University of Southampton

Faculty of Natural and Environmental Sciences

Chemistry

Fundamental Advances in Hybrid Classical and Quantum Free Energy Techniques

by

Christopher Cave-Ayland

Thesis for the degree of Doctor of Philosophy

November 2014

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF NATURAL AND ENVIRONMENTAL SCIENCES
Chemistry
Doctor of Philosophy

Fundamental Advances in Hybrid Classical and Quantum Free Energy Techniques

by Christopher Cave-Ayland

The efficient and accurate calculation of free energies of binding and hydration for small molecules remains a valuable goal in the field of computational chemistry. We focus on free energy calculations as a tool to predict free energies of binding and hydration. Such calculations require the use of an energy model to compute the relative energies of different configurations of a molecular system. Quantum mechanical (QM) energy models offer excellent accuracy but at great computational cost, whereas classical energy models are more approximate but offer several orders of magnitude of computational savings. Here we describe as hybrid those approaches that attempt to combine both classical and quantum energy models to calculate free energies.

We demonstrate a fundamental validation step for the classical to quantum single step free energy perturbation approach (SSFEP), by carrying out the reverse quantum to classical perturbation. This required the generating of extensive trajectories using molecular dynamics at the QM level, a costly and involving process. With this data we demonstrate convergence between forward and reverse perturbations between the different levels of theory, providing support for use of the technique. Such convergence is dependent on the use of interaction energies when evaluating the free energy difference, a practicality rarely alluded to in the literature. We go on to elucidate, from first principles, the approximation engendered by the use of interaction energies with theoretical work and demonstrative calculations in a series of model system.

A number of recently proposed and more sophisticated hybrid techniques are considered and compared with SSFEP and a novel approach, proposed by us. The extensive QM trajectories generated previously are used again to a system within which these techniques can be compared and benchmarked. We also consider extension of each of these techniques to the use of interaction energies and demonstrate results of superior convergence.

Finally, we set up and carry out simulation work with an isoform of the CD1 protein family. This forms preparatory work for the application of hybrid free energy techniques within a protein-ligand system of clinical interest.

A thesis *must* be long. The object, you see, is to bore and stupefy the examiners to such an extent that they will *have* to accept it - only if a thesis is short enough to be read all through word for word is there any danger of failure.

Less than Angels, Barbara Pym

Contents

List of Figures

List of Tables

Academic Thesis: Declaration Of Authorship

I, CHRISTOPHER CAVE-AYLAND [please print name]

declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

FUNDAMENTAL ADVANCES IN HYBID CLASSICAL

AND QUANTUM FREE ENERGY TECHNIQUES

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Either none of this work has been published before submission, or parts of this work have been published as: [please list references below]:

DIRECT VALIDATION OF THE SINGLE STEP CLASSICAL TO
QUANTUM FREE ENERGY PERTURBATION, J. PHYS. CHEM. B
DOI: 10.1021/jp506459v

Signed: ....

Date: ....

Acknowledgements

I would like to express my deep thanks to both my supervisors, Professor Jonathan Essex and Dr Chris-Kriton Skylaris, for their invaluable input and encouragement throughout the last four years. Both their depth of knowledge and dedication are truly inspiring and I look forward to continuing to work with them in the year to come. Amongst my friends and colleagues of the Essex group, I must particularly thank Sam Genheden, for the immeasurably edifying journal club, and the other postdocs for their technical help. A special mention to Professor Seth Bullock for the chance to be a part of the complexity science doctoral training centre. Many of my friends and colleagues within the DTC have been a great inspiration and have always provided a supportive atmosphere.

Outside the academic world, I owe so much to Eleanor for the unstinting love and support and for never letting things seem bleak. I can't imagine having had to go through this process alone. Love you bae.

All my friends for tolerating the unanswered messages and generally being great. Maike, thanks for the sofa and letting me turn up and complain once a week.

Finally, I thank whichever twisted god rules technology and computer hardware for protecting me from any catastrophic technical failures. May my backups remain forever uncorrupted.

