vector, $\epsilon(\mathbf{u}_1, \mathbf{u}_2)$ and $\sigma(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}})$ are orientation dependent strength and range parameters respectively, and $\hat{\mathbf{r}}$ is the unit vector in the $\mathbf{r}$ direction. The strength and range parameters are given respectively by

$$\epsilon(\mathbf{u}_1, \mathbf{u}_2) = \epsilon_0 \left[1 - \chi^2(\mathbf{u}_1, \mathbf{u}_2)^2\right]^{-1/2}, \tag{4.3}$$

$$\sigma(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}) = \sigma_0 \left(1 - \frac{\chi}{2}\left[\frac{(\hat{\mathbf{r}} \cdot \mathbf{u}_1 + \hat{\mathbf{r}} \cdot \mathbf{u}_2)^2}{1 + \chi \mathbf{u}_1 \cdot \mathbf{u}_2} + \frac{(\hat{\mathbf{r}} \cdot \mathbf{u}_1 - \hat{\mathbf{r}} \cdot \mathbf{u}_2)^2}{1 - \chi \mathbf{u}_1 \cdot \mathbf{u}_2}\right]\right)^{-1/2}, \tag{4.4}$$

and χ is defined as

$$\chi = \frac{(\sigma_{\parallel}^2 - \sigma_{\perp}^2)}{(\sigma_{\parallel}^2 + \sigma_{\perp}^2)}. \tag{4.5}$$

For fixed $\mathbf{u}_1$ and $\mathbf{u}_2$ the equipotentials are ellipsoids in $\mathbf{r}$ with principal axes along the eigenvectors of $(\boldsymbol{\gamma}_1 + \boldsymbol{\gamma}_2)$ where $\boldsymbol{\gamma}_i$ is the range matrix $\boldsymbol{\gamma}_i = (\sigma_{i\parallel}^2 - \sigma_{i\perp}^2)\mathbf{u}_i\mathbf{u}_i + \sigma_{i\perp}^2\mathbf{I}$. The overlap potential is obtained by incorporating the strength and range parameters into the Lennard-Jones functional:

$$U(\mathbf{u}_1, \mathbf{u}_2, \mathbf{r}) = \epsilon(\mathbf{u}_1, \mathbf{u}_2) \left\{ \left(\frac{\sigma(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}})}{\mathbf{r}}\right)^{12} - \left(\frac{\sigma(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}})}{\mathbf{r}}\right)^{6} \right\}. \tag{4.6}$$

A key advantage of the GOP is that despite being geometric in nature it is also sufficiently simple in analytic form and therefore differentiable such that calculations of forces and torques are possible for use in dynamical calculations. However there are also some shortcomings associated with the GOP. These can be seen when it is compared to a site-site interaction model for an equivalent system. The most significant of these is that the well depths are the same for all orientations of the two molecules and the well width is wider along the long principal axis compared to the well width along the shorter two principal axes.

The Gay-Berne potential

A modification of the GOP proposed by Gay and Berne [237] replaced the potential energy functional (equation 4.6) by the following

$$U(\mathbf{u}_1, \mathbf{u}_2, \mathbf{r}) = 4\epsilon(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}) \left\{ \left(\frac{1}{\mathbf{r} - \sigma(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}) + 1}\right)^{12} - \left(\frac{1}{\mathbf{r} - \sigma(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}) + 1}\right)^{6} \right\} \tag{4.7}$$

In addition a new strength parameter was defined which is a function of $\hat{\mathbf{r}}$ as well as $\mathbf{u}_1$ and $\mathbf{u}_2$. This parameter has the form

$$\epsilon(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}) = \epsilon^{\nu}(\mathbf{u}_1, \mathbf{u}_2)\epsilon'^{\mu}(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}), \tag{4.8}$$

where ν and μ are exponents, $\epsilon(\mathbf{u}_1, \mathbf{u}_2)$ is the original strength parameter and $\epsilon'(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}})$ is defined as

$$\epsilon'(\mathbf{u}_1, \mathbf{u}_2, \hat{\mathbf{r}}) = 1 - \frac{\chi'}{2}\left[\frac{(\hat{\mathbf{r}} \cdot \mathbf{u}_1 + \hat{\mathbf{r}} \cdot \mathbf{u}_2)^2}{1 + \chi' \mathbf{u}_1 \cdot \mathbf{u}_2} + \frac{(\hat{\mathbf{r}} \cdot \mathbf{u}_1 - \hat{\mathbf{r}} \cdot \mathbf{u}_2)^2}{1 - \chi' \mathbf{u}_1 \cdot \mathbf{u}_2}\right]. \tag{4.9}$$

Given ϵ_s as the desired strength parameter for the side by side configuration and ϵ_e the same for the end to end configuration then χ' is defined as

$$\chi' = \frac{\epsilon_s^{1/\mu} - \epsilon_e^{1/\mu}}{\epsilon_s^{1/\mu} + \epsilon_e^{1/\mu}}. \tag{4.10}$$

ν and μ are then fitted to an appropriate site-site potential.

The effect of these modifications is that the ellipsoidal repulsive core and the well width are more realistic than in the GOP potential. The Gay-Berne potential has gained enormous popularity within the liquid crystal modelling community and is the work-horse potential for liquid crystal research [238]. It has also gained popularity in recent years for studying lipids in biological membranes [239] and as coarse grained models for side-chains in protein structure prediction calculations [240].

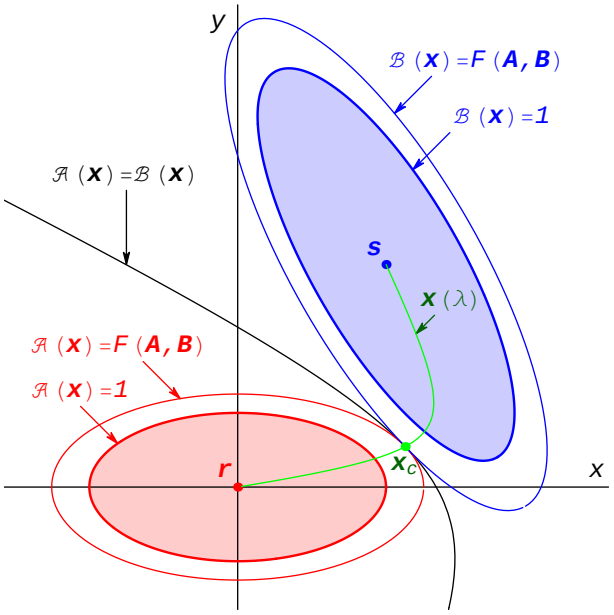


Figure 4.1: The elliptic contact function for two ellipsoidal particles. Two ellipsoidal particles $\mathcal{A}(\mathbf{x})$ and $\mathcal{B}(\mathbf{x})$ centred at points $\mathbf{r}$ and $\mathbf{s}$ respectively considered in the local reference frame of ellipsoid $\mathcal{A}(\mathbf{x})$. The original ellipsoids are scaled up or down by $\sqrt{F(\mathbf{A}, \mathbf{B})}$ until they touch tangentially at the contact point x_c. The curve $\mathbf{x}(\lambda)$ is shown as a green line [2].

The elliptic contact function

An alternative approach uses the elliptic contact function (ECF) which was originally proposed by Wertheim and Perram [241] and Perram et al. [242] which takes a geometrical approach to the problem. Two ellipsoids $\mathcal{A}(\mathbf{x})$ and $\mathcal{B}(\mathbf{x})$ centred at points $\mathbf{r}$ and $\mathbf{s}$ respectively can be represented as quadratic forms given by

$$\begin{aligned} \mathcal{A}(\mathbf{x}) &= (\mathbf{x}-\mathbf{r})\mathbf{A}(\mathbf{x}-\mathbf{r}) \leq 1, \\ \mathcal{B}(\mathbf{x}) &= (\mathbf{x}-\mathbf{s})\mathbf{B}(\mathbf{x}-\mathbf{s}) \leq 1, \end{aligned} \tag{4.11}$$

where $\mathbf{A}$ and $\mathbf{B}$ are positive definite matrices. The distance of closest approach of the two ellipsoids with specified orientations can be calculated through the elliptic contact function (see figure 4.1). The elliptic contact function for two ellipsoids $\mathcal{A}$ and $\mathcal{B}$ is given by the optimisation problem

$$F(\mathbf{A},\mathbf{B}) = \max_{\lambda} \min_{\mathbf{x}} S(\mathbf{x},\lambda), \tag{4.12}$$

where λ is a parameter in the interval $[0,1]$ and $S(\mathbf{x},\lambda)$ is a nonnegative function over the interval $[0,1]$ given by

$$S(\mathbf{x},\lambda) = \lambda\mathcal{A}(\mathbf{x}) + (1-\lambda)\mathcal{B}(\mathbf{x}). \tag{4.13}$$

Obtaining the distance of closest approach from the ECF

A highly efficient approach to obtaining a good approximation to the distance of closest approach between two ellipsoids was recently published by Paramonov and Yaliraki [2], based on the elliptic contact function. They took as an approximation to the true minimum distance the "directional distance of closest approach" (d_r) i.e. the minimum distance between ellipsoids along a vector parallel to the inter-centre vector which approaches the true distance from above and is given by (see figure 4.2)

$$d_r(\mathbf{A},\mathbf{B}) = r(1-F(\mathbf{A},\mathbf{B})^{-1/2}). \tag{4.14}$$

This directional distance can be used to construct an elliptic potential such as a Lennard-Jones radial potential. The following Lennard-Jones-like elliptic potential can be derived using the elliptic contact function

$$U = 4\epsilon_0 \left[\left(\frac{rF(\mathbf{A},\mathbf{B})^{1/2}}{r} \right)^{12} - \left(\frac{rF(\mathbf{A},\mathbf{B})^{1/2}}{r} \right)^{6} \right], \tag{4.15}$$

which is the elliptic contact potential developed by Perram et al. [242]. However this potential suffers from a number of disadvantages,

1. The potential does not become isotropic at large separations.
2. The depth of the potential minimum is constant for all relative orientations of the two particles.
3. The shape of the potential is problematic since it is wider along the longer semi-axis.

To remedy this Paramonov and Yaliraki used the directional distance to define a "shifted form" of the potential

$$U(\mathbf{A},\mathbf{B}) = 4\epsilon_0 \left[\left(\frac{\sigma_0}{d_r(\mathbf{A},\mathbf{B}) + \sigma_0} \right)^{12} - \left(\frac{\sigma_0}{d_r(\mathbf{A},\mathbf{B}) + \sigma_0} \right)^{6} \right]. \tag{4.16}$$

The depth of the wells equals one in any direction. To allow the well depth to vary as a function of direction, they defined different **A** and **B** matrices for the attractive and repulsive parts of the potential.

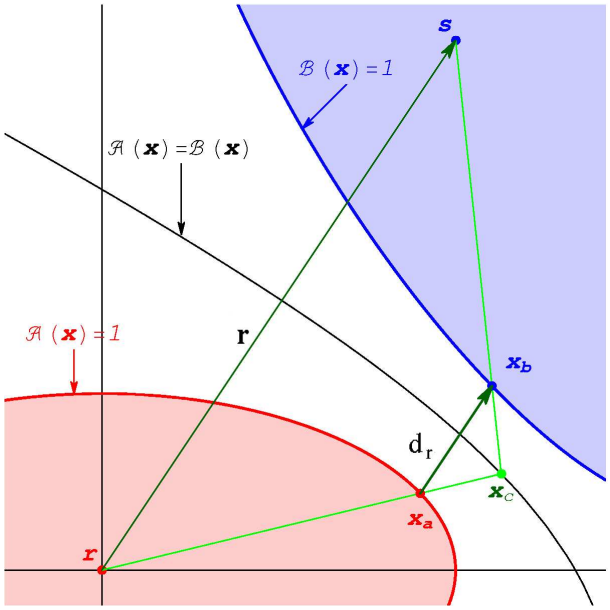


Figure 4.2: The directional distance of closest approach ($d_r(\mathbf{A},\mathbf{B})$), is the minimum distance between the two ellipsoids that is parallel to the inter-centre vector **r**. Adapted from [2].

4.2 Objective

The primary goal of this chapter is to show how a second order cone program can be applied to efficiently produce single-site anisotropic representations of molecules

in a coarse grained molecular dynamics programme. An extension of the approach from ellipsoidal to arbitrary non-convex shapes using a SOS relaxation is proposed. The SOS approach provides a useful general mathematical framework for studying interactions between arbitrarily shaped coarse-grained particles, whether they are convex or non-convex. However it is very computationally expensive to implement and at the moment would not be feasible for use in a dynamical calculation.

4.3 Ellipsoid representations using SOCP

The SOCP representation of the ellipsoidal minimum distance problem however appears to be a very nice example where solving a SOS optimisation problem could be feasible in a dynamical calculation. In the case of representations of ellipsoids, it turns out that the equations presented above simplify greatly since ellipsoids are effectively quadratic forms and can be written in the form described in equation 4.11.

It is useful to examine exactly what the relationship is between the matrices $\mathbf{A}$ and $\mathbf{B}$ and the ellipsoids they define. There are three ways of describing an ellipsoid geometrically.

1. In the description given in equation 4.11, the lengths of the semi-axes of the ellipsoid $\sqrt{\lambda_1}, \sqrt{\lambda_2}, \sqrt{\lambda_3}$ are defined as the square root of the eigenvalues of $\mathbf{A}^{-1}$ and the directions of the semi-axes are given by the eigenvectors (figure 4.3).

2. Another description of an ellipsoid is given by

$$\epsilon = \{\mathbf{A}^{-1/2}\mathbf{a} + \mathbf{r} \mid \| \mathbf{a} \| \leq 1\}, \tag{4.17}$$

   where $\mathbf{r}$ is centre of the ellipsoid. The geometrical interpretation of this description is clear. The set $\{\mathbf{a}\}$ defines a unit sphere in $\mathbb{R}^3$ and a linear transformation $\mathbf{A}^{-1/2}$ distorts the sphere into an ellipsoid of dimensions specified by the eigenvalues and eigenvectors of $\mathbf{A}^{-1/2}$. This is followed by translation to $\mathbf{r}$. Equation 4.17 can also be rewritten as

$$\epsilon = \{\mathbf{a}' + \mathbf{r} \mid \| \mathbf{A}^{1/2}\mathbf{a}' \| \leq 1\}, \tag{4.18}$$

   and $\mathbf{a}'$ is defined as $\mathbf{a}' = \mathbf{A}^{-\mathbf{1/2}}\mathbf{a}$. This description will be used in what follows.

3. Another interesting description is

$$\begin{aligned} \epsilon &= \{\mathbf{x} \mid f(\mathbf{x}) \leq 0\}, \\ f(\mathbf{x}) &= \mathbf{x}^T\mathbf{C}\mathbf{x} + 2\mathbf{d}^T\mathbf{x} + e, \\ &= \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix}^T \begin{bmatrix} C & d \\ d^T & e \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix}. \end{aligned} \tag{4.19}$$

where $\mathbf{C} = \mathbf{C}^T \succ 0$ and $e - \mathbf{d}^T\mathbf{C}^{-1}\mathbf{d} < 0$.

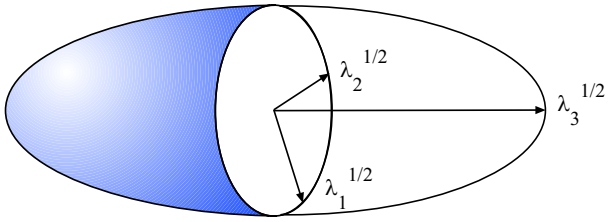


Figure 4.3: The principal axes of an ellipsoid determine its geometry. The most common type of ellipsoid used in coarse grained molecular simulation is one where two of the semi-axis lengths are degenerate i.e. $\lambda_1 = \lambda_2$ shown here. These are typically used to represent straight chain molecules or lipids.

In the case of representing ellipsoidal particles the minimum distance problem defined in equation 4.49 can be written in the form of equation 4.18 as:

$$\begin{aligned} \text{min:} \quad & \| (\mathbf{r} + \mathbf{a}') - (\mathbf{s} + \mathbf{b}') \|_2, \\ \text{such that:} \quad & \| \mathbf{A}^{1/2}\mathbf{a}' \|_2 \leq 1 \quad (\mathbf{a}' \in \mathcal{A}), \\ & \| \mathbf{B}^{1/2}\mathbf{b}' \|_2 \leq 1 \quad (\mathbf{b}' \in \mathcal{B}). \end{aligned} \tag{4.20}$$

Here the search for a sum of squares representation is now straightforward since the square of the norm functions will always be quadratic forms and hence a sum of squares. Minimising the square of the norm is equivalent to minimising the norm itself since the norm is a nonnegative quantity by definition.

In fact equation 4.20 can be redefined as follows:

$$\begin{aligned} \text{minimise:} \quad & d, \\ \text{subject to:} \quad & \| \mathbf{A}^{1/2}\mathbf{a}' \|_2 \leq 1, \\ & \| \mathbf{B}^{1/2}\mathbf{b}' \|_2 \leq 1, \\ & \| (\mathbf{r} + \mathbf{a}') - (\mathbf{s} + \mathbf{b}') \|_2 \leq d, \end{aligned} \tag{4.21}$$

where d is an upper bound on the minimum distance d_{true} This is in the form of a second order cone programme described in equation 2.23. As described in chapter 2 second order cone programmes are convex programmes which can be regarded as a subset of semidefinite programmes. As such they can be solved using interior point methods and therefore are solvable in polynomial time. Being convex optimisation

problems they also possess similar duality properties to both linear and semidefinite programmes.

Indeed interior point SDP solvers such as SeDuMi can also solve SOCPs. However there are also very efficient solvers which also use interior point methods but which were developed specifically for solving SOCP's which are much more efficient and can deal with large systems of many hundreds of constraints. It was deemed that it would be a worthwhile exercise to test the feasibility of using such a solver to calculate the minimum distance between ellipsoidal representations of molecules and evaluate the possibility of incorporating such a routine into a coarse-grained molecular dynamics programme.

The MOSEK SOCP solver

The SOCP solver chosen was MOSEK [243]. One of the advantages of MOSEK is that it provides i) a direct interface for Matlab which is ideal for testing and development and ii) an application programming interface for the C and C++ programming languages, which allows direct calls to MOSEK routines from an external piece of software written in C or C++. The algorithm has been implemented in C and is scalable across parallel machines if required. It also exploits sparsity where possible. Therefore it is ideal for incorporation into a piece of software such as a molecular dynamics programme.

MOSEK uses a homogeneous model which avoids the problem of needing to detect whether the starting point is strictly feasible (see appendix B.3.2 for a brief discussion). The idea is to embed the optimisation problem presented into a slightly larger problem which always has a solution. The solution to the embedded problem either gives a certificate of infeasibility or a (scaled) optimal solution to the original problem. A more detailed description of self-dual or homogeneous embedding can be found in [244]. The interior point method traces the central path and the search direction used is the Nesterov-Todd search direction. MOSEK also preprocesses the initial problem presented to it to remove any obvious redundancies, including fixed variables and redundant constraints. Also if possible linear free variables are substituted out of the problem.

4.4 Initial evaluation

If the minimum distance calculated using these approaches is to be of use in the calculation of forces and torques in a molecular dynamics programme then it is vital that the solver be able to predict the minimum distance to a high degree of numerical accuracy. Otherwise the behaviour of the system will be non-physical and the energy of the system may not be conserved. Although an error may appear

small in a single calculation, when these errors are summed over interactions between many molecules and over many time-steps they can lead to large energy leaks into or out of the system.

Therefore an evaluation of the accuracy of the MOSEK algorithm in calculating the minimum distance between ellipsoids was performed. The accuracy was evaluated as a function of two parameters:

- The angle between the two ellipsoids.
- The aspect ratio of the ellipsoids.

The results were also compared against a similar optimisation carried out using the non-linear programming routine FMINCON in the optimisation toolbox in Matlab. Initially the plan was to use FMINCON as a benchmark to compare with the MOSEK solution, however as will become evident from the results of the test, the MOSEK results were far superior to the FMINCON results.

An important consideration was that the solver be robust enough to produce accurate solutions without user intervention in setting parameters as in a molecular dynamics simulation it would not be feasible to set the solver parameters based on the configuration of the ellipsoids at a particular time-step.

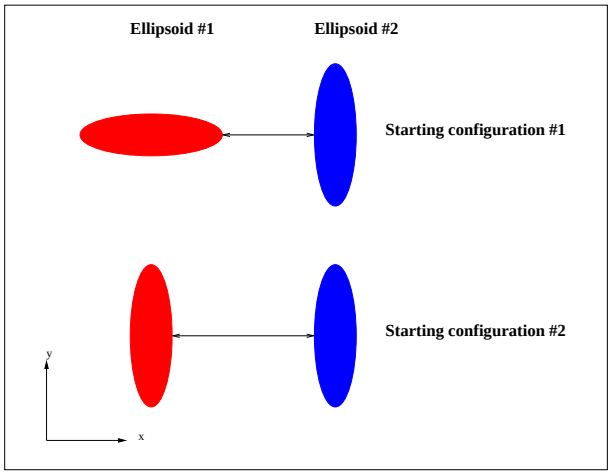


Figure 4.4: Two initial configurations were considered. In the first configuration, ellipsoid 1 was oriented with its largest principal axis parallel to the x-axis while the largest principal axis of ellipsoid 2 was oriented parallel to the y-axis. The second ellipsoid is then incrementally rotated about the z-axis and the minimum distance between the ellipsoids is calculated at each position. In the second configuration ellipsoid 1 was oriented with its largest principal axis parallel to the y-axis.

4.4.1 Test setup

The test was setup with two ellipsoids with semi-axis lengths $[s_x, 1.0, 1.0]$ and $[s_x, 1.0, 1.0]$ where s_x is the largest semi-axis and ranges from 1.0 to 96.0 in increments of 5.0.