List of Abbreviations

AMBER - Assisted Model Building with Energy Refinement
AM1-bcc - Austin Model 1 Bond Charge Correction
BAR - Bennett's Acceptance Ratio
BFGS - Broyden-Fletcher-Goldfarb-Shannon
COX2 - Cyclooxygenase 2
CPU - Central Processing Unit
CUDA - Compute Unified Device Architecture
DFT - Density Function Theory
DNA - Deoxyribose Nucleic Acid
FDTI - Finite Difference Thermodynamic Integration
FEP - Free Energy Perturbation
FFT - Fast Fourier Transform
GAFF - General Amber Force field
GGA - Generalised Gradient Approximation
GPU - Graphics Processing Unit
HF - Hartree-Fock
HRE - Hamiltonian Replica Exchange
LDA - Local Density Approximation
LJ - Lennard-Jones
MC - Monte Carlo
MD - Molecular Dynamics
MHC - Major Histocompatibility Complex
MM - Molecular Mechanics
MMBIF - Molecular Mechanics Based Importance Function
MOE - Molecular Operating Environment
MPM - Mannosyl-β1-Phosphomycoketide
MTMC - Multi 'Timestepping' Monte Carlo
MUE - Mean Unsigned Error
NVE - Number of atoms, Volume, Energy
NVT - Number of atoms, Volume, Temperature
OI - Overlap Integral
OPLS - Optimized Potentials for Liquid Simulations
PBE - Perdew-Burke-Ernzerhof
PEG - Polyethylene glycol

PME - Particle Mesh Ewald
PMF - Potential of Mean Force
PMFC - Potential of Mean Force of Constraint
QM - Quantum Mechanics
RESP - Restrained Electrostatic Potential
RETI - Replica Exchange Thermodynamic Integration
RMSD - Root Mean Squared Deviation
SCF - Self Consistent Field
SSFEP - Single Step Free Energy Perturbation
TCR - T-cell Receptor
TI - Thermodynamic Integration
XC - Exchange Correlation

1. Introduction

The accurate and rapid prediction of free energies of binding and hydration for small molecule targets remains a long sought goal in the field of computational chemistry[1]. A range of different techniques have been developed to tackle this problem, the most accurate of which make use of extensive Molecular Dynamics (MD) or Monte Carlo (MC) sampling and rigorously derived free energy difference estimators[2,3,4]. Two factors limit the accuracy of these free energy techniques: the realism of the energy model used to describe the potential energy surface of interest, and achieving a sufficient degree of sampling of the system to obtain converged ensemble average statistics. Highly realistic i.e. first principles Quantum Mechanics (QM) based energy models are able to accurately model a system's potential energy surface but are prohibitively expensive to undertake sufficient sampling of even moderately sized systems. This practical restriction generally necessitates the use of classically inspired Molecular Mechanics (MM) force fields. Although computationally far cheaper, the approximate and parameterized nature of MM methods places inherent restrictions on the achievable accuracy of calculations using MM potentials.

The dichotomy between MM and QM approaches has led to the development of hybrid methods that attempt to exploit the accuracy of QM models at a fraction of the computational cost, through judicious combination with MM potentials[5,6,7,8,9]. These are known as hybrid free energy techniques and are based around a variety of different approaches described fully in section 2.5.

In chapter 3 we consider the application of the mostly widely employed hybrid approach, the single step free energy perturbation. Despite numerous occurrences[6,10,11,12,13,14,7,15,16,17,9] within the literature, the ability to rigorously calculate converged free energy differences has yet to be demonstrated. We approach this problem directly by generating extensive molecular dynamics trajectories at the QM level of theory. The considerable computational cost of this approach has previously prevented the carrying out of this important validation step.

One of our key findings is the necessity of the use of interaction energies, in the place of total potential energies when employing hybrid free energy techniques. In chapter 4 the theoretical implications of this are fully explored from a statistical mechanics basis and it is shown to introduce an error term. This is explored in a series of model systems of increasing complexity.

Chapter 5 deals with reweighting and ensemble building hybrid free energy techniques. These are applied to the previously considered model systems of chapter 4 and the DNA base pair MD dataset of chapter 3. Particular attention is paid to the combination of these approaches with interaction energies and some theoretical results to this end are given in addition to the computational work.

Finally, we consider preparation and MD work with a proposed crystal structure for an isoform of the CD1 protein family. This serves as preparatory work for the future application of hybrid free energy techniques with a CD1 isoform of interest.

This thesis attempts to place hybrid classical and quantum free energy techniques in a rigorous context. We robustly examine the convergence behaviour and theoretical shortcomings of currently accepted practice and consider extension of more recently proposed approaches.

2. Theoretical Background

2.1. Statistical Mechanics

The foundation of the computational calculation of free energies is rooted in statistical mechanics. This is the theoretical branch of physics that deals with the prediction of high level (macrostate) properties of molecular systems, e.g. temperature/pressure, from low level (microstate) information e.g. atomic velocities and positions[18]. A simple example is the calculation of the temperature of a system from the velocities of its constituent N particles:

$$T = \frac{2K(\mathbf{v})}{3Nk_B} \tag{2.1}$$

where the temperature, T, is proportional to the total kinetic energy of the system, K a function of the system velocities $\mathbf{v}$, over N and the Boltzmann constant k_B. The classical kinetic energy term being given by:

$$K(\mathbf{v}) = \frac{1}{2}\sum_{i=1}^{N} m_i|\mathbf{v}_i|^2 \tag{2.2}$$

the kinetic energy is a sum over all N particles where m_i and $\mathbf{v}_i$ are respectively the mass and Cartesian velocities of the i^{th} particle. Thus microstate knowledge of atomic velocities can be used to calculate the higher level macrostate system property of temperature. Throughout this introduction, as here, we will not attempt any first principles derivations of statistical mechanics results as these may be obtained from a range of texts[19,18,20,21,22]. Instead we will present and attempt to intuitively rationalise the various widely used expressions relevant to the calculation of free energy differences.

A central concept to the statistical mechanics of equilibrium systems is that of ensembles. To explain this we start by considering a thermodynamic system of N particles. In a three dimensional space each particle can be completely described by 6 degrees of freedom, its position in each dimension and its corresponding velocity in each. Thus a system of N particles can be described by 6N degrees of freedom in total. The positional degrees of freedom are limited by the size of the system, whilst the velocity degrees of freedom are unbounded. These 6N degrees of freedom can be represented in a phase space, also referred to as a state space, of 6N dimensions[19]. Any given configuration of positions and velocities forms a point within the phase space, which is bounded as described previously. The phase space describes every possible state of the system at hand. However, equilibrium thermodynamics places restrictions on how the phase space will be occupied. This leads neatly to the concept of an ensemble which defines a set of restrictions or conditions on how a system will occupy its associated phase space[18]. The term ensemble is often also used to refer to the set of system states that conform to the correct properties of the associated thermodynamic ensemble.

The simplest example is known as the microcanonical ensemble, often abbreviated as *NVE* for the three restrictions it imposes. *N*, the number of particles, *V*, the volume and *E*, the energy of the system are all constant. A particular *NVE* ensemble is composed of all the system configurations within the phase space that meet these three conditions. In the case of the *NVE* ensemble there are points of the phase space that are inaccessible due to the restriction on the energy of the system. This can be expressed formally by the probability density[21]:

$$P_{NVE}(\mathbf{x},\mathbf{v}) \propto \delta(H(\mathbf{x},\mathbf{v}) - E) \tag{2.3}$$

where P_{NVE} is the probability density of a particular system configuration given by $\mathbf{v}$ and $\mathbf{x}$ within the *NVE* ensemble. δ is the Dirac delta function, that has a value of infinity where its argument is zero, but a value of zero everywhere else. When integrated over however the Dirac delta gives a value of one. H is the system Hamiltonian, the formal name for the function that takes the system state as input and gives its corresponding energy. The energy of the system is dependent on the positional degrees of freedom, $\mathbf{x}$, and the velocities, $\mathbf{v}$. We will consider the properties of the Hamiltonian later in this section and in section 2.2. Equation (2.3) defines a hyper-surface within the full phase space of the system and displays an important property of the microcanonical ensemble. The Dirac delta function has a constant value for all system states with the correct energy. Therefore, all structures that comprise the *NVE* ensemble have an equal probability of arising. The probability density can be made absolute by introducing a normalisation constant, known in this context as the partition function, Q_{NVE}[22]:

$$P_{NVE}(\mathbf{x},\mathbf{v}) = \frac{\delta(H(\mathbf{x},\mathbf{v}) - E)}{\mathrm{Q}_{NVE}} \tag{2.4}$$

$$\mathrm{Q}_{NVE} = \int_V \delta(H(\mathbf{x},\mathbf{v}) - E)d\mathbf{x}d\mathbf{v} \tag{2.5}$$