The centres of the ellipsoids are positioned at the points $[0, 0, 0]$ and $[2s_x + 10, 0, 0]$ and are therefore always at least 10.0 units apart. Two initial configurations were considered (figure 4.16). In the first configuration, ellipsoid 1 was oriented with its largest principal axis aligned parallel to the x-axis while the largest principal axis of ellipsoid 2 was oriented parallel to the y-axis. The second ellipsoid is then incrementally rotated about the z-axis in increments of $\pi/50$ and over a range from zero to π and the minimum distance between the ellipsoids was calculated at each position. In the second configuration ellipsoid 1 was oriented with its largest principal axis parallel to the y-axis and the rotation of the second ellipsoid was repeated. The reason for choosing two starting configurations was to examine all possible alignments of the ellipsoids. This is important since one of the factors most likely to affect the convergence of the optimisation algorithm to the minimum distance vector is the curvature of the ellipsoids in the region of the minimum distance points. The process was repeated for each of the lengths of the major semi-axis.

4.4.2 Results

At each point the minimum distance and the points on each ellipsoid of the minimum distance vectors **a** and **b** were calculated using the FMINCON and MOSEK solvers. In order to evaluate the accuracy of the minimisation, the equation for the ellipsoid (4.11) was evaluated at the vectors **a** and **b**. The deviation of the result from 1.0 indicates an error in the minimum distance, i.e. the quadratic constraint was being breached if greater than 1.0 and optimality was not achieved if less than 1.0. In configuration 1, the absolute deviation for MOSEK was less than 1.0×10^{-4} with the standard convergence parameters (figure 4.5). In configuration 2, the absolute deviation was within 1×10^{-6} (figure 4.6). This was deemed to be of sufficient accuracy for calculating of forces and torques in a molecular dynamics simulation. The FMINCON solver was not as successful. In configuration 1, FMINCON produced erroneous results over all angles and particularly in ellipsoid 2 (figure 4.7). In configuration 2, FMINCON performed better but still produced some erroneous results near to 0 and π (figure 4.8).

It is clear that configuration 1 produced poorer results from both solvers and this is likely to be related to the fact that minimum distance vector **a** lies on a point of maximum curvature. It is interesting however that the greatest errors for both MOSEK and FMINCON occur not when the two ellipsoids are aligned end to end ($\pi/2$ in both figures 4.5 and 4.7) but when ellipsoid 2 is just slightly perturbed from that point.

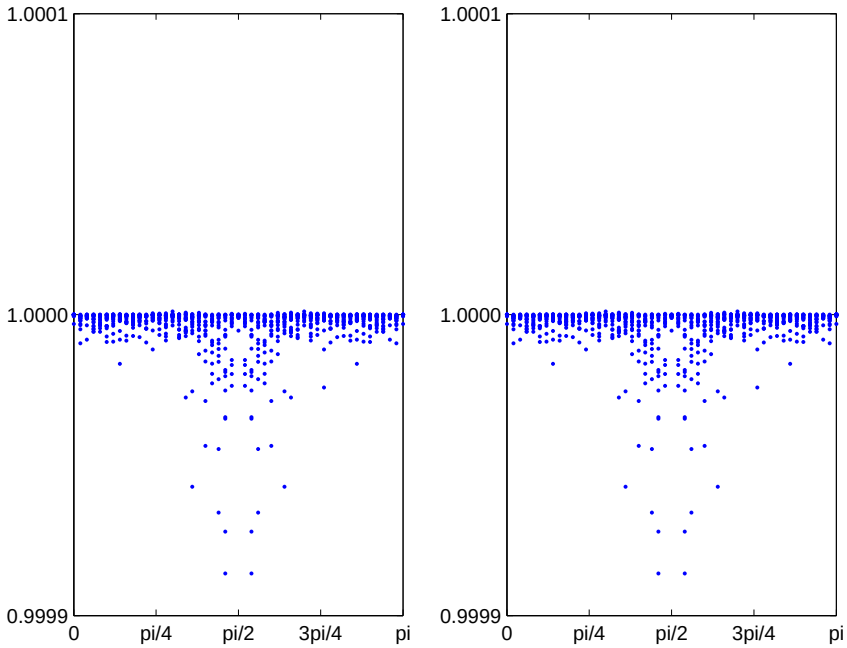


Figure 4.5: Quadratic form evaluations at the minimum distance points obtained using MOSEK for ellipsoid 1 (left) and ellipsoid 2 (right) with starting configuration 1 shown in figure 4.16. All the minimum distances lie within the ellipsoid and are within a tolerance of 1×10^{-4} of the exact boundary.

4.4.3 Conclusion

The results of this test demonstrate that MOSEK is both robust and accurate in calculating the distance of closest approach between ellipsoids within a wide range of aspect ratios and through all orientations. Therefore it was decided to proceed and implement a coarse grained molecular dynamics programme using MOSEK to calculate the distance of closest approach.

4.5 Molecular dynamics implementation

A coarse-grained molecular dynamics programme was written which calls the MOSEK API to calculate the minimum distance between the ellipsoids prior to the force calculation routine. The programme was written primarily in FORTRAN77 with a C function incorporated for calling the MOSEK API. The programme outline is summarised in table 4.1 and can be divided into the following main sections:

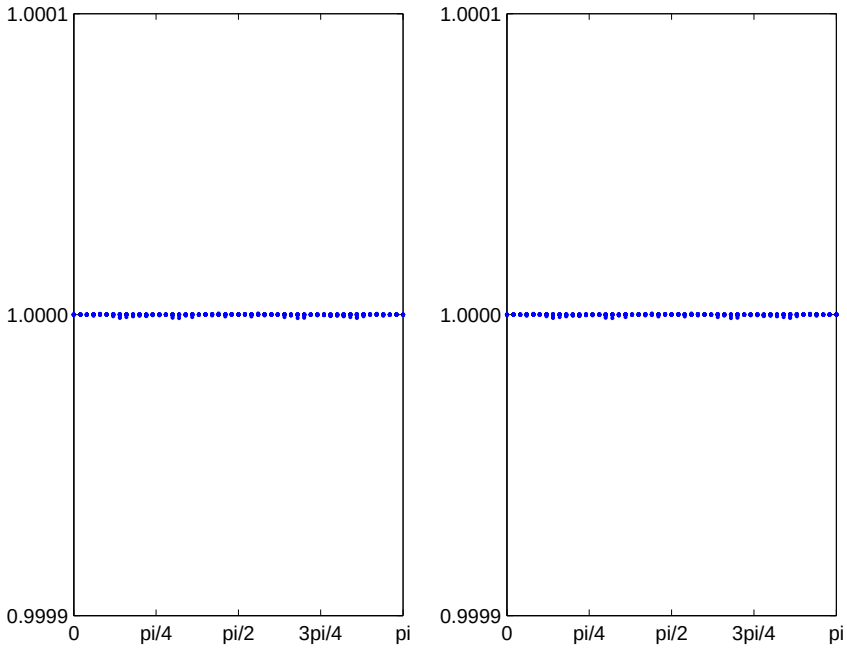


Figure 4.6: Quadratic form evaluations at the minimum distance points obtained using MOSEK for ellipsoid 1 (left) and ellipsoid 2 (right) with starting configuration 2 shown in figure 4.16. All the minimum distances lie within the ellipsoid and are within a tolerance of 1×10^{-6} of the exact boundary.

4.5.1 Initialisation

Simulations were run in reduced units given in table 4.2. The mass of each ellipsoid is 1. Semi-axis lengths $[s_x, s_y, s_z]$ were then defined and from these principal moments of inertia were defined as

$$\begin{aligned} I_{xx} &= (1/5)(s_y^2 + s_z^2), \\ I_{yy} &= (1/5)(s_x^2 + s_z^2), \\ I_{zz} &= (1/5)(s_x^2 + s_y^2). \end{aligned} \tag{4.22}$$

The number of molecules, number of time-steps and the time-step size are also defined at this point. The simulation box dimensions are then defined. The box is rectangular and is typically sized to give the required density of particles.

Particle positions

Particles were assigned positions on a regular rectangular grid.

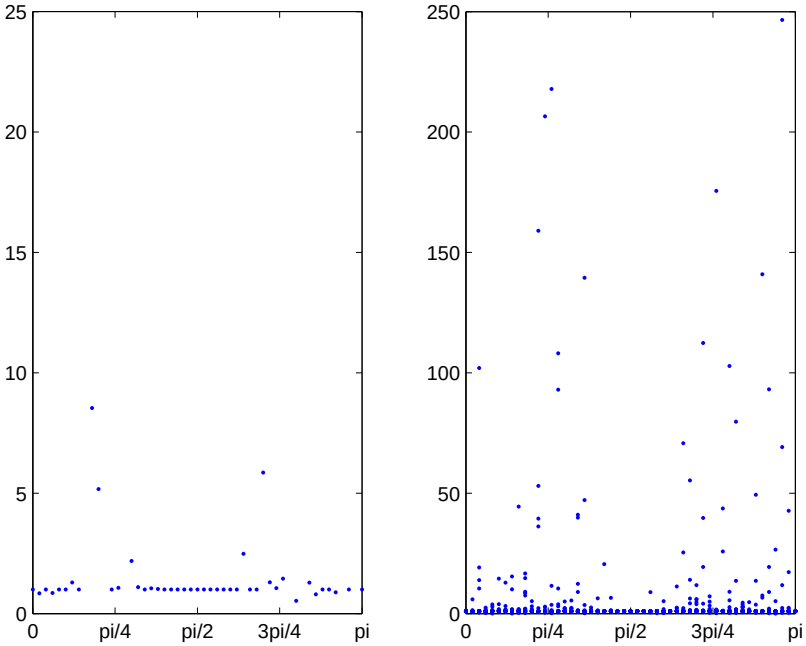


Figure 4.7: Quadratic form evaluations at the minimum distance points obtained using FMINCON for ellipsoid 1 (left) and ellipsoid 2 (right) with starting configuration 1 shown in figure 4.16.

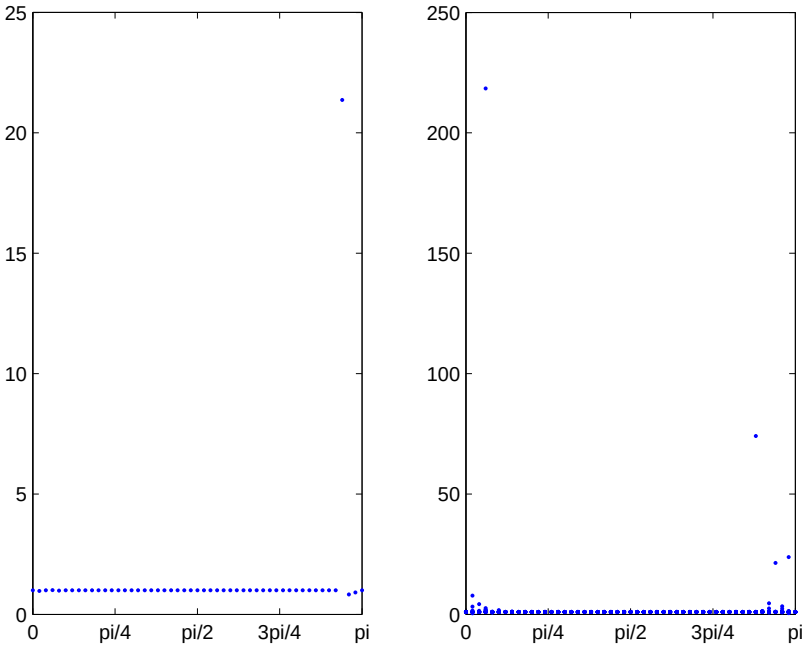


Figure 4.8: Quadratic form evaluations at the minimum distance points obtained using FMINCON for ellipsoid 1 (left) and ellipsoid 2 (right) with starting configuration 2 shown in figure 4.16.

MOLSOCP Coarse grained molecular dynamics programme:

1. **Initialisation**
   (a) Initialise particle positions on grid.
   (b) Initialise particle quaternions.
   (c) Initialise linear and angular velocities.
2. **Main Loop** Repeat while $i \leq$ NSTEP (number of steps)
   (a) Build molecule interaction list.
   (b) Solve SOCP to obtain minimum distances.
   (c) Calculate forces and torques.
   (d) Integrate equations of motion.
   (e) Write data for step.
   (f) Compile Statistics
      i. Write particle positions and velocities.
      ii. Write kinetic, potential and total energies.
   (g) increment time-step: i=i+1.

End Programme

Table 4.1: Summary of the MOLSOCP coarse grained molecular dynamics programme. Each step is described in detail in section 4.5.

Quaternions

In a standard dynamical problem with isotropic particles each particle has 3 degrees of freedom. However in rigid body dynamics the particles are anisotropic and therefore have 6 degrees of freedom, i.e. the original 3 translational and an additional 3 rotational degrees of freedom. The standard way to represent the three rotational degrees of freedom is through Euler angles (ϕ, θ, ψ) which are rotations about the z, x, z axes executed in that order. Using x convention Euler angles, a rotation matrix can be assembled as follows:

$$\mathbf{R} = \begin{bmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & \sin\theta \\ 0 & -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} \cos\phi & \sin\phi & 0 \\ -\sin\phi & \cos\phi & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (4.23)$$

Because the problems of translation and rotation are separable it is possible to propagate the translational and rotational degrees of freedom of the rigid bodies separately with the translational degrees of freedom being propagated according

density	$\rho^* = \rho\sigma^3$
temperature	$T^* = k_B T/\epsilon$
energy	$E^* = E/\epsilon$
pressure	$P^* = P\sigma^3/\epsilon$
time	$t^* = (\epsilon/m\sigma^2)^{1/2}t$
force	$\mathbf{f}^* = \mathbf{f}\sigma/\epsilon$
torque	$\tau^* = \tau/\epsilon$

Table 4.2: Properties of the system are quoted in reduced units.

to Newton's equations of motion and the rotational degrees of freedom propagated according to Euler's equations of motion in a body fixed reference frame:

$$\begin{aligned}
\dot{\omega}_x^b &= \frac{\tau_x^b}{I_{xx}} + \left(\frac{I_{yy} - I_{zz}}{I_{xx}}\right)\omega_y^b\omega_z^b, \\
\dot{\omega}_y^b &= \frac{\tau_y^b}{I_{yy}} + \left(\frac{I_{zz} - I_{xx}}{I_{yy}}\right)\omega_z^b\omega_x^b, \\
\dot{\omega}_z^b &= \frac{\tau_z^b}{I_{zz}} + \left(\frac{I_{xx} - I_{yy}}{I_{zz}}\right)\omega_x^b\omega_y^b,
\end{aligned} \tag{4.24}$$

where ω_i^b is the angular velocity about axis i in the body fixed reference frame, τ_i^b is the torque about axis i in the body fixed reference frame. Conversion between body and space fixed reference frames and vice versa is carried out through the operation of rotation matrices $\mathbf{R}$ (equation 4.23), and $\mathbf{R}^{-1}$. Notice that $\mathbf{R}^{-1} = \mathbf{R}^T$. The equations of motion for the Euler angles can be written as:

$$\dot{\phi} = -\omega_x^s\frac{\sin\phi\cos\theta}{\sin\theta} + \omega_y^s\frac{\cos\phi\cos\theta}{\sin\theta} + \omega_z^s, \tag{4.25}$$

$$\dot{\theta} = \omega_x^s\cos\phi + \omega_y^s\sin\phi, \tag{4.26}$$

$$\dot{\psi} = \omega_x^s\frac{\sin\phi}{\sin\theta} - \omega_y^s\frac{\cos\phi}{\sin\theta}. \tag{4.27}$$

where the s superscript designates the space fixed reference frame. Euler angles suffer from a serious drawback which can be understood from these equations. As θ approaches 0 or π the fractional terms diverge to infinity making Euler angles problematic for dynamical algorithms. A number of workarounds have been devised, however the most elegant and straightforward way is to switch to quaternions to calculate the rotation matrix $\mathbf{R}$. A quaternion is a set of four scalars which can be written as $\mathbf{Q} = (q_0, q_1, q_2, q_3)$. In order for the quaternion to represent orientation in three dimensional space it must be subject to the constraint

$$q_0^2 + q_1^2 + q_2^2 + q_3^2 = 1. \tag{4.28}$$

The relationship between a quaternion and the corresponding set of Euler angles is then given by

$$q_0 = \cos(\theta/2)\cos\left(\frac{\phi+\psi}{2}\right), \tag{4.29}$$

$$q_1 = \sin(\theta/2)\cos\left(\frac{\phi-\psi}{2}\right), \tag{4.30}$$

$$q_2 = \sin(\theta/2)\sin\left(\frac{\phi-\psi}{2}\right), \tag{4.31}$$

$$q_3 = \cos(\theta/2)\sin\left(\frac{\phi+\psi}{2}\right). \tag{4.32}$$

The rotation matrix is then defined as

$$\mathbf{R} = \begin{pmatrix} q_0^2+q_1^2+q_2^2+q_3^2 & 2(q_1q_2+q_0q_3) & 2(q_1q_3-q_0q_2) \\ 2(q_1q_2-q_0q_3) & q_0^2-q_1^2+q_2^2-q_3^2 & 2(q_2q_3+q_0q_1) \\ 2(q_1q_3+q_0q_2) & 2(q_2q_3-q_0q_1) & q_0^2-q_1^2-q_2^2+q_3^2 \end{pmatrix}. \tag{4.33}$$

The time derivatives of quaternions can be written as

$$\frac{d\mathbf{q}}{dt} = \frac{1}{2}\begin{pmatrix} 0 & \omega_z & -\omega_x & -\omega_y \\ -\omega_z & 0 & -\omega_y & \omega_x \\ \omega_x & \omega_y & 0 & \omega_z \\ \omega_y & -\omega_x & -\omega_z & 0 \end{pmatrix}\mathbf{q} \equiv \mathbf{Q}(\omega)\mathbf{q}. \tag{4.34}$$

Initial velocities

Initial linear and angular velocities were randomised from a single Gaussian distribution with mean defined by the required temperature and corrected so that overall linear momentum is zero. After some equilibration steps the velocity distribution reached a Maxwell-Boltzmann distribution.

4.5.2 Main loop

The main loop is repeated for the number of specified time-steps.

Periodic boundary conditions

Periodic boundary conditions were implemented at the edges of the simulation box so molecules leaving through one side of the box enter on the opposite side with the same linear and angular momenta. The minimum image convention was applied so each particles interacts only with the nearest image of the particles it is interacting with. The minimum image calculation is based on the inter-centre distance rather than the minimum distance.

At each time-step each possible interaction between molecule pairs (including their images) was examined and if the inter-centre distance between the molecules was below a parameter r_{cut} then the pair were added to an interaction list. The minimum distance vector for each pair of molecules on the interaction list was calculated in the following way.

Calculate minimum distances

The orientation matrices for each molecule, and vectors of the centre of mass coordinates as well as the interaction list were used to build a SOCP programme statement and this is submitted to MOSEK through the API in a single call which returns a vector of minimum distances between interacting molecules. The response code from MOSEK are also examined to ensure successful execution. Checking that the response code is correct ensures that the duality gap is below 1×10^{-7}. An example with 32 particles in a simulation box is shown in figure 4.9.

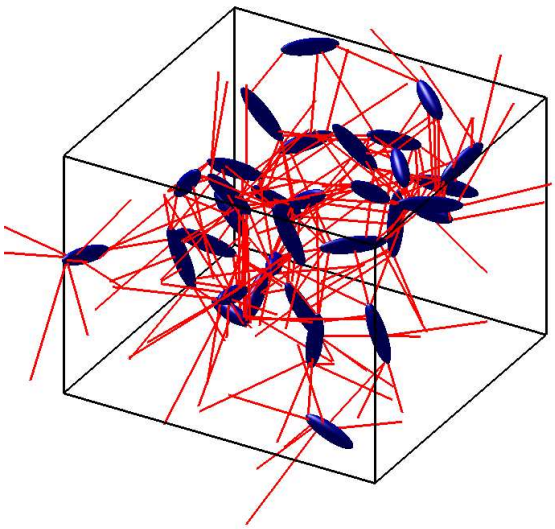


Figure 4.9: Minimum distance vectors (shown as red lines) for 32 ellipsoidal particles in a cubic box calculated using MOSEK as part of a molecular dynamics simulation. Notice that the molecules interact with the nearest image of its interacting partner even if it is outside the box.

Calculate forces and torques

If the minimum distance between molecule i and j is within a second cut-off distance r_{cut2} then the applied force is calculated by substituting the minimum distance into the first derivative of the Lennard-Jones potential energy given by

$$\mathbf{f}_{ij} = \frac{24\epsilon}{r_{ij}^2}\left[2\left(\frac{\sigma}{r_{ij}}\right)^{12} - \left(\frac{\sigma}{r_{ij}}\right)^{6}\right]\mathbf{r}_{ij}. \tag{4.35}$$

where $\mathbf{r}_{ij}$ is the minimum distance vector between molecules i and j. $\mathbf{f}_{ij}$ is then applied to molecule i and $-\mathbf{f}_{ij}$ is applied to molecule j. Truncation of the potential at the cut-off distance r_{cut2} means that the potential contains a discontinuity and this causes stability problems. To overcome this a shift was added to the potential U_c which shifts the potential to zero at the cut-off. In addition a small linear term was added to the potential to ensure that the force remains continuous also,

$$U_s(r_{ij}) = U(r_{ij}) - U_c - \left(\frac{dU(r_{ij})}{dr_{ij}}\right)_{r_{ij}=r_{\text{cut2}}} (r_{ij} - r_{\text{cut2}}) \quad \text{for } r_{ij} \leq r_{\text{cut2}}. \tag{4.36}$$

The force now goes smoothly to zero at r_{cut2} aiding both energy conservation and numerical stability (see figure 4.11). However it should be borne in mind that these modifications mean that the ensemble sampled and hence any thermodynamic properties calculated are not exactly those of a Lennard-Jones fluid. However perturbation approaches to deal with this have been reported in the literature, see for example [245, 246].

In order to carry out the integration of the equations of motion the total intermolecular forces acting on each molecule must be summed and in addition the torque applied to the molecule as a result of the total applied force must be calculated from

$$\tau = \mathbf{f}_{ij} \times \mathbf{a}, \tag{4.37}$$

where $\mathbf{f}_{ij}$ is the applied force and $\mathbf{a}$ is the vector to the point of minimum distance (see figure 4.10).

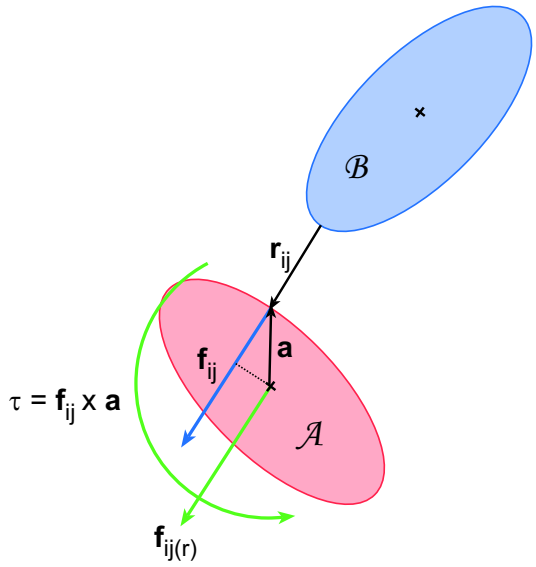


Figure 4.10: The force and torque on molecule $\mathcal{A}$ resulting from interactions between molecule $\mathcal{A}$ and $\mathcal{B}$. The force $\mathbf{f}_{ij}$ is decomposed into a force $\mathbf{f}_{ij(r)}$ (identical in magnitude and direction to $\mathbf{f}_{ij}$) acting radially on the centre of mass and a torque τ acting about the centre of mass.

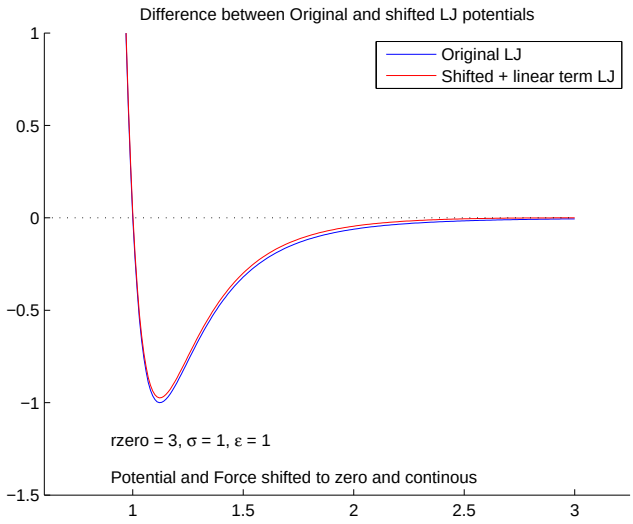


Figure 4.11: Original and shifted Lennard-Jones potential with shift and linear correction described in equation 4.36.

Integrate equations of motion

As described earlier since the linear and angular equations of motion are separate they are integrated separately at each time-step. Time integration of the linear velocities is carried out using the "leap frog" algorithm [247]. The algorithm is as follows

$$\begin{aligned} \mathbf{r}(t+\delta t) &= \mathbf{r}(t) + \delta t \mathbf{v}(t+1/2\delta t), \\ \mathbf{v}(t+1/2\delta t) &= \mathbf{v}(t-1/2\delta t) + \delta t \mathbf{a}(t). \end{aligned} \tag{4.38}$$

The positions $\mathbf{r}(t)$ and accelerations $\mathbf{a}(t)$ are stored at time t while the velocities $\mathbf{v}(t-1/2\delta t)$ are stored at $(t-1/2\delta t)$. The velocity equation is implemented first and this gives the velocities at a half time-step ahead of the positions which are at t. The current velocities which are necessary to calculate kinetic energies are calculated using

$$\mathbf{v}(t) = 1/2\left(\mathbf{v}(t+1/2\delta t) + \mathbf{v}(t-1/2\delta t)\right). \tag{4.39}$$

The positions can then be calculated using the velocities half a step ahead to give the positions one full step ahead. Then the algorithm begins again.

A leapfrog algorithm due to Omelyan [248] was used to integrate the Euler's equations of motion. This is a relatively recently developed algorithm which integrates the angular velocities instead of the angular momenta which leads to significant

simplifications. Using the Euler equations the leap frog algorithm can be written as

$$\begin{aligned}
\omega_x^{(n+1)}\left(t+\frac{\delta t}{2}\right) &= \omega_x\left(t-\frac{\delta t}{2}\right)+\frac{\delta t}{I_{xx}}\left[\tau_x(t)+(I_{yy}-I_{zz})\omega_y^{(n)}(t)\omega_z^{(n)}(t)\right], \\
\omega_y^{(n+1)}\left(t+\frac{\delta t}{2}\right) &= \omega_y\left(t-\frac{\delta t}{2}\right)+\frac{\delta t}{I_{yy}}\left[\tau_y(t)+(I_{zz}-I_{xx})\omega_z^{(n)}(t)\omega_x^{(n)}(t)\right], \\
\omega_z^{(n+1)}\left(t+\frac{\delta t}{2}\right) &= \omega_z\left(t-\frac{\delta t}{2}\right)+\frac{\delta t}{I_{zz}}\left[\tau_z(t)+(I_{xx}-I_{yy})\omega_x^{(n)}(t)\omega_y^{(n)}(t)\right].
\end{aligned} \tag{4.40}$$

The velocities at time t are evaluated using

$$\begin{aligned}
\omega_y^{(n)}(t)\omega_z^{(n)}(t) &= \frac{1}{2}\left[\omega_y\left(t-\frac{\delta t}{2}\right)\omega_z\left(t-\frac{\delta t}{2}\right)+\omega_y\left(t+\frac{\delta t}{2}\right)\omega_z\left(t+\frac{\delta t}{2}\right)\right], \\
\omega_z^{(n)}(t)\omega_x^{(n)}(t) &= \frac{1}{2}\left[\omega_z\left(t-\frac{\delta t}{2}\right)\omega_x\left(t-\frac{\delta t}{2}\right)+\omega_z\left(t+\frac{\delta t}{2}\right)\omega_x\left(t+\frac{\delta t}{2}\right)\right], \\
\omega_x^{(n)}(t)\omega_y^{(n)}(t) &= \frac{1}{2}\left[\omega_x\left(t-\frac{\delta t}{2}\right)\omega_y\left(t-\frac{\delta t}{2}\right)+\omega_x\left(t+\frac{\delta t}{2}\right)\omega_y\left(t+\frac{\delta t}{2}\right)\right].
\end{aligned} \tag{4.41}$$

Equations 4.40 and 4.41 constitute a system of three nonlinear equations in the unknowns $\omega_x(t+\delta t/2),\ \omega_y(t+\delta t/2),\ \omega_z(t+\delta t/2)$ which can be solved iteratively $n=0,\ldots,$ taking

$$\begin{aligned}
\omega_x^{(0)}(t+\delta t/2) &= \omega_x(t-\delta t/2), \\
\omega_y^{(0)}(t+\delta t/2) &= \omega_y(t-\delta t/2), \\
\omega_z^{(0)}(t+\delta t/2) &= \omega_z(t-\delta t/2).
\end{aligned} \tag{4.42}$$

as the initial guesses.

The quaternions are updated using

$$\mathbf{q}(t+\delta t)=\left(\mathbf{I}-\frac{\delta t}{2}\mathbf{Q}\right)^{-1}\left(\mathbf{I}+\frac{\delta t}{2}\mathbf{Q}\right)\mathbf{q}(t). \tag{4.43}$$

The iterative approach to the calculation of the angular velocities allows the propagation to satisfy Euler's equations of motion exactly. In addition the method maintains the norm of the quaternions to 1 as well as othonormality, eliminating the need for the unphysical rescaling employed in other integration algorithms. These two factors are the main contributors to the excellent energy conservation provided by Omelyan's integration algorithm.

Compiling statistics

At the end of each step each particles position, quaternion, linear velocity and angular velocity vectors are written to file while at an output step, the positions and quaternions are written to file for visualisation and animation generation using a specially constructed Matlab script for viewing the simulation. In addition at each step the kinetic, potential, total energies, instantaneous pressure and temperature are written to file. The instantaneous kinetic energy is a sum of the linear and angular components ($\mathcal{K} = \mathcal{K}_{\text{linear}} + \mathcal{K}_{\text{angular}}$) which are calculated respectively as

$$\mathcal{K}_{\text{linear}} = \frac{1}{2}\left(v_x^2 + v_y^2 + v_z^2\right), \tag{4.44}$$

and

$$\mathcal{K}_{\text{angular}} = \frac{1}{2}\left(I_{xx}\omega_x^2 + I_{yy}\omega_y^2 + I_{zz}\omega_z^2\right). \tag{4.45}$$

Instantaneous temperature was calculated from the kinetic energy using

$$\mathcal{T} = \frac{2\mathcal{K}}{3N}, \tag{4.46}$$

where N is the number of molecules. The internal virial $\mathcal{W}$ was calculated using

$$\mathcal{W} = \frac{1}{3}\sum_i \sum_{j>i} \mathbf{r}_{ij} \cdot \mathbf{f}_{ij}. \tag{4.47}$$

The instantaneous pressure $\mathcal{P}$ was calculated from the virial and the instantaneous temperature using

$$\mathcal{P} = \frac{N}{V}\mathcal{T} + \frac{\mathcal{W}}{V}. \tag{4.48}$$

4.6 Model system

To evaluate the method an A-B binary system of ellipsoids was simulated. The system consisted of 240 ellipsoids with semi-axis lengths $[2.5\sigma, 0.5\sigma, 0.5\sigma]$, in a box of dimensions $19.0\sigma \times 21.0\sigma \times 17.0\sigma$. The box dimensions were chosen to maximise the density of ellipsoids with these dimensions that could be fit into the box. The Lennard-Jones interactions were modelled with the following parameters $\sigma_{AA} = \sigma_{AB} = \sigma_{BB} = 1.5$ while $\epsilon_{AA} = \epsilon_{BB} = 1.0$ and $\epsilon_{AB} = 0.2$. The simulation was run for 40000 steps with a time-step $\delta t = 0.01$ to ensure numerical stability. The two cutoff parameters were defined as $r_{\text{cut1}} = 8.0\sigma$ and $r_{\text{cut2}} = 3.0\sigma$. The particles were placed on a rectangular array with randomised A/B properties taken from a uniform probability distribution. Particle positions and velocities were written every

40 time-steps.

The simulation was run on a 3GHz Pentium 4 processor with 3Gb RAM and took approximately 38 hours. By far the slowest part of each time-step was the formulation and solution of the SOCP taking typically 3 seconds. The number of minimum distances between ellipsoids that MOSEK calculated in each time-step was around 9000. It is expected that this number could be reduced significantly by using linked lists etc.

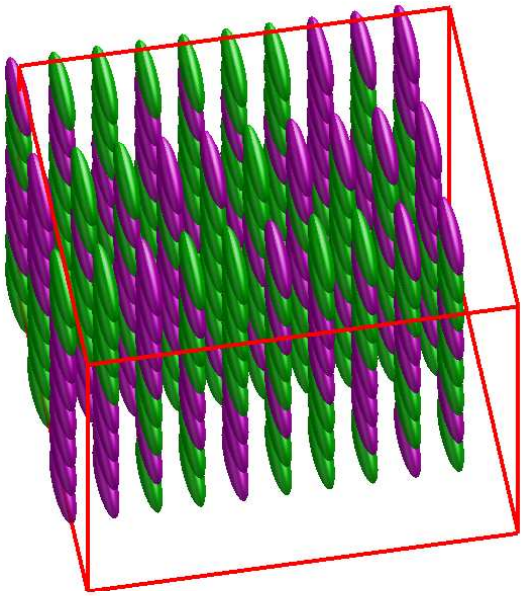


Figure 4.12: Simulation of a binary system of Lennard-Jones anisotropic molecules: initial configuration. The A type molecules are labelled green and the B type molecules are labelled purple. The molecules are arranged on a rectangular grid. The configuration was rendered using a Matlab script specifically developed for rendering of anisotropic molecules since standard molecular visualisation software tools are not well suited to visualising coarse-grained anisotropic molecules. Notice that the hard sphere core of the molecules are rendered and not the Lennard-Jones surface.

4.7 Results

The simulation was run for 20000 equilibration steps followed by a further 20000 production steps. The initial configuration is shown in figure 4.12. Notice that at first glance there appears to be ordering in the distribution of the particles. This in fact is not the case since a random process should produce sequences of A and B molecules of differing lengths. If roughly every second molecule were of type A and B then although that may look more random it would indicate correlation in the random number generator.

By step 500 the orientational correlation is beginning to disappear (figure 4.13) but the distribution of particles remains random. However by step 20000 the system has equilibrated into two distinct phases (figure 4.14) one consisting primarily of A type molecules and the other consisting of B type molecules. The system remains in a biphasic state for the remainder of the simulation.

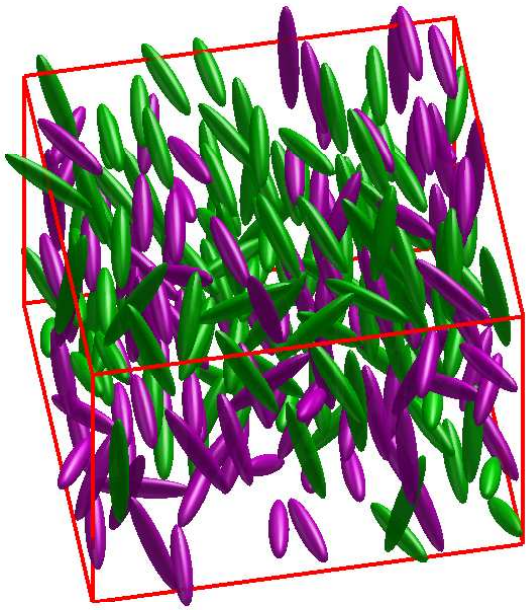


Figure 4.13: Simulation of a binary system of Lennard-Jones anisotropic molecules: Step 500. At this stage the molecules have lost much of their orientational correlation, however the spatial distribution of the molecules remains quite random.

The conservation of total energy throughout the simulation is good (figure 4.15) although there is a very small growth in total energy which requires further investigation. It is clear from the slope of the potential and kinetic energy plots that the system is completely equilibrated by step 20000. The temperature of the system has also settled to approximately 0.85 and the pressure to 0.08 (both in reduced units). The system remains stable for the remaining 20000 steps.

4.8 General SOS formulation for arbitrary geometries

In some cases the geometrical shape of the molecule to be represented is not at all well approximated by an ellipsoid or indeed even by a convex body. Two arbitrary geometrical shapes can again be represented as subsets $\mathcal{A}$ and $\mathcal{B}$ of $\mathbb{R}^3$ and once again a potential may be defined based upon a minimum distance or distance of closest approach of the two molecules. The minimum distance problem can be defined as

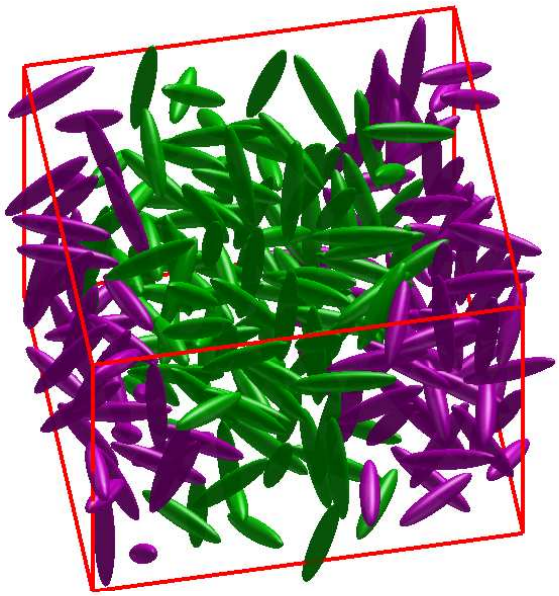


Figure 4.14: Simulation of a binary system of Lennard-Jones anisotropic molecules: step 20000. At this point the system is equilibrated and two distinct phases have formed, one rich in A type molecules and the other rich in B type molecules. Orientational correlation has completely disappeared but spatial correlation is now high. The system remains in a biphasic configuration for the remainder of the simulation.

finding $\mathbf{a} \in \mathcal{A}$ and $\mathbf{b} \in \mathcal{B}$ such that the distance $\| \mathbf{a} - \mathbf{b} \|_2$ is a minimum:

$$\begin{aligned} \text{minimise:} \quad & \| \mathbf{a} - \mathbf{b} \|_2, \\ \text{subject to:} \quad & \mathbf{a} \in \mathcal{A}, \quad \mathbf{b} \in \mathcal{B}. \end{aligned} \tag{4.49}$$

The arbitrary sets $\mathcal{A}$ and $\mathcal{B}$ can be defined using polynomial inequalities $g_i(\mathbf{a}) \geq 0,\ i = 1, \ldots, n_1$ and $g_j(\mathbf{b}) \geq 0,\ j = 1, \ldots, n_2$ where n_1 and n_2 are the number of constraints necessary to accurately represent $\mathcal{A}$ and $\mathcal{B}$ as closed sets (see figure 4.16). Finding the minimum distance of closest approach between $\mathcal{A}$ and $\mathcal{B}$ in this case can be formulated as a SOS optimisation problem, with a Positivstellensatz

$$\| \mathbf{a} - \mathbf{b} \|_2 - \sum_i s_i(\mathbf{a}, \mathbf{b}) g_i(\mathbf{a}) - \sum_j s_j(\mathbf{a}, \mathbf{b}) g_j(\mathbf{b}) \quad \text{is SOS}, \tag{4.50}$$

where $s_i(\mathbf{a}, \mathbf{b}),\ s_j(\mathbf{a}, \mathbf{b})$ are SOS polynomials in the components of $\mathbf{a}$ and $\mathbf{b}$ and there are i and j inequality constraints $g_i(\mathbf{a}) \geq 0,\ g_j(\mathbf{b}) \geq 0$ which define $\mathcal{A}$ and $\mathcal{B}$,

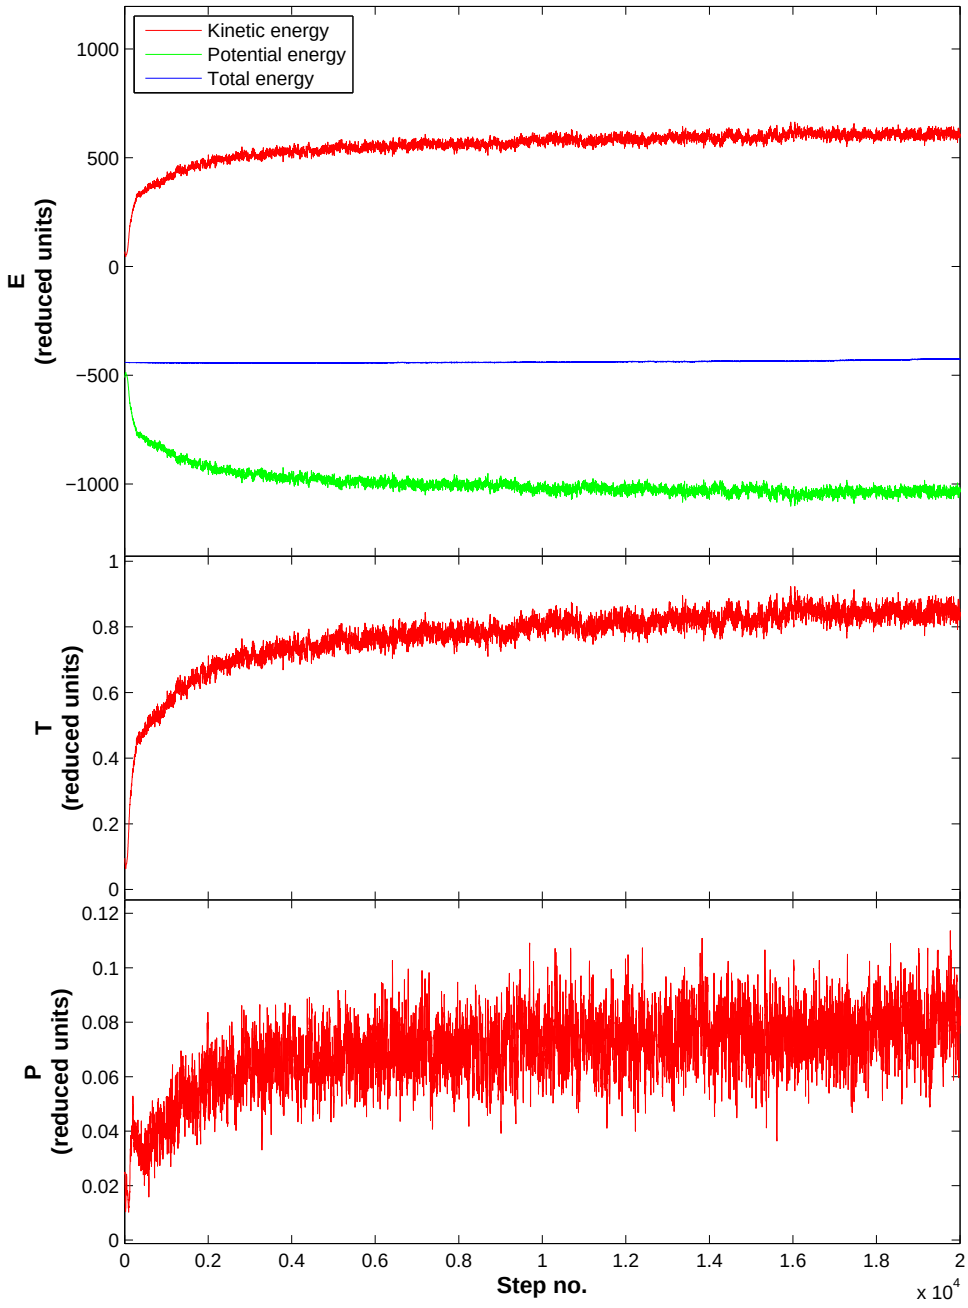


Figure 4.15: Trajectory data for the equilibration phase (20000 steps) of a simulation of a binary system of Lennard-Jones anisotropic molecules. The top graph shows kinetic, potential and total energy in reduced units. The overall conservation of total energy is quite good. The middle graph shows temperature in reduced units while the bottom graph shows pressure in reduced units.

respectively. The above problem (equation 4.49) can therefore be written as:

$$\begin{aligned} \text{max:} \quad & \lambda, \\ \text{such that:} \quad & \| \mathbf{a} - \mathbf{b} \|_2 - \sum_i s_i(\mathbf{a}, \mathbf{b}) g_i(\mathbf{a}) - \sum_j s_j(\mathbf{a}, \mathbf{b}) g_j(\mathbf{b}) - \lambda \quad \text{is SOS}, \end{aligned} \tag{4.51}$$

where λ is a real number. The minimum distance can now be obtained by solving a semidefinite programme. Duality theory provides a proof that the minimum distance between the two geometrical shapes has been obtained.