the partition function is simply given as an integral of all the degrees of freedom within the system over the relative probability density. Here V denotes the limits of integration imposed by the volume of the system. The microcanonical partition function corresponds to the volume of phase space containing structures that comprise the ensemble. Although of fundamental significance, the NVE ensemble rarely corresponds to systems of interest. Biological systems in particular are not isolated from fluctuations in energy. Rather these systems are often assumed to be accurately modelled as being in temperature equilibrium with their surroundings by the canonical or NVT ensemble. In addition to the particle number and volume restrictions, T signifies that the temperature of the system is constant. As given in equation (2.1) the instantaneous temperature is shown to be a function of particle velocities. However, the temperature of a system is a time-averaged property such that the kinetic energy need not be restricted to being the same value throughout the ensemble. In contrast to the microcanonical ensemble where certain system states are inaccessible for particular values of E, in the canonical ensemble any position within the entire phase space can be occupied with the following probability density:

$$P_{NVT}(\mathbf{x},\mathbf{v}) = \frac{\exp[-\beta H(\mathbf{x},\mathbf{v})]}{\mathrm{Q}_{NVT}} \tag{2.6}$$

$$\mathrm{Q}_{NVT} = \int_V \exp[-\beta H(\mathbf{x},\mathbf{v})] d\mathbf{x} d\mathbf{v} \tag{2.7}$$

where $\beta = \frac{1}{k_B T}$, a common abbreviation that will be used extensively throughout this work. This is the famous Boltzmann distribution, with the numerator of equation (2.6) referred to as the Boltzmann factor. The canonical partition function is similarly given as an integral over the relative probability density of the system. One particularly desirable property of the partition function is that it can be directly related to the free energy of a system[19]:

$$A = -\frac{1}{\beta} \ln Q_{NVT} \tag{2.8}$$

where A is defined as the Helmholtz free energy. The partition function in essence is the term that links microstate data of the system to the macrostate property free energy. Owing to the high dimensionality of the problem however, direct calculation of the partition function is impossible for all but the simplest of systems. Techniques that circumvent direct calculation of free energies are considered in detail in section 2.4, but we shall first consider some simplifications that can be applied to the problem. Returning to the definition of the system Hamiltonian, it is possible to give the decomposition:

$$H(\mathbf{x},\mathbf{v}) = U(\mathbf{x}) + K(\mathbf{v}) \tag{2.9}$$

where K is the kinetic energy of the system as defined in equation (2.2) and U is the potential energy. The calculation of U based on the atomic configuration of the system is an interesting problem, the details of which are discussed extensively in section 2.2. For now, observing the ability to split the Hamiltonian into separate potential and kinetic terms is sufficient. Using equation (2.9), we can reconsider our definition of the canonical partition function:

$$\begin{aligned}
\mathrm{Q}_{NVT} &= \int_V \exp\left[-\beta H(\mathbf{x},\mathbf{v})\right] d\mathbf{x}d\mathbf{v} && (2.10)\\
&= \int_V \exp\left[-\beta (U(\mathbf{x}) + K(\mathbf{v}))\right] d\mathbf{x}d\mathbf{v} && (2.11)\\
&= \int_V \exp\left[-\beta U(\mathbf{x})\right] \exp\left[-\beta K(\mathbf{v}))\right] d\mathbf{x}d\mathbf{v} && (2.12)\\
&= \int_V \exp\left[-\beta U(\mathbf{x})\right] d\mathbf{x} \int \exp\left[-\beta K(\mathbf{v})\right] d\mathbf{v} && (2.13)\\
&= Z \times Q_K && (2.14)
\end{aligned}$$

the first of these integrals, over the atomic positions of the system is often referred to as the configurational integral and denoted by Z. The integral over the velocities on the other hand can be treated analytically. This is straightforward as the kinetic energy term, equation (2.2), is a sum over the contributions of each of the atoms. The analytical result is given by[19]:

$$\begin{aligned}
Q_K &= \int \exp\left[-\beta \sum_i^N \frac{m_i |\mathbf{v}_i|^2}{2}\right] d\mathbf{v} && (2.15)\\
&= \int \prod_i^N \exp\left[-\beta \frac{m_i |\mathbf{v}_i|^2}{2}\right] d\mathbf{v} && (2.16)\\
&= \left(\prod_i^N 2\pi m_i \frac{1}{\beta}\right)^{\frac{3}{2}} && (2.17)
\end{aligned}$$