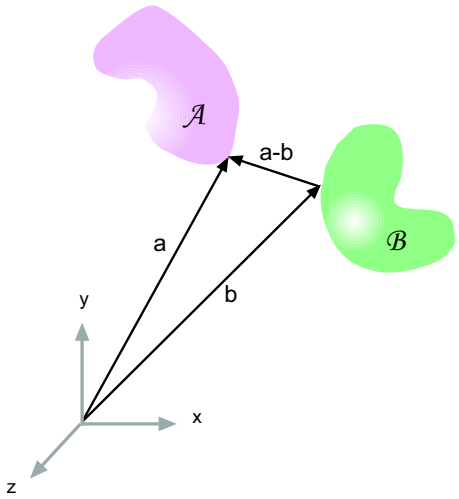


Figure 4.16: The minimum distance vector $\mathbf{a}-\mathbf{b}$ between two arbitrarily shaped particles $\mathcal{A}$ and $\mathcal{B}$ can be obtained in general by solving a SOS optimisation problem of the kind defined in equation 4.51. The particles need not be convex or even identical.

4.9 Discussion

The main objective of this work as outlined in section 4.2 was to demonstrate how the SOS approach could be used to calculate interaction energies between anisotropic coarse-grained particles in a molecular dynamics simulation. This goal has been achieved for ellipsoidal representations of molecules in an effective and potentially useful way. In addition an outline of how the approach could in future be extended even to non-convex bodies (given sufficient computational resources) was presented. However it is likely that this will require considerable development to become useful in a practical sense.

The formulation of the minimum distance problem as an SOCP however does have possibilities for future applications since ellipsoidal representations are already a topic of considerable interest within the simulation community. To the author's knowledge this is the first application of the minimum distance approach to calculating the intermolecular interactions in which the true minimum distance has been used. But the SOCP approach also has room for improvement. The most obvious area is the time take to compute the minimum distance to the required accuracy. Compared with the elliptic contact function approach, the SOCP approach is quite slow with the difference being an order of magnitude. However there is considerable scope for improving in the calculation times for the SOCP.

The first improvement that could be made is to remove the need for a preprocessing (presolve) phase in the SOCP call. This would involve removing any redundant constraints and reformulating the problem in a more efficient way. The presolve phase can take up to 20% of the solution time for a call to MOSEK. The second improvement that would be worth applying is to implement the SOCP optimisation on a cluster. MOSEK is fully configured for operating in a multiprocessor environment and the potential for increases in speed is significant. Another worthwhile improvement would be to investigate the possibility of using a "hot start" in MOSEK. This means the SOCP formulation remains in memory between calculations, with only the orientation matrices and centre to centre distance vectors being updated at each step. At present the entire SOCP is reformulated at each step.

The remainder of the molecular dynamics programme is quite fast (although not configured to benefit from a multiprocessor environment at present) since it is written in FORTRAN 77. However the subroutines as they stand have been written for readability and not for speed and therefore could be optimised for faster computation. These improvements may serve to bring the performance of the method significantly closer to that of the ECF method.

The SOCP approach has some advantages over the ECF approach since it calculates both the magnitude and direction of the true minimum distance vector accurately. Therefore it should provide a more realistic and physical representation of the dynamical behaviour of the system and therefore the ensemble sampled. The only time that the ECF approximation is exact is when the semi-axes of the two ellipsoids in question are parallel. In the case where there is no alignment of the semi-axes $d_{\text{true}} - d_R \to 0$ as $d_{\text{true}} \to 0$ and $d_{\text{true}} - d_R \to$constant when $d \to \infty$ [249]. The difference also degrades as s_x (the principal semi-axis) tends to infinity. How important this is in a practical sense remains to be seen. To establish this will require considerable testing of both the ECF and SOCP methods on various systems and a comparison made of thermodynamic data calculated from the two approaches with that obtained from experimental data. This is not a trivial task since the accuracy of the thermodynamic data obtained will depend strongly on the parameterisation of the Lennard-Jones functionals which at present is done in a relatively ad-hoc manner for coarse-grained particles. The advantages that the ECF approach has over the other single-site approaches also hold for the SOCP approach since they are both based on calculation of the force from the minimum distance vector.

An area that requires thorough investigation is the small increase in total energy. It may be that the time-step of 0.01 is a little too large for this particular set of simulation parameters. This can be examined by rerunning the simulation with a time-step of 0.005 for 80000 steps. The overall level of the energy conservation is a testament to the quality of the Omelyan integration algorithm for integrating the

angular velocities and to the accuracy of the MOSEK solver.

There are also issues which are completely unrelated to the use of the SOCP approach to calculate the minimum distance. The most significant of these is that the physical rationale for calculating the potential energy of interaction between two molecules as a pairwise functional based on the distance of closest approach is not very strong (see [250] for a discussion) and it can be argued that it neglects non-negligible interactions between non-adjacent parts of molecules. This criticism applies also to the ECF approach. However the ability of the minimum distance approach to reproduce behaviour seen using Gay-Berne type potentials without having to deal with some of the associated limitations means it will probably be a popular choice for practical applications given the widespread acceptance of the Gay-Berne potential within the liquid-crystal simulation community.

The approach described here can also be used in a Monte Carlo simulation if appropriate measures are taken to eliminate the possibility of particle overlap. Such measures have been developed and are standard for spherical particles and may be extendable to anisotropic particles.

It would be worthwhile at this stage to consider implementing standard interactions between spherical particles as well as sphere-ellipsoid type interactions. This would allow for the modelling of solvent based systems. The computational expense of this would be minimal since the majority of interactions would be sphere-sphere for which the interaction distance is based simply on the centres of mass. In addition sphere-ellipsoid interactions would be only half as expensive (in terms of the SOCP needing to be solved) as ellipsoid-ellipsoid interactions.

The implementation of rigid constraints between ellipsoids and between ellipsoids and spheres thereby allowing the representation of a single molecule by composite assemblies consisting of multiple ellipsoids and/or spheres would also be useful. This is particularly the case for modelling systems such as lipids where the zwitterionic head group of the molecule interacts with the solvent in a very different manner to the hydrocarbon tail. Therefore head-groups and tails need to be represented by different geometrical shapes with significantly different interaction parameters.

4.10 Conclusion

1. It has been demonstrated that the minimum distance problem in coarse grained molecular dynamics reduces to solving a SOCP. This has been implemented in a coarse grained molecular dynamics programme. To the author's knowledge this is the first such application in which the true minimum distance is employed.

2. The approach as it currently stands is considerably slower than the elliptic con-

tact function approach. However this is compensated by the fact that the true minimum distance is employed to calculate the intermolecular interactions.

3. A coarse grained molecular dynamics simulation employing the SOCP solver MOSEK to calculate the minimum distances has been shown to conserve energy with sufficient accuracy to complete a stable simulation of a binary system of anisotropic Lennard-Jones particles.

4. Possible avenues of further development include:

   (a) Elimination of the requirement for a presolve phase in the SOCP.

   (b) Evaluation of the benefits of running MOSEK in parallel on a linux cluster.

   (c) Investigate the possibility of using a "hot start" in MOSEK.

   (d) Implementation of a standard treatment of interactions between spherical molecules and treatment of sphere-ellipsoid interactions with a view to adding spherical solvent molecules.

   (e) Implementation of constraints to build composite molecules consisting of spheres and ellipsoids for lipid head-tail representations.

Appendices

Appendix A

How to obtain the global minimum and a SOS decomposition in practice

This section describes in detail the process for obtaining a SOS decomposition as well as the global minimum of the function from the SOS decomposition in order to demonstrate how the method works in practice. Note that the decomposition and optimisation are carried out simultaneously. A simple quartic polynomial function is chosen to exemplify the method by showing explicitly all the steps involved:

$$U = 4x^4 + 4x^3y + 5x^2y^2 + 2xy^3 + y^4 + 4y^2 - 4xy + x^2. \tag{A.1}$$

For the purposes of this example constraints are not considered. The global minimizer (0,0) is straightforward to obtain but the decomposition is not immediate by standard techniques. The emphasis here is in the implementation of the method in the available software (SOSTools).

1. The first step is to rewrite the problem as a polynomial optimization problem in the form of equation 1 in the flowchart of figure 2.5, where U is the function to be minimized subject to all the constraints:

$$\begin{aligned} \min: \quad & U(\mathbf{x}), \\ \text{subject to}: \quad & g(\mathbf{x}) \geq 0, \\ & h(\mathbf{x}) = 0. \end{aligned} \tag{A.2}$$

   In this example this is straightforward (U is given above and there are no constraints).

2. The next step is to introduce an additional real variable, $\lambda \in \mathbb{R}$, such that the

minimization problem becomes a maximization (equation 2, figure 2.5):

$$\begin{aligned} \max: \quad & \lambda, \\ \text{subject to}: \quad & U(\mathbf{x}) - \lambda \geq 0, \\ & g(\mathbf{x}) \geq 0, \\ & h(\mathbf{x}) = 0. \end{aligned} \tag{A.3}$$

Note that the function to be maximized now is a linear function while all the nonlinear functions are constraints. There is no effect on the underlying problem itself but the benefit is that solutions to the problem (optimal points) that may have been on the interior of the domain space are now on the surface. This is a key requirement for convex optimisation. The example now becomes:

$$\begin{aligned} \max: \quad & \lambda, \\ \text{subject to}: \quad & 4x^4 + 4x^3y + 5x^2y^2 + 2xy^3 + y^4 \\ & +4y^2 - 4xy + x^2 - \lambda \geq 0. \end{aligned} \tag{A.4}$$

3. The global nonnegativity constraint is now relaxed to a sum of squares constraint (see equation 3, figure 2.5),

$$\begin{aligned} \min: \quad & -\lambda, \\ \text{subject to}: \quad & 4x^4 + 4x^3y + 5x^2y^2 + 2xy^3 + y^4 \\ & +4y^2 - 4xy + x^2 - \lambda \quad \text{is SOS}. \end{aligned} \tag{A.5}$$

4. Once the problem is written in the form above, it can be submitted to the software SOSTools to find the lower bound and the sum of squares decomposition. To proceed, it is necessary to define the degree of the SOS basis to use. The maximum degree of the basis is d where $2d$ is the maximum degree of the function of interest including the constraints. (In this example, $d = 2$.) Now the entire vector of monomials with total degree less than d can be written:

$$\mathbf{z} = [x_1, x_2, \ldots, x_n, x_1^2, \ldots, x_n^2, x_1x_2, \ldots, x_{n-1}x_n, \ldots, x_1^d \ldots, x_n^d], \tag{A.6}$$

which in the example is :

$$\mathbf{z} = [1, x, y, x^2, xy, y^2]. \tag{A.7}$$

The ordering isn't particularly important but the ordering that allows easiest interpretation of results from SOSTools has been chosen here.

5. The basis of monomials allows the construction a matrix $\boldsymbol{Q}$. Finding the SOS decomposition is equivalent to finding a suitable positive semidefinite matrix,

$\mathbf{Q} \succeq 0$, such that :

$$U - \lambda = \mathbf{z}^T \mathbf{Q} \mathbf{z}. \tag{A.8}$$

This $\mathbf{Q}$ matrix is an $r \times r$ symmetric matrix in unknown variables q_{ij}. The problem now becomes (equation 4, figure 2.5)

$$\begin{aligned} \min: \quad & -\lambda \\ \text{subject to}: \quad & \mathbf{z}^T \mathbf{Q} \mathbf{z} \geq 0. \end{aligned} \tag{A.9}$$

So the SOS decomposition involves finding the entries of $\mathbf{Q}$ such that

$$U - \lambda = \begin{bmatrix} 1 \\ x \\ y \\ x^2 \\ xy \\ y^2 \end{bmatrix}^T \begin{bmatrix} q_{11} - \lambda & q_{12} & q_{13} & q_{14} & q_{15} & q_{16} \\ q_{12} & q_{22} & q_{23} & q_{24} & q_{25} & q_{26} \\ q_{13} & q_{23} & q_{33} & q_{34} & q_{35} & q_{36} \\ q_{14} & q_{24} & q_{34} & q_{44} & q_{45} & q_{46} \\ q_{15} & q_{25} & q_{35} & q_{45} & q_{55} & q_{56} \\ q_{16} & q_{26} & q_{36} & q_{46} & q_{56} & q_{66} \end{bmatrix} \begin{bmatrix} 1 \\ x \\ y \\ x^2 \\ xy \\ y^2 \end{bmatrix}. \tag{A.10}$$

and $\mathbf{Q}$ remains positive semidefinite. The full vector of monomials that can be expressed as a quadratic form in the basis $\mathbf{z}$ is

$$[1, x, x^2, x^3, x^4, y, xy, x^2y, x^3y, y^2, xy^2, x^2y^2, y^3, xy^3, y^4]. \tag{A.11}$$

6. The entries of $\mathbf{Q}$ and the coefficients of the monomials of the function of interest define a set of linear constraints (one for each entry in set A.11), given by the following set of linear equations:

$$\begin{array}{lllll} q_{11} - \lambda = 0, & q_{12} = 0, & q_{22} = 1, & q_{24} = 0, & q_{44} = 4, \\ q_{13} = 0, & 2q_{23} = -4, & q_{34} = 0, & q_{45} = 2, & q_{33} = 4, \\ q_{26} + q_{35} = 0, & 2q_{46} + q_{55} = 0, & q_{36} = 0, & q_{56} = 1, & q_{66} = 1. \end{array} \tag{A.12}$$

Therefore the problem can be rewritten as

$$\begin{aligned} \min: \quad & -\lambda \\ \text{subject to}: \quad & \mathbf{Tr}(\mathbf{A}_i, \mathbf{Q}) = b_i, \quad \forall i = 1, \ldots, 15, \\ & \mathbf{Q} \succeq 0, \end{aligned} \tag{A.13}$$

where the $\mathbf{A}_i$ are matrices representing the coefficients in the linear constraints.

For example, $\mathbf{A}_{12}$ is

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}, \tag{A.14}$$

and $b_{12} = 0$. The linear objective function to be minimised, the linear constraints and the constraint $\mathbf{Q} \succeq 0$ form a statement of a semidefinite programme. If $\mathbf{C}$ is defined to be a 6×6 matrix with zeros everywhere except a -1 in the $(1,1)$ entry and $\mathbf{b}$ as

$$b = [0, 0, 1, 0, 4, 0, -4, 0, 4, 4, 0, 0, 0, 2, 1],$$

(from the right hand sides of (A.12)), then this is a primal form SDP.

$$\begin{aligned} \min: \quad & \mathbf{Tr}(\mathbf{C}, \mathbf{Q}), \\ \text{subject to}: \quad & \mathbf{Tr}(\mathbf{A}_i, \mathbf{Q}) = b_i, \quad i = 1, \ldots, m, \\ & \mathbf{Q} \succeq 0. \end{aligned} \tag{A.15}$$

The dual form is formulated as:

$$\begin{aligned} \max: \quad & b^T y, \\ \text{subject to}: \quad & \mathbf{C} - \textstyle\sum_i \mathbf{A}_i y_i \succeq 0. \end{aligned} \tag{A.16}$$

SOSTools now invokes an SDP solver to obtain the solution. For this example, solutions are reported from SeDuMi. Once the SDP solver has completed its solution a number of checks must be made to ensure proper convergence. The checks are more or less standard for most SDP solvers such as SeDuMi or SDPT3.

(a) Check that the primal solution is feasible.

(b) Check that the dual solution is feasible.

(c) Check the duality gap.

In this example the primal solution is -1.5061×10^{-9}, the dual solution -1.0535×10^{-9} and the gap is -4.5263×10^{-10}.

7. Upon solving the semidefinite programme, a primal and dual solution matrix are produced. The primal matrix provides the lower bound on the global

minimum and the SOS decomposition while the dual matrix provides the minimiser. In this example, the solution matrices are

$$\text{Primal:} \begin{bmatrix} 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 4.0000 & 0.0000 & -2.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 1.0000 & 0.0000 & 1.0000 & 0.2460 \\ 0.0000 & -2.0000 & 0.0000 & 1.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 1.0000 & 0.0000 & 4.5080 & 2.0000 \\ 0.0000 & 0.0000 & 0.2460 & 0.0000 & 2.0000 & 4.0000 \end{bmatrix}, \quad \text{(A.17)}$$

$$\text{Dual:} \begin{bmatrix} 1.0000 & 0.0000 & -0.0000 & 0.0000 & -0.0000 & 0.0000 \\ 0.0000 & 0.0000 & -0.0000 & -0.0000 & -0.0000 & 0.0000 \\ -0.0000 & -0.0000 & 0.0000 & -0.0000 & 0.0000 & -0.0000 \\ 0.0000 & -0.0000 & -0.0000 & 0.0000 & -0.0000 & 0.0000 \\ -0.0000 & -0.0000 & 0.0000 & -0.0000 & 0.0000 & -0.0000 \\ 0.0000 & 0.0000 & -0.0000 & 0.0000 & -0.0000 & 0.0000 \end{bmatrix}. \quad \text{(A.18)}$$

From the primal solution matrix the global minimum of the function is $\lambda = q_{11} = 0.000$. Interpretation of the dual solution matrix is obtained by noting that the dual solution matrix in a sum of squares SDP solution should correspond to the matrix $\mathbf{z}\mathbf{z}^T$. The first row and the first column are equal to $\mathbf{z}$ provided the first entry in $\mathbf{z}$ is 1. Therefore the location of the global minimum is $(x, y) = (0, 0)$. It is important to check the rank of the dual solution matrix which should be 1 if the relaxation is exact.

8. Finding the SOS decomposition: To see the SOS decomposition, take the Cholesky factorisation of the $\mathbf{Q}$ matrix:

$$\boldsymbol{L} = \begin{bmatrix} 0.0003 & -0.0000 & -0.3140 & 0.0000 & -1.4840 & -1.6770 \\ 0.0000 & 2.0000 & -0.0000 & -1.0000 & -0.0000 & -0.0000 \\ 0.0000 & 0.0000 & 0.9494 & 0.0000 & 0.5625 & -0.2953 \\ 0.0000 & 0.0000 & 0.0000 & 0.0113 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 1.4103 & -0.2287 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 1.0238 \end{bmatrix}. \quad \text{(A.19)}$$

Each row contains a single polynomial within the SOS with each entry in each row being the coefficient of the corresponding entry in $\mathbf{z}$. So the SOS

decomposition is:

$$\begin{aligned} f^{SOS} = &(0.000256 - 0.31398y - 1.4840xy - 1.6770y^2)^2 + \\ &(2x - x^2)^2 + (0.94943y + 0.56249xy - 0.29530y^2)^2 + \\ &(0.0113x^2)^2 + (1.4103xy - 0.22875y^2)^2. + 1.0481y^4. \end{aligned} \tag{A.20}$$

The SOS decomposition is not unique and other decompositions may be also be found for example an alternative decomposition for this problem is

$$f^{SOS} = (2x^2 + 2xy + y^2)^2 + (2y - x)^2. \tag{A.21}$$

However the global minimum and location will be unique provided that this is the case for the energy function being minimised.

Appendix B

Interior point methods

With a view to having this thesis as self-contained as possible and bearing in mind that its intended audience are physical scientists, this appendix gives an overview of the main approaches for solving SDP's in sum of squares decompositions, namely interior point methods. The polynomial complexity bounds in interior point methods for semidefinite programming are a result of the rigourous characterisation of the algorithms which can only be understood by examining the mathematical foundations the approach. Although many SDP solvers are robust enough to be used in a black box approach, an understanding of the theory behind interior point methods is vital to appreciate both the polynomial complexity advantages of the approach as well as its practical limitations which were discussed in section 3.8. This description of interior point methods is based on the excellent books of Boyd and Vandenberghe [194] and Ben-Tal and Nemirovski [195]. By necessity it is very brief but should be sufficient for the non-specialist to appreciate the points discussed in section 3.8.

Algorithms for solving convex optimisation problems can be viewed as a hierarchy in terms of how they solve the KKT conditions. For the simplest problems, linear equality constrained quadratic optimisation problems, the KKT equations are a system of linear equations which can be solved analytically. Newton's method finds a local minimim for any linear equality constrained optimisation problem with twice continously differentiable objective. To do this it solves a sequence of linear equality constrained quadratic optimisation problems. It can also solve convex optimisation problems globally and is therefore next in the hierarchy. At the top of the hierarchy are interior point methods which can solve linear equality and inequality constrained optimisation problems by reducing them to a sequence of linear equality constrained optimisation problems.

Interior point (IP) methods are currently the algorithm of choice for solving both SOCP's and SDP's. This is primarily due to their provable polnomial efficiency and accuracy. Alternative approaches are currently the subject of considerable ongoing research within the optimisation community but to date none have been able

to match both the practical performance (for small-medium problems) or provable complexity bounds of IP methods. They are of course not perfect and one considerable disadvantage in using them to solve semidefinite programmes is related to the size of the problem that they can deal with.

B.1 Unconstrained minimisation

In the initial discussion the problem being focussed on is an unconstrainted minimisation where the objective $f(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}$ is a continuously twice differentiable convex function. The types of algorithm being considered are iterative in the sense that a sequence of points called a minimising sequence begining at some initial point $f(\mathbf{x}^{(0)})$ converges to p^* as the number of iterations $k \to \infty$.

Example

Analytic centre of a Linear matrix inequality Consider the optimisation problem

$$\text{minimise } f(\mathbf{x}) = \log\det \mathbf{F}(\mathbf{x})^{-1}, \tag{B.1}$$

where $\mathbf{F} : \mathbb{R}^n \to \mathbb{S}^p$ is an affine function of the form

$$\mathbf{F}(\mathbf{x}) = \mathbf{F}_0 + x_1\mathbf{F}_1 + \cdots + x_n\mathbf{F}_n, \tag{B.2}$$

and $\mathbf{dom} f = \{\mathbf{x} \mid \mathbf{F}(\mathbf{x}) \succeq 0\}$. The function f is the logarithmic barrier for the linear matrix inequality $\mathbf{F}(\mathbf{x}) \succeq 0$ and plays a very important role in the application of interior point methods to SDP's.

B.1.1 Methods for descending

Steepest descent direction in the Euclidean norm

The first order Taylor approximation of $f(\mathbf{x}+\mathbf{v})$ around $\mathbf{x}$ is given by

$$f(\mathbf{x}+\mathbf{v}) \approx f(\mathbf{x}) + \nabla f(\mathbf{x})^T\mathbf{v}, \tag{B.3}$$

where $\nabla f(\mathbf{x})^T\mathbf{v}$ is the directional derivative of f at $\mathbf{x}$ in the direction of $\mathbf{v}$. If the directional derivative is negative then $\mathbf{v}$ is a descent direction. Because the directional derivative is linear in $\mathbf{v}$ it can be made to approach $-\infty$ by making $\mathbf{v}$ very large. Therefore a normalised steepest descent direction with repect to a

particular norm $\|\cdot\|$ is defined as

$$\Delta \mathbf{x}_{nsd} = \text{argmin}\{\nabla f(\mathbf{x})^T \mathbf{v} \mid \|\mathbf{v}\| \leq 1\}. \tag{B.4}$$

If the norm $\|\cdot\|$ is taken to be the Euclidean norm, the steepest descent direction is the negative of the gradient.

Steepest descent direction in the quadratic norm

Instead consider the quadratic norm

$$\|\mathbf{z}\|_{\mathbf{P}} = (\mathbf{z}^T \mathbf{P} \mathbf{z})^{1/2} = \|\mathbf{P}^{1/2} \mathbf{z}\|_2, \tag{B.5}$$

where $\mathbf{P} \in \mathbb{S}^n_{++}$. The normalised steepest descent direction is given by

$$\Delta \mathbf{x}_{nsd} = -(\nabla f(\mathbf{x})^T \mathbf{P}^{-1} \nabla f(\mathbf{x}))^{-1/2} \mathbf{P}^{-1} \nabla f(\mathbf{x}). \tag{B.6}$$

The dual norm is given by $\|\mathbf{P}^{-1/2}\mathbf{z}\|_2$, so the steepest descent step is

$$\Delta \mathbf{x}_{sd} = -\mathbf{P}^{-1} \nabla f(\mathbf{x}). \tag{B.7}$$

The choice of norm is important because it can dramatically impact the convergence rate of an iterative opimtisation algorithm. For example in the case of a quadratic norm the choice of $\mathbf{P}$ should be made such that the sub-level sets of f are well conditioned when transformed by $\mathbf{P}^{-1/2}$. Another very good choice is an approximation $\hat{\mathbf{H}}$ to the Hessian at the optimal point p^* if it is known since

$$\hat{\mathbf{H}}^{-1/2} \nabla^2 f(\mathbf{x}^*) \hat{\mathbf{H}}^{-1/2} \approx \mathbf{I}, \tag{B.8}$$

and will therefore have a very low condition number.

B.2 Newton's method for unconstrained minimisation

B.2.1 The Newton step

The Newton step is defined as

$$\Delta \mathbf{x}_{nt} = -\nabla^2 f(\mathbf{x})^{-1} \nabla f(\mathbf{x}) \tag{B.9}$$

for $\mathbf{x} \in \mathbf{dom} f$. The Newton step can be interpreted as the steepest descent direction at $\mathbf{x}$, for the quadratic norm defined by the Hessian $\nabla^2 f(\mathbf{x})$.

The Newton decrement The Newton decrement $\lambda(\mathbf{x})$ at $\mathbf{x}$ is defined by $f(\mathbf{x}) - \hat{f}(\mathbf{x} + \Delta\mathbf{x}_{nt}) = (1/2)\lambda(\mathbf{x})^2$, where $\hat{f}$ is the quadratic approximation of f at $\mathbf{x}$ and therefore λ is useful as a stopping criterion. It is also defined as

$$\begin{aligned} \lambda(\mathbf{x}) &= (\nabla f(\mathbf{x})^T \nabla^2 f(\mathbf{x})^{-1} \nabla f(\mathbf{x}))^{1/2} \\ &= (\Delta\mathbf{x}_{nt} \nabla^2 f(\mathbf{x})^{-1} \Delta\mathbf{x}_{nt})^{1/2}. \end{aligned} \tag{B.10}$$

λ is therefore the norm of the Newton step defined in the Hessian norm.

Algorithm B.2.1 (Newton's method) *Given a starting point* $\mathbf{x} \in \mathbf{dom} f$*, and a tolerence* $\epsilon > 0$.
Repeat:

1. *Calculate Newton step.*

$$\Delta\mathbf{x}_{nt} = -\nabla^2 f(\mathbf{x})^{-1} \nabla f(\mathbf{x}). \tag{B.11}$$

2. *Calculate the Newton decrement*

$$\lambda(\mathbf{x}) = (\nabla f(\mathbf{x})^T \nabla^2 f(\mathbf{x})^{-1} \nabla f(\mathbf{x}))^{1/2}. \tag{B.12}$$

3. *Stop if* $\lambda^2/2 \leq \epsilon$.

4. *Line search: Choose step size* t *by backtracking linesearch.*

5. *Update* $\mathbf{x} = \mathbf{x} + t\Delta\mathbf{x}_{nt}$.

The computation of the Newton step is the most expensive calculation in the algorithm. First the Hessian $\mathbf{H} := \nabla^2 f(\mathbf{x})$ and the gradient $\mathbf{g} = \nabla f(\mathbf{x})$ at $\mathbf{x}$ are computed. Then a system of linear equations (the Newton system) $\mathbf{H}\Delta\mathbf{x}_{nt} = -\mathbf{g}$ is solved for $\Delta\mathbf{x}_{nt}$. Usually since the Hessian is symmetric positive definite, a Cholesky factorisation of $\mathbf{H}$ is used to compute the lower triangular matrix $\mathbf{L}$ such that $\mathbf{L}\mathbf{L}^T = \mathbf{H}$. Then $\mathbf{L}\omega = -\mathbf{g}$ is solved by forward substitution giving $-\mathbf{L}^{-1}\mathbf{g}$. Finally $\mathbf{L}^T\Delta\mathbf{x}_{nt} = \omega$ is solved by back substitution. The Newton decrement is calculated as $\lambda^2 = \|\omega\|_2^2$. The cost for such a computation is dominated by the Cholesky factorisation ($(1/3)n^3$ flops) for a dense unstructured Hessian. Although it is often possible to improve on this computational cost if the Hessian is banded or has sparse structure, this is not usually the case, in for example, the solution of semidefinite programmes.

B.3 Newton's method for equality constrained minimisation

Here the problem under consideration is

$$\begin{aligned} \text{minimise:} \quad & f(\mathbf{x}) \\ \text{subject to:} \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \end{aligned} \tag{B.13}$$

The method is very similar for equality constrained optimisation except for two differences. First the initial point must be feasible and secondly the definition of the Newton step is modified to ensure that it is in a feasible direction, i.e. $\mathbf{A}\Delta\mathbf{x}_{nt} = \mathbf{0}$.

The Newton step for the equality constrained optimisation problem is derived by replacing the objective with its second order Taylor approximation near $\mathbf{x}$

$$\begin{aligned} \text{minimise:} \quad & \hat{f}(\mathbf{x}+\mathbf{v}) = f(\mathbf{x}) + \nabla f(\mathbf{x})^T\mathbf{v} + 0.5\mathbf{v}^T\nabla^2 f(\mathbf{x})\mathbf{v} \\ \text{subject to:} \quad & \mathbf{A}(\mathbf{x}+\mathbf{v}) = \mathbf{b} \end{aligned} \tag{B.14}$$

which is a convex quadratic optimisation problem.

Linearly constrained convex quadratic optimisation problems For a linearly constrained convex quadratic optimisation problem

$$\begin{aligned} \text{minimise:} \quad & 0.5\mathbf{x}^T\mathbf{P}\mathbf{x} + \mathbf{q}^T\mathbf{x} + r \\ \text{subject to:} \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \end{aligned} \tag{B.15}$$

where $\mathbf{P} \in \mathbb{S}^n_+$ and $\mathbf{A} \in \mathbb{R}^{p\times n}$, the KKT optimality conditions can be written as

$$\begin{bmatrix} \mathbf{P} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x}^* \\ \boldsymbol{\nu}^* \end{bmatrix} = \begin{bmatrix} -\mathbf{q} \\ \mathbf{b} \end{bmatrix}. \tag{B.16}$$

This is a set of $n+p$ linear equations called the KKT system and the coefficient matrix is often called the KKT matrix. If the KKT matrix is nonsingular, there is a unique optimal primal dual pair $(\mathbf{x}^*, \boldsymbol{\nu}^*)$. If the KKT system is unsolvable the quadratic optimisation problem is unbounded below. If it is non-singular then any solution yeilds an optimal primal-dual pair.

B.3.1 The Newton step

The Newton step $(\Delta\mathbf{x}_{nt})$ for the second-order Taylor expansion problem is therefore defined by

$$\begin{bmatrix} \nabla^2 f(\mathbf{x}) & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x}_{nt} \\ \boldsymbol{\omega} \end{bmatrix} = \begin{bmatrix} -\nabla f(\mathbf{x}) \\ \mathbf{0} \end{bmatrix}, \tag{B.17}$$

where $\boldsymbol{\omega}$ is the dual optimal variable for the quadratic problem. This can be interpreted as a solution to the linearised KKT conditions.

The Newton decrement The Newton decrement is exactly the same as in the unconstrained case (B.10) and so remains the norm of the Newton step defined in the Norm of the Hessian.

B.3.2 Infeasible starts

Sometimes it is difficult to have an initial point that is feasible and therefore Newton's method as it stands will fail. This can be overcome by writing the KKT system as

$$\begin{bmatrix} \nabla^2 f(\mathbf{x}) & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x}_{nt} \\ \boldsymbol{\omega} \end{bmatrix} = - \begin{bmatrix} \nabla f(\mathbf{x}) \\ \mathbf{A}\mathbf{x} - \mathbf{b} \end{bmatrix}. \tag{B.18}$$

The second block in the right-hand-side vector is known as the residual and when $\mathbf{x}$ is feasible the residual vanishes and (B.18) becomes (B.17). Interestingly the infeasible Newton step need not be a descent direction. However it can be shown that if a step length of one is taken using the infeasible Newton step, the following iterate will be feasible. Once a feasible point has been reached all further steps will be feasible.

An appealing alternative to using an infeasible start method for an SOCP or SDP is to use a self-dual embedding technique [251]. This will not be discussed here but it is of practical importance in conic programming software packages such as SeDuMi [211] and MOSEK [252].

Example

Analytic centre of a linear matrix inequality Given the problem

$$\begin{aligned} \text{minimise:} \quad & f(\mathbf{X}) = -\log(\det \mathbf{X}) \\ \text{subject to:} \quad & \mathbf{Tr}(\mathbf{A}_i\mathbf{X}) = b_i, \quad i = 1, \ldots, p, \end{aligned} \tag{B.19}$$

where $\mathbf{X}, \mathbf{A}_i \in \mathbb{S}^n, b_i \in \mathbb{R}$. The KKT conditions are

$$-\mathbf{X}^{*-1} + \sum_{i=1}^{m} \nu_i^* \mathbf{A}_i = 0 \tag{B.20}$$

$$\mathbf{Tr}(\mathbf{A}_i\mathbf{X}^*) = b_i, \quad i = 1, \ldots, p. \tag{B.21}$$

To derive the KKT system for $\Delta\mathbf{X}_{nt}$, $\mathbf{X}^*$ is replaced by $\mathbf{X}+\Delta\mathbf{X}_{nt}$ and $\boldsymbol{\nu}^*$ is replaced by $\boldsymbol{\omega}$. The first equation in the KKT system is then linearised through a first order

approximation

$$(\mathbf{X} + \Delta\mathbf{X}_{nt})^{-1} \approx \mathbf{X}^{-1} - \mathbf{X}^{-1}\Delta\mathbf{X}_{nt}\mathbf{X}^{-1}. \tag{B.22}$$

This leads to the KKT system

$$\begin{aligned} -\mathbf{X}^{-1} + \mathbf{X}^{-1}\Delta\mathbf{X}_{nt}\mathbf{X}^{-1} + \sum_{i=1}^{p} \omega_i \mathbf{A}_i = 0, \\ \mathbf{Tr}(\mathbf{A}_i \Delta\mathbf{X}_{nt}) = 0, \quad i = 1, \ldots, p. \end{aligned} \tag{B.23}$$

B.4 Interior point methods

Interior point methods are used to solve equality and inequality constrained convex optimisation problems where the objective function is twice continuously differentiable. What is new here is how to deal with the inequality constraints. The idea is to make the inequality constraints implicit in the objective function. To do this a function known as a logarithmic barrier function is employed which behaves like an indicator function, i.e. is zero when the inequality constraints are satisfied and infinity when they are not. The logarithmic barrier function in general is defined as

$$\phi(\mathbf{x}) = -\sum_{i=1}^{m} \log(-f_i(\mathbf{x})). \tag{B.24}$$

Using this logarithmic barrier problem B.13 can be written as

$$\begin{aligned} \text{minimise:} \quad & f_0(\mathbf{x}) + (1/t)\phi(\mathbf{x}), \\ \text{subject to:} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{aligned} \tag{B.25}$$

where t is a parameter. This is only an approximation of a true indicator function but the approximation improves as the parameter t grows. However when t is large the Hessian varies very dramatically near the boundary of the feasible set which is where the solution is guaranteed to lie. This difficulty can be overcome by solving a sequence of problems with t increasing in each problem. The initial point in each case is the solution point of the previous value of t.

B.4.1 The central path

At this stage is it useful to rewrite (B.25) as

$$\begin{aligned} \text{minimise:} \quad & t f_0(\mathbf{x}) + \phi(\mathbf{x}) \\ \text{subject to:} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned} \tag{B.26}$$

The set of points $\mathbf{x}^*(t)$, $t > 0$ define the central path. The central path is fundamental to interior point methods as an interior point optimisation will trace the central path to reach the optimal point of the problem to be solved. Points on the central path are defined by the neccessary and sufficient conditions

$$\mathbf{A}\mathbf{x}^*(t) = \mathbf{b}, \quad f_i(\mathbf{x}^*(t)) < 0, \quad i = 1, \ldots, m, \tag{B.27}$$

and that there exists $\hat{\boldsymbol{\nu}} \in \mathbb{R}^p$ such that

$$t\nabla f_0(\mathbf{x}^*(t)) + \nabla\phi(\mathbf{x}^*(t)) + \mathbf{A}^T\hat{\boldsymbol{\nu}} = 0. \tag{B.28}$$

The gradient of a barrier function is given by

$$\nabla\phi(\mathbf{x}) = \sum_{i=1}^{m} \frac{1}{-f_i(\mathbf{x})}\nabla f_i(\mathbf{x}). \tag{B.29}$$

Hence (B.28) becomes,

$$t\nabla f_0(\mathbf{x}^*(t)) + \sum_{i=1}^{m} \frac{1}{-f_i(\mathbf{x}^*(t))}\nabla f_i(\mathbf{x}^*(t)) + \mathbf{A}^T\hat{\boldsymbol{\nu}} = 0. \tag{B.30}$$

Every point on the central path has a dual feasible point which provides a lower bound on p^*. By defining

$$\lambda_i^*(t) = -\frac{1}{tf_i(\mathbf{x}^*(t))}, \quad i = 1, \ldots, m, \quad \boldsymbol{\nu}^*(t) = \hat{\boldsymbol{\nu}}/t, \tag{B.31}$$

equation (B.30) can be written as

$$\nabla f_0(\mathbf{x}^*(T)) + \sum_{i=1}^{m} \lambda_i^*(t)\nabla f_i(\mathbf{x}^*(t)) + \mathbf{A}^T\boldsymbol{\nu}^*(t) \;=\; 0. \tag{B.32}$$

The Lagrangian for the central path is

$$L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = f_0(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i f_i(\mathbf{x}) + \boldsymbol{\nu}^T(\mathbf{A}\mathbf{x} - \mathbf{b}), \tag{B.33}$$

for $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*(t)$ and $\boldsymbol{\nu} = \boldsymbol{\nu}^*(t)$. The dual function is therefore

$$\begin{aligned} g(\boldsymbol{\lambda}^*(t), \boldsymbol{\nu}^*(t)) &= f_0(\mathbf{x}^*(t)) + \textstyle\sum_{i=1}^{m} \lambda_i^*(t) f_i(\mathbf{x}^*(t)) + \boldsymbol{\nu}^*(t)^T(\mathbf{A}\mathbf{x}^*(t) - \mathbf{b}) \\ &= f_0(\mathbf{x}^*(t)) - m/t. \end{aligned} \tag{B.34}$$

As a result the duality gap at $(\mathbf{x}^*(t), \boldsymbol{\lambda}^*(t), \boldsymbol{\nu}^*(t))$ is m/t, i.e.

$$f_0(\mathbf{x}^*(t)) - p^* \leq m/t, \tag{B.35}$$

and $\mathbf{x}^*(t)$ is no more than m/t suboptimal. This is quite important as it provides an upper bound on the duality gap which depends only on the number of inequality constraints and the parametric position t on the central path.

Points on the central path are characterised by the following centrality conditions

$$\begin{aligned} f_i(\mathbf{x}) &\leq 0 \quad i = 1, \ldots, m, \\ \mathbf{Ax} &= \mathbf{b}, \\ \boldsymbol{\lambda} &\geq \mathbf{0}, \\ \nabla f_0(\mathbf{x}) + \textstyle\sum_{i=1}^m \lambda_i \nabla f_i(\mathbf{x}) + \mathbf{A}^T\boldsymbol{\nu} &= 0, \\ -\lambda_i f_i(\mathbf{x}) &= 1/t, \quad i = 1, \ldots, m. \end{aligned} \tag{B.36}$$

The centrality conditions are almost identical to the KKT conditions with the exception that the complementary slackness condition has been augmented by $1/t$.

B.4.2 The path-following method

By solving a sequence of linearly constrained minimisation problems, using the last point as the starting point for the next minimisation it is possible to compute $\mathbf{x}^*(t)$ for a sequence of increasing values of t, until $t \leq m/\epsilon$ guaranteeing that an ϵ-suboptimal solution has been obtained.

Algorithm B.4.1 (Path-following method) *Given a strictly feasible* $\mathbf{x}$, $t^{(0)} > 0, \mu > 1$, *and* $\epsilon > 0$
Repeat:

1. *Centering Step: Minimise* $tf_0 + \phi$ *subject to* $\mathbf{Ax} = \mathbf{b}$ *starting at* $\mathbf{x}$ *to compute* $\mathbf{x}^*(t)$.
2. *Update* $\mathbf{x} = \mathbf{x}^*(t)$.
3. *Stop if* $m/t < \epsilon$.
4. $t = \mu t$.

Step 1 is refered to as the centering step or iteration. The Newton iterations within step 1 are referred to as inner iterations. Note that computing $\mathbf{x}^*(t)$ exactly is not necessary except at the final centering step. The choice of μ is a trade off between the number of centering steps versus the number of inner iterations required. A small μ will result in a small number of inner iterations per centering step, while the opposite is true for larger steps. In the large step case the solution path can

vary quite considerably from the central path, however in practical terms this make relatively little difference.

In the path-following method, the Newton step $\Delta\mathbf{x}_{nt}$ and its associated dual variable are given by the linear equations

$$\begin{bmatrix} t\nabla^2 f_0(\mathbf{x}) + \nabla^2\phi(\mathbf{x}) & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x}_{nt} \\ \boldsymbol{\nu}_{nt} \end{bmatrix} = - \begin{bmatrix} t\nabla f_0(\mathbf{x}) + \nabla\phi(\mathbf{x}) \\ \mathbf{0} \end{bmatrix}. \tag{B.37}$$

These linear equations can be seen as defining the Newton steps for the modified KKT equations

$$\begin{aligned} \nabla f_0(\mathbf{x}) + \textstyle\sum_{i=1}^m \lambda_i \nabla f_i(\mathbf{x}) + \mathbf{A}^T\boldsymbol{\nu} &= 0, \\ -\lambda_i f_i(\mathbf{x}) &= 1/t, \quad i = 1, \ldots, m, \\ \mathbf{A}\mathbf{x} &= \mathbf{b}. \end{aligned} \tag{B.38}$$

Central path methods require a feasible starting point but often a feasible point is not known. In this case a variant known as a Phase 1 method is employed which uses an infeasible start Newton method. This will not be described here for the sake of brevity but the basic concept is that the following problem is solved instead

$$\begin{aligned} \text{minimise:} \quad & s \\ \text{subject to:} \quad & f_i(\mathbf{x}) \leq s, \quad i = 1, \ldots, m \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{aligned} \tag{B.39}$$

where $s \in \mathbb{R}$.

B.4.3 Logarithmic barriers for generalised inequalities

$\psi : \mathbb{R}^q \to \mathbb{R}$ is a generalised logarithm for a proper cone K if

1. ψ is concave, closed, twice continuously differentiable, $\mathbf{dom}\psi = \mathbf{int}K$ and $\nabla^2\psi(\mathbf{y}) \prec 0$ for $\mathbf{y} \in \mathbf{int}K$.

2. There is a constant $\theta > 0$ such that for all $\mathbf{y} \succ_K 0$ and all $s > 0$,

$$\psi(s\mathbf{y}) = \psi(\mathbf{y}) + \theta \log s. \tag{B.40}$$

θ is often referred to as the degree of ψ.