the last term being derived from the general form of integration over a Gaussian function. The ability to treat components of the free energy analytically helps to simplify the problem of calculating free energy differences by removing a source of statistical error. In the rest of this work we will no longer consider the inclusion of the kinetic energy component in the calculation of free energy differences but proceed with almost all derivations and explanations in terms of the potential energy component and the configurational integral. Excluding the kinetic component, the free energy difference between states can be given as:

$$\Delta A_{1\to2} = A_2 - A_1 = \ln \frac{Z_2}{Z_1} \qquad (2.18)$$

thus the ratio of the partition functions of the two states is often the sought quantity in evaluating $\Delta A_{1\to2}$. Here, as throughout this work, the configurational contribution to the free energy is taken as being synonymous with the full free energy. This implicitly assumes cancellation of the kinetic terms. Where this is not the case appropriate corrections can be given using the analytical kinetic result.

Finally, we introduce the concept of an ensemble average:

$$\langle X \rangle = \frac{\int X(\mathbf{x}) \exp\left[-\beta U(\mathbf{x})\right] d\mathbf{x}}{\int \exp\left[-\beta U(\mathbf{x})\right] d\mathbf{x}} \tag{2.19}$$

$$= \frac{\int X(\mathbf{x}) \exp\left[-\beta U(\mathbf{x})\right] d\mathbf{x}}{Z} \tag{2.20}$$

this may be better explained as the expectation value of the system property X, which is some function of $\mathbf{x}$. In principle X could relate to any possible property of the system, for instance a particular bond length or an energy term. Unlike the partition function, evaluating ensemble averages is computationally feasible and methods for this are considered in section 2.3. The uses of ensemble averages in calculating free energy differences are explored in section 2.4.

A free energy calculation is dependent on three separate components working together[23]. First, an energy function or model that realistically describes the potential energy surface of the system of interest is combined with a sampling protocol that generates an ensemble of configurations with the appropriate Boltzmann distribution and finally a expression estimating the free energy difference is required. The following sections provide a comprehensive overview of modern approaches to all of these components.

2.2. Energy Models

An energy model or potential energy function should comprise a function that takes the atomic positions of the system as input and gives a value for the potential energy. This is often used synonymously with the term Hamiltonian to distinguish thermodynamic states as the kinetic energy term is identical between different energy models.

For theoretical calculations to be correct the energy model of a system must accurately recreate the energies of the real system, i.e. it must be realistic. Speaking broadly, energy models can be divided into two groups, those that employ classical theory (also known as molecular mechanics or MM) and those that make use of quantum mechanics (QM), also known as *ab initio.* These employ different underlying theories to describe the potential energy surface of a system and the combination of MM and QM energy models within free energy calculations forms the basis of the ‘hybrid’ nature of the work presented in this thesis.

2.2.1. Quantum Mechanics

QM theory provides access to a class of potential energy functions of very high achievable accuracy but at great computational expense. In contrast to classical theory QM calculations deal explicitly with the electronic component of a molecular system. The field of electronic structure theory is vast, with a huge array of techniques proposed displaying varying degrees of accuracy and computational expense. What follows is a very restricted discussion of modern QM approaches that will focus on laying out the problem at hand followed by an examination of early Hartree-Fock and modern plane wave DFT approaches.

At its core, quantum chemistry can be reduced to a constrained minimisation problem[24] for the electronic component of the system. For a given configuration of the system nuclei, the task at hand is to find the corresponding structure of the electronic degrees of freedom that minimise the system energy, subject to certain constraints. It is by now a well-known phenomenon that the QM description of systems, includes many effects that seem counter-intuitive from the classical perspective. We will develop therefore an appropriate mathematical description of the electronic system components such that the underlying physics is reflected. As such we will not attempt to rationalise the physical properties of particles at the QM level but take them as criteria to be met by our description. Once a suitable mathematical form has been developed we shall consider methods to find minimised configurations within the electronic configuration space.