Examples

Logarithmic barrier for the second order cone

$$\psi = \log\left(x_{n+1}^2 - \sum_{i=1}^{n} x_i^2\right) \tag{B.41}$$

is the generalised logarithm for the second order cone (2.4). The gradient of ψ is given as

$$\frac{\partial\psi(\mathbf{x})}{\partial x_j} = \frac{-2x_j}{\left(x_{n+1}^2 - \sum_{i=1}^{n} x_i^2\right)}, \quad j = 1, \ldots, n, \tag{B.42}$$

$$\frac{\partial\psi(\mathbf{x})}{\partial x_{n+1}} = \frac{-2x_{n+1}}{\left(x_{n+1}^2 - \sum_{i=1}^{n} x_i^2\right)}. \tag{B.43}$$

Logarithmic barrier for the positive semidefinite cone

$$\psi(\mathbf{x}) = \log(\det \mathbf{X}) \tag{B.44}$$

is the generalised logarithm for the positive semidefinite cone $\mathbb{S}_{++}^p$ and the gradient of ψ at a point $\mathbf{X} \in \mathbb{S}_{++}^p$ is

$$\nabla\psi(\mathbf{X}) = \mathbf{X}^{-1}. \tag{B.45}$$

The central path The key properties of the central path generalise to the case of generalised inequalities,

1. To find a point on the central path it is necessary to minimize a twice differentiable function subject to equality constraints which can be done via Newton's method.

2. Each point on the central path $\mathbf{x}^*(t)$ has an associated dual feasible point $(\boldsymbol{\lambda}^*(t), \boldsymbol{\nu}^*(t))$ and the corresponding duality gap is $\bar{\theta}/t$, where $\bar{\theta} = \sum_{i=1}^{m} \theta_i$. Hence $\mathbf{x}^*(t)$ is no more than $\bar{\theta}/t$ suboptimal.

Example

The central path for a semidefinite programme For the inequality form SDP (2.45) and its corresponding dual problem (2.48), the barrier function is given by

$$\phi(\mathbf{x}) = \log\left(\det(-\mathbf{F}(\mathbf{x})^{-1})\right), \tag{B.46}$$

with $\mathbf{dom}\phi = \{\mathbf{x} \mid \mathbf{F}(\mathbf{x}) \preceq 0\}$. The gradient $\nabla\phi(\mathbf{x})$ is

$$\frac{\partial\phi(\mathbf{x})}{\partial x_i} = \mathbf{Tr}(\mathbf{F}(\mathbf{x})^{-1}\mathbf{F}_i), \quad i = 1, \ldots, n, \tag{B.47}$$

so the optimality condition which defines the central path can be written as

$$tc_i + \mathbf{Tr}(\mathbf{F}(\mathbf{x}^*(t))^{-1}\mathbf{F}_i) = 0, \quad i = 1, \ldots, n \tag{B.48}$$

Therefore the matrix

$$\mathbf{Z}^*(t) = \frac{1}{t}(-\mathbf{F}(\mathbf{x}^*(t)))^{-1} \tag{B.49}$$

is strictly dual feasible and the duality gap associated with $\mathbf{X}^*(t)$ and $\mathbf{Z}^*(t)$ is p/t.

B.4.4 Primal-dual interior point methods

Path-following methods solve the primal problem only. The dual variable is only included in the Newton system to ensure that the Newton step be in a feasible direction. The approach taken in primal-dual interior point methods (PDIPMs) is somewhat different. Although like path-following methods, primal-dual interior point methods use the central path concept, the primal and dual problems are given equal importance in the Newton system. In addition no distinction is made between inner and outer iterations. Another important difference is that in the PDIPM a given point at an intermediate iterate $(\mathbf{x}(t), \boldsymbol{\lambda}(t), \boldsymbol{\nu}(t))$ is not guananteed to be feasible. To start with the modified KKT conditions (B.38) are expressed as $\boldsymbol{r}_t(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = 0$ where

$$\boldsymbol{r}_t(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = \begin{bmatrix} \nabla f_0(\mathbf{x}) + Df(\mathbf{x})^T\boldsymbol{\lambda} + \mathbf{A}^T\boldsymbol{\nu} \\ -\mathbf{diag}(\boldsymbol{\lambda})f(\mathbf{x}) - (1/t)\mathbf{1} \\ \mathbf{Ax} - \mathbf{b} \end{bmatrix} \tag{B.50}$$

where $t > 0$ and $f : \mathbb{R}^n \to \mathbb{R}^m$ and Df are defined as

$$f(\mathbf{x}) = \begin{bmatrix} f_1(\mathbf{x}) \\ \vdots \\ f_m(\mathbf{x}) \end{bmatrix}, \qquad Df(\mathbf{x}) = \begin{bmatrix} \nabla f_1(\mathbf{x}) \\ \vdots \\ \nabla f_m(\mathbf{x}) \end{bmatrix}. \tag{B.51}$$

The first block of $\boldsymbol{r}_t(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu})$ is called the dual residual ($\boldsymbol{r}_{dual}$), the second block is called the centrality residual ($\boldsymbol{r}_{cent}$) and the third block is the primal residual ($\boldsymbol{r}_{pri}$) with these terms being self-explanatory.

The next step is to produce the Newton step for solving $\boldsymbol{r}_t(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = 0$ for a given t at a point $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu})$ for $f(\mathbf{x}) \prec 0$, $\boldsymbol{\lambda} \succ 0$. The current point is denoted

$\mathbf{y} = (\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu})$ and the Newton step is defined as $\Delta\mathbf{y} = (\Delta\mathbf{x}, \Delta\boldsymbol{\lambda}, \Delta\boldsymbol{\nu})$ and is given by the following set of linear equations

$$\boldsymbol{r}_t(\mathbf{y} + \Delta\mathbf{y}) \approx \boldsymbol{r}_t(\mathbf{y}) + D\boldsymbol{r}_t(\mathbf{y})\Delta\mathbf{y} = 0, \tag{B.52}$$

i.e. $\Delta\mathbf{y} = -D\boldsymbol{r}_t(\mathbf{y})^{-1}\boldsymbol{r}_t(\mathbf{y})$. The KKT system expressed in terms of $\mathbf{x}, \boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ is

$$\begin{bmatrix} \nabla^2 f_0(\mathbf{x}) + \sum_{i=1}^m \lambda_i \nabla^2 f_i(\mathbf{x}) & Df(\mathbf{x})^T & \mathbf{A}^T \\ -\mathbf{diag}(\boldsymbol{\lambda})Df(\mathbf{x}) & -\mathbf{diag}(f(\mathbf{x})) & \mathbf{0} \\ \mathbf{A} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x} \\ \Delta\boldsymbol{\lambda} \\ \Delta\boldsymbol{\nu} \end{bmatrix} = - \begin{bmatrix} \boldsymbol{r}_{dual} \\ \boldsymbol{r}_{cent} \\ \boldsymbol{r}_{pri} \end{bmatrix}, \tag{B.53}$$

and the primal-dual search direction is defined as the solution. The primal and dual search directions are coupled, through both the coefficient matrix and the residuals. Note that once the primal feasibility residual is zero then $\mathbf{A}\Delta\mathbf{x}_{pd} = \mathbf{0}$ and $\Delta\mathbf{x}_{pd}$ defines a primal feasible direction.

Because in the PDIPM, the iterates $\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}, \boldsymbol{\nu}^{(k)}$ are not guaranteed to be feasible except in the limit where the algorithm converges, it is not always possible to direcly evaluate a duality gap $\eta^{(k)}$ at step k as in the path-following method. Instead the "surrogate" duality gap which would be the duality gap if $\mathbf{x}$ were primal feasible ($\boldsymbol{r}_{pri} = 0$) and $\boldsymbol{\lambda}, \boldsymbol{\nu}$ were dual feasible ($\boldsymbol{r}_{dual} = 0$) is used

$$\hat{\eta}(\mathbf{x}, \lambda) = -f(\mathbf{x})^T\lambda. \tag{B.54}$$

Algorithm B.4.2 (Primal-dual interior point method) *given* $\mathbf{x}$ *that satisfies* $f_1(\mathbf{x}) < 0, \ldots, f_m(\mathbf{x}) < 0, \boldsymbol{\lambda} \succeq 0, \mu > 1, \epsilon_{feas} > 0, \epsilon > 0$.
Repeat:

1. *Set* $t := \mu m/\hat{\eta}$.
2. *Solve the primal-dual KKT system (B.53) for the primal-dual search direction* $\Delta\mathbf{y}_{pd}$.
3. *Determine step length* $s > 0$ *and set* $\mathbf{y} := \mathbf{y} + s\Delta\mathbf{y}_{pd}$. *(Line search)*
4. *Stop if* $\|\boldsymbol{r}_{pri}\|_2 \leq \epsilon_{feas}, \|\boldsymbol{r}_{dual}\|_2 \leq \epsilon_{feas}$ *and* $\hat{\eta} \leq \epsilon$.

In step 1 t is defined using the surrogate duality gap $\hat{\eta}$. If $\mathbf{x}, \boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ were central, the update would be exactly the same as in the path-following method since t would be increased by a factor of μ. The line search procedure is a standard backtracking procedure.

B.5 Primal-dual interior point methods applied to SDP

There are a variety of algorithms for solving SDP's which fall under the umbrella of PDIPM's but which can produce quite different performance on different categories of SDP. The main difference in the construction of these algorithms is related to how they calculate the Newton search directions. It is important to note that the advantage gained from using a patricular search direction may not only come from inceased efficiency in the practical sense but also from improvements in the provable polynomial complexity bounds for the algorithm.

However it is difficult to understand what differentiates these different search directions by looking at the SDP in the inequality representation given in (2.45) and (2.48). The disadvantage of looking at the SDP problem in this form is that the geometrical symmetry which is inherent in the problem because the cones are self-dual is not apparent. A more symmetrical way of presenting this pair of problems is as follows:

$$\begin{array}{lll} \text{Primal:} & & \text{Dual:} \\ & & \\ \text{minimise:} & \mathbf{Tr}(\mathbf{CX}) & \text{minimise:} \quad \mathbf{Tr}(-\mathbf{GZ}) \\ \text{subject to:} & \mathbf{X} \in (\mathcal{L} + \mathbf{G}) \cap \mathbb{S}^n_+ & \text{subject to:} \quad \mathbf{Z} \in (\mathcal{L}^\perp + \mathbf{C}) \cap \mathbb{S}^n_+ \end{array} \tag{B.55}$$

where $\mathbf{C}$ satisfies $\mathbf{Tr}(\mathbf{F}_i\mathbf{C}) = c_i$ and $\mathcal{L}$ is a linear subspace of $\mathbb{S}^n$ and $\mathcal{L}^\perp$ is the orthogonal complement to $\mathcal{L}$ in $\mathbb{S}^n$. A detailed description of the geometry of a semidefinite programme is given in [195]. In the symmetric presentation if the primal and dual variables in (B.55) are strictly feasible then the duality gap is nonzero and is given by

$$p^* - d^* = -\mathbf{Tr}(\mathbf{C}(-\mathbf{G})). \tag{B.56}$$

In this form the modified complementary slackness condition in the centrality equations (B.36) can be summarised as

$$\mathbf{X} + t^{-1}\nabla\phi(\mathbf{Z}) = 0, \tag{B.57}$$

which can be rewritten as

$$\mathbf{XZ} = t^{-1}\mathbf{I}. \tag{B.58}$$

using (B.49). The Hessian of the barrier function (B.46) taken at a point $\mathbf{Y} \in \mathbb{S}^n_+$ is

positive definite and its inverse can be used to define a norm on $\mathbb{S}^n_+$

$$\|\mathbf{H}\|^*_{\mathbf{Y}} = \sqrt{\langle \nabla^2\phi(\mathbf{Y})^{-1}\mathbf{H}, \mathbf{H}\rangle}. \tag{B.59}$$

This norm provides a useful way to measure the distance from a strictly feasible point $\mathbf{W} = (\mathbf{X}, \mathbf{Z})$ to a central point $(\mathbf{W}^*(t) = \mathbf{W}(\mathbf{X}^*(t), \mathbf{Z}^*(t)))$

$$\begin{aligned} \mathrm{dist}(\mathbf{W}, \mathbf{W}^*(t)) &= \|t\mathbf{Z} + \nabla\phi(\mathbf{X})\|^*_{\mathbf{X}} \\ &= \sqrt{\langle (\nabla^2\phi(\mathbf{X}))^{-1}(t\mathbf{Z} + \nabla\phi(\mathbf{X})), t\mathbf{Z} + \nabla\phi(\mathbf{X})\rangle}. \end{aligned} \tag{B.60}$$

The norm is also symmetric, i.e.

$$\|t\mathbf{Z} + \nabla\phi(\mathbf{X})\|^*_{\mathbf{X}} = \|t\mathbf{X} + \nabla\phi(\mathbf{Z})\|^*_{\mathbf{Z}}. \tag{B.61}$$

If $\mathbf{Z} = (\mathbf{X}, \mathbf{Z})$ is a primal-dual strictly feasible pair such that

$$\mathrm{dist}(\mathbf{Z}, \mathbf{Z}^*(t)) \leq 1, \tag{B.62}$$

then $\mathbf{Z}$ is effectively as good a solution as $\mathbf{Z}^*(t)$. What is meant by this is

$$\text{Duality Gap}(\mathbf{X}, \mathbf{Z}) = \langle \mathbf{X}, \mathbf{Z}\rangle \leq 2 \times (\text{Duality Gap}(\mathbf{Z}^*(t))) = \frac{2\theta(\phi)}{t}. \tag{B.63}$$

So there is little difference between moving along the primal-dual central path exactly or to trace its path remaining within its neighbourhood defined by

$$\mathcal{N}_\kappa = \left\{(t, \mathbf{X}, \mathbf{Z}) \mid \mathbf{X} \in (\mathcal{L} + \mathbf{G}), \mathbf{Z} \in (\mathcal{L}^\perp + \mathbf{C}), t > 0, (\mathbf{X}^*(t), \mathbf{Z}^*(t)) \leq \kappa\right\},$$

with $\kappa \leq 1$. Methods that remain within neighbourhoods $\mathcal{N}_{\mathcal{O}(1)}$ are known as short-step methods while algorithms that work in wider neighbourhoods are known as long-step methods. The advantage of long step methods is that they avoid slowing down at highly curved segments of the central path.

To find the Newton step the modified complementary slackness system for an SDP (B.58) is rewritten as

$$\mathbf{Z} - t^{-1}\mathbf{X}^{-1} = 0. \tag{B.64}$$

However because the matrix $\mathbf{XZ}$ is not a symmetric matrix this causes problems in the calculation of the Newton step. These problems are circumvented by rewriting equation B.58 in a symmetric form which can be done in a number of ways. One such way is to note that it is equivalent to each of the following equations

$$\mathbf{XZ} = t^{-1}\mathbf{I}, \qquad \mathbf{ZX} = t^{-1}\mathbf{I}. \tag{B.65}$$

And adding these together produces a symmetric equation

$$\mathbf{XZ} + \mathbf{ZX} = 2t^{-1}\mathbf{I}. \tag{B.66}$$

Using this equation instead of equation B.58 leads to the Alizadeh-Haeberly-Overton search direction [253]. It is also possible to use a scaled version of equation B.66. The modified complementary slackness equation (B.58) is also equivalent to the following two equations

$$\mathbf{QXZQ}^{-1} = t^{-1}\mathbf{I}, \qquad \mathbf{Q}^{-1}\mathbf{ZXQ} = t^{-1}\mathbf{I}, \tag{B.67}$$

where $\mathbf{Q}$ is a positive definite scaling matrix. Summing these equations gives

$$\mathbf{QXZQ}^{-1} + \mathbf{Q}^{-1}\mathbf{ZXQ} = 2t^{-1}\mathbf{I}. \tag{B.68}$$

This leads to what is known as the Zhang search direction [254]. The scaling matrices can vary from iteration to iteration and the choice of scaling matrix can have a dramatic effect on the analysis of the method. For example when the scaling matrices are chosen such that

$$\widetilde{\mathbf{Z}}(t) = \mathbf{Q}(t)^{-1}\mathbf{Z}(t)\mathbf{Q}(t)^{-1} \quad \text{and} \quad \widehat{\mathbf{X}}(t) = \mathbf{Q}(t)\mathbf{X}(t)\mathbf{Q}(t), \tag{B.69}$$

commute, these are known as commutative scaling approaches. A particularly attractive commutative scaling approach is that the Nesterov-Todd (NT) approach [255] which requires that $\mathbf{Q}$ is chosen such that $\widetilde{\mathbf{Z}} = \widehat{\mathbf{X}}$. Nesterov and Todd proved that $\mathbf{Q}$ is guaranteed to exist and be unique and this leads to an efficient way of finding $\mathbf{Q}$. The Nesterov-Todd commutative scaling approach is widely used in state of the art interior-point solvers such as SeDuMi and SDPT3 [256]. Another popular search direction is the Helmberg-Kojima-Monteiro (HKM) direction [257–259] used in SDPA [212], CSDP [213] and SDPT3 [256].

References

[1] P. Parrilo and B. Sturmfels, *Algorithmic and quantitative real algebraic geometry* (AMS, 2003), vol. 60, chap. Minimising polynomial functions, 83–99.

[2] L. Paramonov and S. N. Yaliraki, *The directional contact distance of two ellipsoids: Coarse-grained potentials for anisotropic interactions*, Journal of Chemical Physics **123**, 194111 (2005).

[3] D. J. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses* (Cambridge Molecular Science, 2004).

[4] J. D. Bryngelson and P. G. Wolynes, *Spin-Glasses and the Statistical-Mechanics of Protein Folding*, Proceedings of the National Academy of Sciences of the United States of America **84**, 7524–7528 (1987).

[5] J. D. Bryngelson and P. G. Wolynes, *Intermediates and Barrier Crossing in a Random Energy-Model (With Applications to Protein Folding)*, Journal of Physical Chemistry **93**, 6902–6915 (1989).

[6] J. D. Bryngelson and P. G. Wolynes, *A Simple Statistical Field-Theory of Heteropolymer Collapse with Application to Protein Folding*, Biopolymers **30**, 177–188 (1990).

[7] P. E. Leopold, M. Montal and J. N. Onuchic, *Protein Folding Funnels - a Kinetic Approach to the Sequence Structure Relationship*, Proceedings of the National Academy of Sciences of the United States of America **89**, 8721–8725 (1992).

[8] J. N. Onuchic, P. G. Wolynes, Z. Luthey-Schulten and N. D. Socci, *Toward an Outline of the Topography of a Realistic Protein-Folding Funnel*, Proceedings of the National Academy of Sciences of the United States of America **92**, 3626–3630 (1995).

[9] J. D. Bryngelson, J. N. Onuchic, N. D. Socci and P. G. Wolynes, *Funnels, Pathways, and the Energy Landscape of Protein-Folding - a Synthesis*, Proteins-Structure Function and Genetics **21**, 167–195 (1995).

[10] J. N. Onuchic, Z. LutheySchulten and P. G. Wolynes, *Theory of protein folding: The energy landscape perspective*, Annual Review of Physical Chemistry **48**, 545–600 (1997).

[11] J. Simons, *Introduction to Theoretical Chemistry* (Cambridge University Press, 2003).

[12] D. J. Wales, J. P. K. Doye, M. A. Miller, P. N. Mortenson and T. R. Walsh, *Energy landscapes: From clusters to biomolecules*, Advances in Chemical Physics **115**, 1–111 (2000).

[13] R. Hoffmann, *An Extended Hückel Theory. I. Hydrocarbons*, Journal of Chemical Physics **39**, 1397 (1963).

[14] C. Møller and M. S. Plesset, *Note on an Approximation Treatment for Many-Electron Systems*, Physical Review **46**, 618 (1934).

[15] R. G. Parr and W. Yang, *Density-functional theory of atoms and molecules*, vol. 16 (Oxford University Press; Clarendon Press, New York; Oxford England, 1989).

[16] B. L. Hammond, W. A. Lester and P. J. Reynolds, *Monte Carlo methods in Ab Initio Quantum Chemistry* (World Scientific, Singapore, 1994).

[17] U. Burkert and N. L. Allinger, *Molecular mechanics*, vol. 177 (American Chemical Society, Washington, D.C., 1982).

[18] N. L. Allinger, *Conformational-Analysis .130. Mm2 - Hydrocarbon Force-Field Utilizing V1 and V2 Torsional Terms*, Journal of the American Chemical Society **99**, 8127–8134 (1977).

[19] W. D. Conrell, P. Cieplak, C. I. Bayly, I. R. Gould, K. M. Merz, D. M. Ferguson, D. C. Spellmeyer, T. Fox, J. W. Caldwell and P. A. Kollman, *A 2nd Generation Force-Field for the Simulation of Proteins, Nucleic-Acids, and Organic-Molecules*, Journal of the American Chemical Society **117**, 5179–5197 (1995).

[20] B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan and M. Karplus, *Charmm - a Program for Macromolecular Energy, Minimization, and Dynamics Calculations*, Journal of Computational Chemistry **4**, 187–217 (1983).

[21] A. D. MacKerell, D. Bashford, M. Bellott, R. L. Dunbrack, J. D. Evanseck, M. J. Field, S. Fischer, J. Gao, H. Guo, S. Ha, D. Joseph-McCarthy, L. Kuchnir, K. Kuczera, F. T. K. Lau, C. Mattos, S. Michnick, T. Ngo, D. T. Nguyen,

B. Prodhom, W. E. Reiher, B. Roux, M. Schlenkrich, J. C. Smith, R. Stote, J. Straub, M. Watanabe, J. Wiorkiewicz-Kuczera, D. Yin and M. Karplus, *All-atom empirical potential for molecular modeling and dynamics studies of proteins*, Journal of Physical Chemistry B **102**, 3586–3616 (1998).

[22] F. A. Momany, R. F. McGuire, A. W. Burgess and H. A. Scheraga, *Energy Parameters in Polypeptides .7. Geometric Parameters, Partial Atomic Charges, Nonbonded Interactions, Hydrogen-Bond Interactions, and Intrinsic Torsional Potentials for Naturally Occurring Amino-Acids*, Journal of Physical Chemistry **79**, 2361–2381 (1975).

[23] M. J. Sippl, G. Nemethy and H. A. Scheraga, *Intermolecular Potentials from Crystal Data. 6. Determination of Empirical Potentials for O-H=o=c Hydrogen-Bonds from Packing Configurations*, Journal of Physical Chemistry **88**, 6231–6233 (1984).

[24] W. F. van Gunsteren and H. J. C. Berendsen, *The GROMOS package*, http://www.igc.ethz.ch/gromos/ (1987).

[25] W. L. Jorgensen and J. Tirado-Rives, *The OPLS Potential Functions for Proteins - Energy Minimizations for Crystals of Cyclic-Peptides and Crambin*, Journal of the American Chemical Society **110**, 1657–1666 (1988).

[26] J. Lee, A. Liwo and H. A. Scheraga, *Energy-based de novo protein folding by conformational space annealing and an off-lattice united-residue force field: Application to the 10-55 fragment of staphylococcal protein A and to apo calbindin D9K*, Proceedings of the National Academy of Sciences of the United States of America **96**, 2025–2030 (1999).

[27] J. Pillardy, C. Czaplewski, A. Liwo, W. J. Wedemeyer, J. Lee, D. R. Ripoll, P. Arlukowicz, S. Oldziej, Y. A. Arnautova and H. A. Scheraga, *Development of physics-based energy functions that predict medium-resolution structures for proteins of the alpha,beta and alpha/beta structural classes*, Journal of Physical Chemistry B **105**, 7299–7311 (2001).

[28] J. N. Murrell and K. J. Laidler, *Symmetries of activated complexes*, Transactions of the Faraday Society **64**, 371 (1968).

[29] H. Eyring, *The Activated Complex and the Absolute Rate of Chemical Reactions*, Chemical Reviews **17**, 65–77 (1935).

[30] M. G. Evans and M. Polanyi, *Some applications of the transition state method to the calculation of reaction velocities, especially in solution*, Transactions of the Faraday Society **31**, 875–894 (1935).

[31] W. Forst, *Theory of unimolecular reactions*, vol. 30 (Academic Press, New York, 1973).

[32] K. J. Laidler, *Chemical kinetics* (HarperCollins, New York, 1987).

[33] B. C. Garrett and D. G. Truhlar, *Criterion of Minimum State Density in the Transition-State Theory of Bimolecular Reactions*, Journal of Chemical Physics **70**, 1593–1598 (1979).

[34] D. G. Truhlar and B. C. Garrett, *Variational Transition-State Theory*, Accounts of Chemical Research **13**, 440–448 (1980).

[35] B. C. Garrett and D. G. Truhlar, *Generalized Transition-State Theory - Canonical Variational Calculations using the Bond Energy Bond Order Method for Bimolecular Reactions of Combustion Products*, Journal of the American Chemical Society **101**, 5207–5217 (1979).

[36] B. C. Garrett and D. G. Truhlar, *Variational Transition-State Theory - Primary Kinetic Isotope Effects for Atom Transfer-Reactions*, Journal of the American Chemical Society **102**, 2559–2570 (1980).

[37] D. J. Wales and R. S. Berry, *Limitations of the Murrell-Laidler Theorem*, Journal of the Chemical Society-Faraday Transactions **88**, 543–544 (1992).

[38] C. J. Tsai and K. D. Jordan, *Use of the Histogram and Jump-Walking Methods for Overcoming Slow Barrier Crossing Behavior in Monte-Carlo Simulations - Applications to the Phase-Transitions in the (Ar)13 and (H2o)8 Clusters*, Journal of Chemical Physics **99**, 6957–6970 (1993).

[39] F. H. Stillinger and T. A. Weber, *Hidden Structure in Liquids*, Physical Review A **25**, 978–989 (1982).

[40] F. H. Stillinger and T. A. Weber, *Packing Structures and Transitions in Liquids and Solids*, Science **225**, 983–989 (1984).

[41] F. H. Stillinger, *Exponential multiplicity of inherent structures*, Physical Review E **59**, 48–51 (1999).

[42] J. P. K. Doye and D. J. Wales, *Saddle points and dynamics of Lennard-Jones clusters, solids, and supercooled liquids*, Journal of Chemical Physics **116**, 3777–3788 (2002).

[43] R. S. Berry and R. Breitengraserkunz, *Topography and Dynamics of Multidimensional Interatomic Potential Surfaces*, Physical Review Letters **74**, 3951–3954 (1995).

[44] R. E. Kunz and R. S. Berry, *Statistical Interpretation of Topographies and Dynamics of Multidimensional Potentials*, Journal of Chemical Physics **103**, 1904–1912 (1995).

[45] P. G. Mezey, *Catchment Region Partitioning of Energy Hypersurfaces. 1.*, Theoretica Chimica Acta **58**, 309–330 (1981).

[46] K. D. Ball, R. S. Berry, R. E. Kunz, F. Y. Li, A. Proykova and D. J. Wales, *From topographies to dynamics on multidimensional potential energy surfaces of atomic clusters*, Science **271**, 963–966 (1996).

[47] O. M. Becker and M. Karplus, *The topology of multidimensional potential energy surfaces: Theory and application to peptide structure and kinetics*, Journal of Chemical Physics **106**, 1495–1517 (1997).

[48] R. Czerminski and R. Elber, *Reaction-Path Study of Conformational Transitions in Flexible Systems - Applications to Peptides*, Journal of Chemical Physics **92**, 5580–5601 (1990).

[49] J. P. K. Doye, M. A. Miller and D. J. Wales, *Evolution of the potential energy surface with size for Lennard-Jones clusters*, Journal of Chemical Physics **111**, 8417–8428 (1999).

[50] S. F. Chekmarev and S. V. Krivov, *Confinement of the molecular dynamics trajectory to a specified catchment area on the potential surface: Application to clusters*, European Physical Journal D **9**, 201–204 (1999).

[51] S. F. Chekmarev and S. V. Krivov, *Confinement of the molecular dynamics trajectory to a specified catchment area on the potential surface*, Chemical Physics Letters **287**, 719–724 (1998).

[52] S. F. Chekmarev and S. V. Krivov, *Total and fractional densities of states from caloric relations*, Physical Review E **57**, 2445–2448 (1998).

[53] H. B. Schlegel, Modern Electronic Structure Theory (World Scientific, Singapore, 1995), 459–500.

[54] H. B. Schlegel (John Wiley and Sons, New York, 1998), vol. 2 of *Encyclopedia of Computational Chemistry*, 1136.

[55] C. G. Broyden, *Convergence of Single-Rank Quasi-Newton Methods*, Mathematics of Computation **24**, 365 (1970).

[56] R. Fletcher, *A New Approach to Variable Metric Algorithms*, Computer journal **13**, 317 (1970).

[57] D. Goldfarb, *A Family of Variable-Metric Methods Derived by Variational Means*, Mathematics of Computation **24**, 23 (1970).

[58] D. F. Shanno, *Conditioning of Quasi-Newton Methods for Function Minimization*, Mathematics of Computation **24**, 647 (1970).

[59] J. Nocedal, *Updating Quasi-Newton Matrices with Limited Storage*, Mathematics of Computation **35**, 773–782 (1980).

[60] D. C. Liu and J. Nocedal, *On the Limited Memory BFGS Method for Large-Scale Optimization*, Mathematical Programming **45**, 503–528 (1989).

[61] W. H. Press, *Numerical recipes in FORTRAN : the art of scientific computing* (Cambridge University Press, Cambridge England ; New York, NY, USA, 1992).

[62] R. Fletcher, *Practical methods of optimization* (Wiley, Chichester ; New York, 1987).

[63] M. Page and J. W. McIver, *On Evaluating the Reaction-Path Hamiltonian*, Journal of Chemical Physics **88**, 922–935 (1988).

[64] L. J. Munro and D. J. Wales, *Defect migration in crystalline silicon*, Physical Review B **59**, 3969–3980 (1999).

[65] K. Schiele and R. Hemmecke, *Migration effects in driven multiple pendula*, Zeitschrift Fur Angewandte Mathematik Und Mechanik **81**, 291–303 (2001).

[66] W. Quapp, *Comment "On the quadratic reaction path evaluated in a reduced potential energy surface model and the problem to locate transition states"*, Journal of Computational Chemistry **22**, 537–540 (2001).

[67] J. M. Bofill and J. M. Anglada, *Finding transition states using reduced potential-energy surfaces*, Theoretical Chemistry Accounts **105**, 463–472 (2001).

[68] R. Crehuet, J. M. Bofill and J. M. Anglada, *A new look at the reduced-gradient-following path*, Theoretical Chemistry Accounts **107**, 130–139 (2002).

[69] M. Dallos, H. Lischka, E. V. D. Monte, M. Hirsch and W. Quapp, *Determination of energy minima and saddle points using multireference configuration interaction methods in combination with reduced gradient following: The S-0 surface of H2CO and the T-1 and T-2 surfaces of acetylene*, Journal of Computational Chemistry **23**, 576–583 (2002).

[70] W. Quapp and D. Heidrich, *Exploring the potential energy surface of the ethyl cation by new procedures*, Journal of Molecular Structure-Theochem **585**, 105–117 (2002).

[71] K. Muller and L. D. Brown, *Location of Saddle Points and Minimum Energy Paths by a Constrained Simplex Optimization Procedure*, Theoretica Chimica Acta **53**, 75–93 (1979).

[72] M. J. Rothman and L. L. Lohr, *Analysis of an Energy Minimization Method for Locating Transition-States on Potential-Energy Hypersurfaces*, Chemical Physics Letters **70**, 405–409 (1980).

[73] K. Müller, *Reaction Paths on Multidimensional Energy Hypersurfaces*, Angewandte Chemie-International Edition in English **19**, 1–13 (1980).

[74] P. Scharfenberg, *Theoretical-Analysis of Constrained Minimum Energy Paths*, Chemical Physics Letters **79**, 115–117 (1981).

[75] P. Scharfenberg, *An Improved Method for the Evaluation of Transition-States*, Journal of Computational Chemistry **3**, 277–282 (1982).

[76] G. M. Maggiora and I. H. Williams, *Intermolecular Interaction Energies from Minimal-Basis SCF Calculations - Interactions Pertinent to Formaldehyde Hydration*, Theochem-Journal of Molecular Structure **5**, 23–35 (1982).

[77] R. Steckler and D. G. Truhlar, *Reaction-Path Power-Series Analysis of NH3 Inversion*, Journal of Chemical Physics **93**, 6570–6577 (1990).

[78] J. Cioslowski, A. P. Scott and L. Radom, *Catastrophes, bifurcations and hysteretic loops in torsional potentials of internal rotations in molecules*, Molecular Physics **91**, 413–420 (1997).

[79] M. Cernohorsky, S. Kettou and J. Koca, *VADER: New software for exploring interconversions on potential energy surfaces*, Journal of chemical information and computer sciences **39**, 705–712 (1999).

[80] S. Fischer and M. Karplus, *Conjugate Peak Refinement - an Algorithm for Finding Reaction Paths and Accurate Transition-States in Systems with Many Degrees of Freedom*, Chemical Physics Letters **194**, 252–261 (1992).

[81] L. R. Pratt, *A Statistical-Method for Identifying Transition-States in High Dimensional Problems*, Journal of Chemical Physics **85**, 5045–5048 (1986).

[82] R. Elber and M. Karplus, *A Method for Determining Reaction Paths in Large Molecules - Application to Myoglobin*, Chemical Physics Letters **139**, 375–380 (1987).

[83] L. L. Stacho and M. I. Ban, *A Global Strategy for Determining Reaction Paths .1. General-Theory of a Procedure Finding Fukui Intrinsic Reaction Coordinate*, Theoretica Chimica Acta **83**, 433–440 (1992).

[84] I. V. Ionova and E. A. Carter, *Ridge Method for Finding Saddle-Points on Potential-Energy Surfaces*, Journal of Chemical Physics **98**, 6377–6386 (1993).

[85] L. L. Stacho and M. I. Ban, *An Algorithm for Determining Dynamically Defined Reaction Paths (DDRP)*, Theoretica Chimica Acta **84**, 535–543 (1993).

[86] G. Domotor, M. I. Ban and L. L. Stacho, *Experiences and Practical Hints on using the DDRP Method, Illustrated by the Example of the H(2)+H Reaction*, Journal of Computational Chemistry **14**, 1491–1497 (1993).

[87] A. Matro, D. L. Freeman and J. D. Doll, *Locating Transition-States using Double-Ended Classical Trajectories*, Journal of Chemical Physics **101**, 10458–10463 (1994).

[88] O. S. Smart, *A New Method to Calculate Reaction Paths for Conformational Transitions of Large Molecules*, Chemical Physics Letters **222**, 503–512 (1994).

[89] P. Bolhuis, D. Chandler, C. Dellago and P. Geissler, *Transition path sampling: Throwing ropes over rough mountain passes, in the dark*, Annual Review of Physical Chemistry **53**, 291–318 (2002).

[90] G. Mills and H. Jonsson, *Quantum and Thermal Effects in H-2 Dissociative Adsorption - Evaluation of Free-Energy Barriers in Multidimensional Quantum-Systems*, Physical Review Letters **72**, 1124–1127 (1994).

[91] G. Henkelman, B. P. Uberuaga and H. Jonsson, *A climbing image nudged elastic band method for finding saddle points and minimum energy paths*, Journal of Chemical Physics **113**, 9901–9904 (2000).

[92] G. Henkelman and H. Jonsson, *Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points*, Journal of Chemical Physics **113**, 9978–9985 (2000).

[93] P. Maragakis, S. A. Andreev, Y. Brumer, D. R. Reichman and E. Kaxiras, *Adaptive nudged elastic band approach for transition state calculation*, Journal of Chemical Physics **117**, 4651–4658 (2002).

[94] I. R. McDonald and K. Singer, Discussions of the Faraday Society **43**, 40 (1967).

[95] A. M. Ferrenberg and R. H. Swendsen, *New Monte-Carlo Technique for Studying Phase-Transitions*, Physical Review Letters **61**, 2635–2638 (1988).

[96] A. M. Ferrenberg and R. H. Swendsen, *Optimized Monte-Carlo Data-Analysis*, Physical Review Letters **63**, 1195–1198 (1989).

[97] C. Bichara, J. P. Gaspard and J. C. Mathieu, *Monte-Carlo Calculation of the Configurational Entropy and Combinatory Factor on Binary-Alloys*, Physics Letters A **119**, 462–468 (1987).

[98] G. M. Torrie and J. P. Valleau, *Monte-Carlo Free-Energy Estimates using Non-Boltzmann Sampling - Application to Subcritical Lennard-Jones Fluid*, Chemical Physics Letters **28**, 578–581 (1974).

[99] J. E. Shea and C. L. Brooks, *From folding theories to folding proteins: A review and assessment of simulation studies of protein folding and unfolding*, Annual Review of Physical Chemistry **52**, 499–535 (2001).

[100] D. D. Frantz, D. L. Freeman and J. D. Doll, *Reducing Quasi-Ergodic Behavior in Monte-Carlo Simulations by J-Walking - Applications to Atomic Clusters*, Journal of Chemical Physics **93**, 2769–2784 (1990).

[101] D. D. Frantz, D. L. Freeman and J. D. Doll, *Extending J-Walking to Quantum-Systems - Applications to Atomic Clusters*, Journal of Chemical Physics **97**, 5713–5731 (1992).

[102] R. H. Swendsen and J. S. Wang, *Replica Monte-Carlo Simulation of Spin-Glasses*, Physical Review Letters **57**, 2607–2609 (1986).

[103] K. Hukushima and K. Nemoto, *Exchange Monte Carlo method and application to spin glass simulations*, Journal of the Physical Society of Japan **65**, 1604–1608 (1996).

[104] M. C. Tesi, E. J. J. vanRensburg, E. Orlandini and S. G. Whittington, *Monte Carlo study of the interacting self-avoiding walk model in three dimensions*, Journal of Statistical Physics **82**, 155–181 (1996).

[105] U. H. E. Hansmann, *Parallel tempering algorithm for conformational studies of biological molecules*, Chemical Physics Letters **281**, 140–150 (1997).

[106] Q. L. Yan and J. J. de Pablo, *Hyper-parallel tempering Monte Carlo: Application to the Lennard-Jones fluid and the restricted primitive model*, Journal of Chemical Physics **111**, 9509–9516 (1999).

[107] Y. Sugita and Y. Okamoto, *Replica-exchange molecular dynamics method for protein folding*, Chemical Physics Letters **314**, 141–151 (1999).

[108] K. Hukushima, *Domain-wall free energy of spin-glass models: Numerical method and boundary conditions*, Physical Review E **60**, 3606–3613 (1999).

[109] D. Gront, A. Kolinski and J. Skolnick, *Comparison of three Monte Carlo conformational search strategies for a proteinlike homopolymer model: Folding thermodynamics and identification of low-energy structures*, Journal of Chemical Physics **113**, 5065–5071 (2000).

[110] Y. Zhang and J. Skolnick, *Parallel-hat tempering: A Monte Carlo search scheme for the identification of low-energy structures*, Journal of Chemical Physics **115**, 5027–5032 (2001).

[111] R. H. Zhou, B. J. Berne and R. Germain, *The free energy landscape for beta hairpin folding in explicit water*, Proceedings of the National Academy of Sciences of the United States of America **98**, 14931–14936 (2001).

[112] D. A. Kofke, *On the acceptance probability of replica-exchange Monte Carlo trials*, Journal of Chemical Physics **117**, 6911–6914 (2002).

[113] D. A. Evans and D. J. Wales, *Free energy landscapes of model peptides and proteins*, Journal of Chemical Physics **118**, 3891–3897 (2003).

[114] H. R. Lewis and C. H. Papadimitriou, *The Efficiency of Algorithms*, Scientific American **238**, 96–109 (1978).

[115] A. Turing, *On Computable Numbers, With an Application to the Entscheidungsproblem*, Proceedings of the London Mathematical Society **42** (1936).

[116] P. A. Parrilo, *Exploiting structure in sum of squares programs*, in *Proceedings for the 42nd IEEE Conference on Decision and Control* (Maui, Hawaii, 2003).

[117] F. Barahona, *On the Computational-Complexity of Ising Spin-Glass Models*, Journal of Physics A-Mathematical and General **15**, 3241–3253 (1982).

[118] I. Bieche, R. Maynard, R. Rammal and J. P. Uhry, *On the Ground-States of the Frustration Model of a Spin-Glass by a Matching Method of Graph-Theory*, Journal of Physics A-Mathematical and General **13**, 2553–2576 (1980).

[119] J. T. Ngo and J. Marks, *Computational-Complexity of a Problem in Molecular-Structure Prediction*, Protein engineering **5**, 313–321 (1992).

[120] R. Unger and J. Moult, *Finding the Lowest Free-Energy Conformation of a Protein is an NP-Hard Problem - Proof and Implications*, Bulletin of mathematical biology **55**, 1183–1198 (1993).

[121] W. E. Hart and S. Istrail, *Robust proofs of NP-hardness for protein folding: General lattices and energy potentials*, Journal of Computational Biology **4**, 1–22 (1997).

[122] I. P. Androulakis, C. D. Maranas and C. A. Floudas, *alpha BB: A global optimization method for general constrained nonconvex problems*, Journal of Global Optimization **7**, 337–363 (1995).

[123] D. J. Wales and J. P. K. Doye, *Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms*, Journal of Physical Chemistry a **101**, 5111–5116 (1997).

[124] L. T. Wiile, *Minimum-Energy Configurations of Atomic Clusters - New Results obtained by Simulated Annealing*, Chemical Physics Letters **133**, 405–410 (1987).

[125] J. A. Northby, *Structure and Binding of Lennard-Jones Clusters - 13≤N≤147*, Journal of Chemical Physics **87**, 6166–6177 (1987).

[126] T. Coleman, D. Shalloway and Z. J. Wu, *A Parallel Buildup Algorithm for Global Energy Minimizations of Molecular Clusters using Effective Energy Simulated Annealing*, Journal of Global Optimization **4**, 171–185 (1994).

[127] G. L. Xue, *Improvement on the Northby Algorithm for Molecular-Conformation - Better Solutions*, Journal of Global Optimization **4**, 425–440 (1994).

[128] J. Pillardy and L. Piela, *Molecular-Dynamics on Deformed Potential-Energy Hypersurfaces*, Journal of Physical Chemistry **99**, 11805–11812 (1995).

[129] J. P. K. Doye, D. J. Wales and R. S. Berry, *The Effect of the Range of the Potential on the Structures of Clusters*, Journal of Chemical Physics **103**, 4234–4249 (1995).

[130] J. P. K. Doye and D. J. Wales, *Magic numbers and growth sequences of small face-centred-cubic and decahedral clusters*, Chemical Physics Letters **247**, 339–347 (1995).

[131] D. M. Deaven, N. Tit, J. R. Morris and K. M. Ho, *Structural optimization of Lennard-Jones clusters by a genetic algorithm*, Chemical Physics Letters **256**, 195–200 (1996).

[132] C. Barron, S. Gomez and D. Romero, *Lower energy icosahedral atomic clusters with incomplete core*, Applied Mathematics Letters **10**, 25–28 (1997).

[133] R. H. Leary, *Global optima of Lennard-Jones clusters*, Journal of Global Optimization **11**, 35–53 (1997).

[134] C. Barron, S. Gomez, D. Romero and A. Saavedra, *A genetic algorithm for Lennard-Jones atomic clusters*, Applied Mathematics Letters **12**, 85–90 (1999).

[135] A. L. Mackay, *A dense non-crystallographic packing of equal spheres*, Acta Crystallographica **15**, 916 (1962).

[136] J. P. K. Doye, *Effect of compression on the global optimization of atomic clusters*, Physical Review E **62**, 8753–8761 (2000).

[137] J. P. K. Doye, *Effect of compression on the global optimization of atomic clusters*, Physical Review E **62**, 8753–8761 (2000).

[138] S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi, *Optimization by Simulated Annealing*, Science **220**, 671–680 (1983).

[139] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller, *Equation of State Calculations by Fast Computing Machines*, Journal of Chemical Physics **21**, 1087 (1953).

[140] S. Geman and D. Geman, *Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images*, IEEE Transactions on Pattern Analysis and Machine Intelligence **6**, 721–741 (1984).

[141] P. Salamon, J. D. Nulton, J. R. Harland, J. Pedersen, G. Ruppeiner and L. Liao, *Simulated Annealing with Constant Thermodynamic Speed*, Computer Physics Communications **49**, 423–428 (1988).

[142] J. D. Nulton and P. Salamon, *Statistical-Mechanics of Combinatorial Optimization*, Physical Review a **37**, 1351–1356 (1988).

[143] K. H. Hoffmann and P. Salamon, *The Optimal Simulated Annealing Schedule for a Simple-Model*, Journal of Physics A-Mathematical and General **23**, 3511–3523 (1990).

[144] J. Y. Yi, J. Bernholc and P. Salamon, *Simulated Annealing Strategies for Molecular-Dynamics*, Computer Physics Communications **66**, 177–180 (1991).

[145] G. A. Huber and J. A. McCammon, *Weighted-ensemble simulated annealing: Faster optimization on hierarchical energy surfaces*, Physical Review E **55**, 4822–4825 (1997).

[146] D. M. Ferguson and D. G. Garrett, *Simulated annealing - Optimal histogram methods*, Monte Carlo Methods in Chemical Physics **105**, 311–336 (1999).

[147] R. P. White and H. R. Mayne, *Optimal annealing schedules for two-, three-, and four-level systems using a genetic algorithm approach*, Journal of Chemical Physics **112**, 7964–7978 (2000).

[148] L. Piela, J. Kostrowicki and H. A. Scheraga, *The Multiple-Minima Problem in the Conformational-Analysis of Molecules - Deformation of the Potential-Energy Hypersurface by the Diffusion Equation Method*, Journal of Physical Chemistry **93**, 3339–3346 (1989).

[149] C. Barron, S. Gomez and D. Romero, *Archimedean polyhedron structure yields a lower energy atomic cluster*, Applied Mathematics Letters **9**, 75–78 (1996).

[150] C. Barron, S. Gomez and D. Romero, *Lower energy icosahedral atomic clusters with incomplete core*, Applied Mathematics Letters **10**, 25–28 (1997).

[151] G. L. Xue, *Molecular-Conformation on the Cm-5 by Parallel 2-Level Simulated Annealing*, Journal of Global Optimization **4**, 187–208 (1994).

[152] D. J. Wales and H. A. Scheraga, *Global optimization of clusters, crystals, and biomolecules*, Science (1999).

[153] J. P. K. Doye and D. J. Wales, *Thermodynamics of global optimization*, Physical Review Letters **80**, 1357–1360 (1998).

[154] J. P. K. Doye, D. J. Wales and M. A. Miller, *Thermodynamics and the global optimization of Lennard-Jones clusters*, Journal of Chemical Physics **109**, 8143–8153 (1998).

[155] D. J. Wales, J. P. K. Doye, A. Dullweber, M. P. Hodges, F. Y. Naumkin, F. Calvo, J. Hernandez-Rojas and T. F. Middleton, *The Cambridge Cluster Database*.

[156] J. P. K. Doye and D. J. Wales, *Structural consequences of the range of the interatomic potential - A menagerie of clusters*, Journal of the Chemical Society-Faraday Transactions **93**, 4233–4243 (1997).

[157] J. P. K. Doye, D. J. Wales, W. Branz and F. Calvo, *Modeling the structure of clusters of C-60 molecules*, Physical Review B **64**, 235409 (2001).

[158] L. Piela and H. A. Scheraga, *On the Multiple-Minima Problem in the Conformational-Analysis of Polypeptides .1. Backbone Degrees of Freedom for a Perturbed Alpha-Helix*, Biopolymers **26**, S33–S58 (1987).

[159] D. R. Ripoll and H. A. Scheraga, *On the Multiple-Minima Problem in the Conformational-Analysis of Polypeptides .2. an Electrostatically Driven Monte-Carlo Method - Tests on Poly(l-Alanine)*, Biopolymers **27**, 1283–1303 (1988).

[160] D. R. Ripoll and H. A. Scheraga, *On the Multiple-Minima Problem in the Conformational-Analysis of Polypeptides .4. Application of the Electrostatically Driven Monte-Carlo Method to the 20-Residue Membrane-Bound Portion of Melittin*, Biopolymers **30**, 165–176 (1990).

[161] D. R. Ripoll, M. J. Vasquez and H. A. Scheraga, *The Electrostatically Driven Monte-Carlo Method - Application to Conformational-Analysis of Decaglycine*, Biopolymers **31**, 319–330 (1991).

[162] F. Y. Naumkin and D. J. Wales, *Influence of the atom-atom interaction anisotropy on the structure and stability of ArnCl2 clusters*, Chemical Physics Letters **290**, 164–170 (1998).

[163] F. Y. Naumkin and D. J. Wales, *Structure and properties of Ne-N(+) clusters from a diatomics-in-molecules approach*, Molecular Physics **93**, 633–648 (1998).

[164] P. N. Mortenson and D. J. Wales, *Energy landscapes, global optimization and dynamics of the polyalanine Ac(ala)(8)NHMe*, Journal of Chemical Physics **114**, 6443–6454 (2001).

[165] P. N. Mortenson, D. A. Evans and D. J. Wales, *Energy landscapes of model polyalanines*, Journal of Chemical Physics **117**, 1363–1376 (2002).

[166] Z. Q. Li and H. A. Scheraga, *Monte-Carlo-Minimization Approach to the Multiple-Minima Problem in Protein Folding*, Proceedings of the National Academy of Sciences of the United States of America **84**, 6611–6615 (1987).

[167] Z. Q. Li and H. A. Scheraga, *Structure and Free-Energy of Complex Thermodynamic Systems*, Theochem-Journal of Molecular Structure **48**, 333–352 (1988).

[168] J. H. Holland, *Adaptation in natural and artificial systems : an introductory analysis with applications to biology, control, and artificial intelligence* (University of Michigan Press, Ann Arbor, 1975).

[169] D. E. Goldberg, *Genetic algorithms in search, optimization, and machine learning* (Addison-Wesley Pub. Co., Reading, Mass., 1989).

[170] B. Hartke, *Global Geometry Optimization of Clusters using Genetic Algorithms*, Journal of Physical Chemistry **97**, 9973–9976 (1993).

[171] Y. Zeiri, *Prediction of the Lowest Energy Structure of Clusters using a Genetic Algorithm*, Physical Review E **51**, R2769–R2772 (1995).

[172] Y. Zeiri, *Study of the lowest energy structure of atomic clusters using a genetic algorithm*, Computer Physics Communications **103**, 28–42 (1997).

[173] B. Hartke, *Global cluster geometry optimization by a phenotype algorithm with niches: Location of elusive minima, and low-order scaling with cluster size*, Journal of Computational Chemistry **20**, 1752–1759 (1999).

[174] D. M. Deaven and K. M. Ho, *Molecular-Geometry Optimization with a Genetic Algorithm*, Physical Review Letters **75**, 288–291 (1995).

[175] R. P. White, J. A. Niesse and H. R. Mayne, *A study of genetic algorithm approaches to global geometry optimization of aromatic hydrocarbon microclusters*, Journal of Chemical Physics **108**, 2208–2218 (1998).

[176] R. L. Johnston and C. Roberts, *Soft Computing Approaches in Chemistry* (Physica-Verlag, Heidelberg, 2003).

[177] *Birmingham cluster web site, http://www.tc.bham.ac.uk/bcweb.*

[178] J. Lee, H. A. Scheraga and S. Rackovsky, *New optimization method for conformational energy calculations on polypeptides: Conformational space annealing*, Journal of Computational Chemistry **18**, 1222–1232 (1997).

[179] J. Lee, H. A. Scheraga and S. Rackovsky, *Conformational analysis of the 20-residue membrane-bound portion of melittin by conformational space annealing*, Biopolymers **46**, 103–115 (1998).

[180] A. Liwo, J. Lee, D. R. Ripoll, J. Pillardy and H. A. Scheraga, *Protein structure prediction by global optimization of a potential energy function*, Proceedings of the National Academy of Sciences of the United States of America **96**, 5482–5485 (1999).

[181] *Third community wide experiment on the crticial assessment of the structural prediction of proteins*, http://predictioncenter.org/casp3/Casp3.html.

[182] Y. Sugita and Y. Okamoto, *Replica-exchange multicanonical algorithm and multicanonical replica-exchange method for simulating systems with rough energy landscape*, Chemical Physics Letters **329**, 261–270 (2000).

[183] K. W. Foreman, A. T. Phillips, J. B. Rosen and K. A. Dill, *Comparing search strategies for finding global optima on energy landscapes*, Journal of Computational Chemistry **20**, 1527–1532 (1999).

[184] J. Barhen, V. Protopopescu and D. Reister, *TRUST: A deterministic algorithm for global optimization*, Science **276**, 1094–1097 (1997).

[185] J. L. Klepeis and C. A. Floudas, *Free energy calculations for peptides via deterministic global optimization*, Journal of Chemical Physics **110**, 7491–7512 (1999).

[186] I. P. Androulakis, C. D. Maranas and C. A. Floudas, *Prediction of oligopeptide conformations via deterministic global optimization*, Journal of Global Optimization **11**, 1–34 (1997).

[187] K. M. Westerberg and C. A. Floudas, *Dynamics of peptide folding: Transition states and reaction pathways of solvated and unsolvated tetra-alanine*, Journal of Global Optimization **15**, 261–297 (1999).

[188] J. L. Klepeis, C. A. Floudas, D. Morikis and J. D. Lambris, *Predicting peptide structures using NMR data and deterministic global optimization*, Journal of Computational Chemistry **20**, 1354–1370 (1999).

[189] C. D. Maranas and C. A. Floudas, *A Global Optimization Approach for Lennard-Jones Microclusters*, Journal of Chemical Physics **97**, 7667–7678 (1992).

[190] M. Goemans and D. Williamson, *Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems using Semidefinite Programming*, Journal of the A.C.M. **42**, 1115–1145 (1995).

[191] P. Parrilo, *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimisation*, Ph.D. thesis, California Institute of Technology,, Pasadena, CA., U.S.A. (2000).

[192] J. B. Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM Journal on optimization **11**, 796 – 817 (2001).

[193] N. Z. Shor, *Class of Global Minimum Bounds of Polynomial Functions*, Cybernetics **23**, 731–734 (1987).

[194] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).

[195] A. Ben-Tal and A. Nemirovski, *Lectures on modern convex optimization, analysis, algorithms and engineering applications* (MPS-SIAM, 2001).

[196] D. Jibetean and E. de Klerk, *Global optimization of rational functions: a semidefinite programming approach*, Mathematical Programming **106**, 93–109 (2006).

[197] B. Reznick, *Some concrete aspects of Hilbert's 17th Problem*, in F. Delon, M. Dickmann and D. Gondard (eds.), *Seminaire de structures algebriques ordonnees* (Equipe de logique mathematique, 1996), vol. 56 of *Math. Univ. Paris VII*.

[198] E. Artin, *Über die Zerlegung definiter Funktionen in Quatrate*, Hamb. Abh. **5**, 110–115 (1927).

[199] P. A. Parrilo, *Semidefinite programming relaxations for semialgebraic problems*, Mathematical Programming **96**, 293–320 (2003).

[200] G. Blekherman, *Convexity properties of the cone of nonnegative polynomials*, Discrete & Computational Geometry **32**, 345–371 (2004).

[201] D. Cantor, B. Gordon and B. Rothschild (eds.), *Theodore S. Motzkin: Selected Papers* (Birkhäuser, Boston, 1983).

[202] M. D. Choi, T. Y. Lam and B. Reznick, *Sum of squares of real polynomials*, Proceedings of symposia in pure mathematics **58**, 103–126 (1995).

[203] V. Powers and T. Wörmann, *An algorithm for sum of squares of real polynomials*, Journal of pure and applied algebra **127**, 99–104 (1998).

[204] P. A. Parrilo and S. Lall, *Semidefinite programming relaxations and algebraic optimization in control*, Lecture notes presented at IEEE Conference on decision and control, Mauii, Hawaii, U.S.A. (2003).

[205] D. Cox, J. Little and D. O'Shea, *Ideals, Varieties, and Algorithms* (Springer-Verlag, Berlin-Heidelberg, New York., 1997), 2nd edn.

[206] I. Z. Emiris, E. D. Fritzilas and D. Manocha, *Algebraic algorithms for structure determination in biological chemistry*, International Journal of Quantum Chemistry (2006).

[207] J. Bochnak, M. Coste and M.-F. Roy, *Real Algebraic Geometry* (Springer, 1998).

[208] G. Stengle, *A Nullstellensatz and a Positivstellensatz in semialgebraic geometry*, Math. Ann **207**, 87–97 (1974).

[209] *Doctoral school on: Optimization over Polynomials and Semidefinite Programming*, University of Klagenfurt (Austria), Sept. 12-16 (2005).

[210] S. Prajna, A. Papachristodoulou and P. Parrilo, *SOSTOOLS, Sum of Squares Optimization Toolbox for MATLAB, User's Guide, Version 1.00*, Dept. of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125 - USA (2002).

[211] J. Sturm, *Using SEDUMI 1.02, A Matlab Toolbox for Optimisation Over Symmetric Cones*, Dept. of Econometrics, Tilburg University, Tilburg, The Netherlands (2001).

[212] M. Yamashita, K. Fujisawa and M. Kojima, *Implementation and evaluation of SDPA 6.0 (SemiDefinite Programming algorithm 6.0)*, Optimization methods & software **18**, 491–505 (2003).

[213] B. Borchers, *CSDP, a C library for semidefinite programming*, Optimization methods & software **11**, 613–623 (1999).

[214] K. Fukuda, *CDD software implementing the Double Description Method of Motzkin*, http://www.ifor.math.ethz.ch/ fukuda/cdd_home/cdd.html.

[215] R. Olender and R. Elber, *Calculation of classical trajectories with a very large time step: Formalism and numerical examples*, J. Chem. Phys. **105**, 9299–9315 (1996).

[216] D. Passerone and M. Parrinello, *Action-derived molecular dynamics in the study of rare events*, Phys. Rev. Lett. **87**, 108302 (2001).

[217] M. G. Burke and S. N. Yaliraki, *Exploring model energy and geometry surfaces using sum of squares decompositions*, Journal of Chemical Theory and Computation **2**, 575–587 (2006).

[218] J. A. Nelder and R. Mead, *A Simplex Method for Function Minimization*, Comp. J. **7**, 308–313 (1965).

[219] K. Müller, *Reaction paths on multidimensional energy hypersurfaces*, Angew. Chem. Int. Ed. Engl. **19**, 1–78 (1980).

[220] J. Wang, R. Wolf, J. Caldwell, P. Kollman and D. Case, *Development and testing of a general amber force field*, Journal of Computational Chemistry **25**, 1157–1174 (2004).

[221] W. Humphrey, A. Dalke and K. Schulten, *VMD – Visual Molecular Dynamics*, Journal of Molecular Graphics **14**, 33–38 (1996).

[222] K. C. Toh, *Solving large scale semidefinite programs via an iterative solver on the augmented systems*, Tech. Rep. B-388, Dept. of Mathematical Sciences, Tokyo Institute of Technology, Toyko, Japan (2002).

[223] *CPLEX Mathematical Programming*, http://www.ilog.com/products/cplex/.

[224] B. Borchers and J. Young, *How Far Can We Go With Primal-Dual Interior Point Methods for SDP?*, Computational Optimization and Applications (2005), accepted for publication.

[225] K. Nakata, M. Yamashita, K. Fujisawa and M. Kojima, *A parallel primal-dual interior-point method for semidefinite programs using positive definite matrix completion*, Parallel Computing (2006).

[226] C. Lemarechal and C. Sagastizabal, *Variable metric bundle methods: From conceptual to implementable forms*, Mathematical Programming **76**, 393–410 (1997).

[227] H. C. and F. F. Rendl, *A Spectral Bundle Method for Semidefinite Programming*, SIAM Journal on Optimization **10**, 673–696 (2000).

[228] F. Rendl, *Semidefinite Programming Algorithms*, Presentation at ADONET doctoral school on sum of squares optimization, University of Klagenfurt, Austria (2005).

[229] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren and J. Hermans, *Intermolecular Forces* (Reidel, Dordrecht, 1981), 331.

[230] M. W. Mahoney and W. L. Jorgensen, *A five-site model for liquid water and the reproduction of the density anomaly by rigid, nonpolarizable potential functions*, J. Chem. Phys. **112**, 8910–8922 (2000).

[231] W. Still, A. Tempczyk, R. Hawley and T. Hendrickson, *Semianalytical Treatment of Solvation for Molecular Mechanics and Dynamics*, J. Am. Chem. Soc. **112**, 6127–6129 (1990).

[232] M. R. Wilson, *Progress in computer simulations of liquid crystals*, International Reviews in Physical Chemistry **24**, 421–455 (2005).

[233] J. Corner, Proceedings of the Royal Society of London A **192**, 275 (1948).

[234] T. Kihara, Reviews of Modern Physics **25**, 831 (1953).

[235] B. J. Berne and P. Pechukas, *Gaussian Model Potentials for Molecular-Interactions*, Journal of Chemical Physics **56**, 4213 (1972).

[236] S. H. Walmsley, *Gaussian Model Potentials for Molecular-Interactions and their Development*, Chemical Physics Letters **49**, 390–392 (1977).

[237] J. G. Gay and B. J. Berne, *Modification of the Overlap Potential to Mimic a Linear Site-Site Potential*, Journal of Chemical Physics **74**, 3316–3319 (1981).

[238] G. R. Luckhurst and P. S. J. Simmonds, *Computer-Simulation Studies of Anisotropic Systems. 21. Parametrization of the Gay-Berne Potential for Model Mesogens*, Molecular Physics **80**, 233–252 (1993).

[239] G. Ayton, S. G. Bardenhagen, P. McMurtry, D. Sulsky and G. A. Voth, *Interfacing continuum and molecular dynamics: An application to lipid bilayers*, Journal of Chemical Physics **114**, 6913–6924 (2001).

[240] A. Liwo, S. Oldziej, M. R. Pincus, R. J. Wawak, S. Rackovsky and H. A. Scheraga, *A united-residue force field for off-lattice protein-structure simulations .1. Functional forms and parameters of long-range side-chain interaction potentials from protein crystal data*, Journal of Computational Chemistry **18**, 849–873 (1997).

[241] J. W. Perram and M. S. Wertheim, *Statistical-Mechanics of Hard Ellipsoids .1. Overlap Algorithm and the Contact Function*, Journal of Computational Physics **58**, 409–416 (1985).

[242] J. W. Perram, J. Rasmussen, E. Praestgaard and J. L. Lebowitz, *Ellipsoid contact potential: Theory and relation to overlap potentials*, Physical Review E **54**, 6565–6572 (1996).

[243] *MOSEK SOCP software*, http://www.mosek.com/.

[244] E. D. Andersen, C. Roos and T. Terlaky, *On implementing a primal-dual interior-point method for conic quadratic optimization*, Mathematical Programming **95**, 249–277 (2003).

[245] J. J. Nicholas, K. E. Gubbins, W. B. Streett and D. J. Tildesley, *Equation of State for the Lennard-Jones Fluid*, Molecular Physics **37**, 1429–1454 (1979).

[246] J. G. Powles, *The Liquid Vapor Coexistence Line for Lennard-Jones-Type Fluids*, Physica a **126**, 289–299 (1984).

[247] R. W. Hockney, *The potential calculation and some applications*, Methods in computational physics **9**, 136–211 (1970).

[248] I. P. Omelyan, *Algorithm for numerical integration of the rigid-body equations of motion*, Physical Review E **58**, 1169–1172 (1998).

[249] L. Paramonov and S. N. Yaliraki, *Application of New Elliptic Contact Potential to Modelling of Amphiphilic Membranes.*, submitted to Journal of Chemical Physics (2006).

[250] B. Kumar, N. Chandra and G. Singh, *New 12-6 intermolecular potential for ellipsoidal molecules*, Physics Letters A **262**, 174–178 (1999).

[251] J. F. Sturm, *Implementation of interior point methods for mixed semidefinite and second-order cone optimization problems*, Optimization Methods and Software **17**, 1105–1154 (2002).

[252] E. D. Andersen, C. Roos and T. Terlaky, *On implementing a primal-dual interior-point method for conic quadratic optimization.*, Tech. Rep. W-274, Helsinki School of Economics and Business Administration, Department of Economics and Management Science,, PL 1210, FIN-00101 Helsinki, Finland (2000.).

[253] F. Alizadeh, J. P. Haeberly and M. L. Overton, *Primal-dual interior-point methods for semidefinite programming: Convergence rates, stability and numerical results*, SIAM Journal on Optimization **8**, 746–768 (1998).

[254] Y. Zhang, *On extending some primal-dual interior-point algorithms from linear programming to semidefinite programming*, SIAM Journal on Optimization **8**, 365–386 (1998).

[255] Y. E. Nesterov and M. J. Todd, *Primal-dual interior-point methods for self-scaled cones*, SIAM Journal on Optimization **8**, 324–364 (1998).

[256] R. H. Tutuncu, K. C. Toh and M. J. Todd, *Solving semidefinite-quadratic-linear programs using SDPT3*, Mathematical Programming **18**, 491–505 (2003).

[257] R. D. Monteiro, *Primal-dual path-following algorithms for semidefinite programming*, SIAM Journal on Optimization **7**, 663–678 (1997).

[258] M. Kojima, S. Shindoh and S. Hara, *Interior-point methods for the monotone semidefinite linear complementarity problem in symmetric matrices*, SIAM Journal on Optimization **7**, 86–125 (1997).

[259] C. Helmberg, F. Rendl, R. J. Vanderbei and H. Wolkowicz, *An interior-point method for semidefinite programming*, SIAM Journal on Optimization **6**, 342–361 (1996).