2.2.1.1. Electron wavefunctions

We start with the appropriate description for an isolated electron. Although classically understood as a particle, the electron is known to obey the Heisenberg uncertainty principle. This states that an electron does not simultaneously have a well defined position and momentum. The treatment of electrons as particles as in classical theory is therefore precluded[25]. Instead electrons are described in terms of a wavefunction, denoted by $\psi(\mathbf{r})$. Here $\mathbf{r}$ corresponds to the Cartesian spatial degrees of freedom, within which the electrons reside. An exact physical interpretation of the wavefunction is still an open question in QM theory. An intuitive one that we shall consider here is the Born interpretation, that the wavefunction of an electron can be related to its probability density:

$$P(\mathbf{r}) = |\psi(\mathbf{r})|^2 \tag{2.21}$$

If $P(\mathbf{r})$ is a probability density then the constraint is placed upon the wavefunction such that:

$$\int P(\mathbf{r})d\mathbf{r} = \int |\psi(\mathbf{r})|^2 d\mathbf{r} = 1 \tag{2.22}$$

i.e. as the square of the wavefunction is a probability density, the integral over all possible states must give a total probability of 1. Our description of the single electron is still not complete however, as we have yet to account for the property of spin. The spin state of an electron can be either up or down, this can be elegantly handled by introducing the auxiliary variable ω, and the two spin functions $\alpha(\omega)$ and $\beta(\omega)$. Such that:

$$\int |\alpha(\omega)|^2 d\omega = \int |\beta(\omega)|^2 d\omega = 1 \tag{2.23}$$

$$\int \alpha^*(\omega)\beta(\omega)d\omega = \int \beta^*(\omega)\alpha(\omega)d\omega = 0 \tag{2.24}$$

where * denotes the complex conjugate of a wavefunction. This can be stated as the requirement that $\alpha(\omega)$ and $\beta(\omega)$ are orthonormal. A spin orbital can now be defined based on the combination of a spatial wavefunction, given by $\psi(\mathbf{r})$ and a spin function[24]:

$$\chi(\mathbf{x}) = \psi(\mathbf{r})\alpha(\omega) \tag{2.25}$$

note we have here used $\mathbf{x}$ to denote the combination of $\mathbf{r}$ and ω. This is in contrast to the rest of this work, where atomic positions are denoted by $\mathbf{x}$. However, this has been adopted for the alternative purpose here to conform to existing conventions. A spin orbital can be made with either spin function and any spatial wavefunction.

Having defined an appropriate description of the single electron we now consider the extension of this to a many-electron wavefunction, represented as:

$$\Psi(\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N) = \frac{1}{\sqrt{(N!)}} \begin{vmatrix} \chi_1(\mathbf{x}_1) & \chi_2(\mathbf{x}_1) & ... & \chi_n(\mathbf{x}_1) \\ \chi_1(\mathbf{x}_2) & \chi_2(\mathbf{x}_2) & ... & \chi_n(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \chi_1(\mathbf{x}_N) & \chi_2(\mathbf{x}_N) & ... & \chi_n(\mathbf{x}_N) \end{vmatrix} \tag{2.26}$$

where Ψ is an n-electron wavefunction, n being the number of electrons in the system. In practice this is one of various ways to form a many-electron wavefunction from single electron wavefunctions and is known as a Slater determinant[19]. Its form is based on the need for many electron wavefunctions to meet the anti-symmetry principle. This requires that:

$$\Psi(\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N) = -\Psi(\mathbf{x}_2, \mathbf{x}_1, ..., \mathbf{x}_N) \tag{2.27}$$

i.e. exchanging the position of any two electrons should change the sign of the wavefunction. This statement reduces to the well known Pauli exclusion principle, that two electrons cannot occupy the same quantum state. For two electrons to occupy the same spatial orbital they must have different spins. The combination of a Slater determinant and the previously defined spin functions satisfies this requirement.

Up to this point we have discussed the properties of wavefunctions but have yet to describe the form that they take in practice. One common choice is to represent single-electron wavefunctions within molecules as a linear combination of so-called atomic orbitals. These are a collection of simple wavefunctions, of tractable form, referred to as a basis and centred on the nuclei of the system. The spatial component of single-electron molecular orbitals, ψ, are then expressed as:

$$\psi_1(\mathbf{x}) = \sum_{i=1}^{N_b} c_{i1}\phi_i(\mathbf{x}) \tag{2.28}$$

where the basis of atomic orbitals is denoted by ϕ, N_b is the number of atomic orbitals and each atomic orbital is weighted by c. This allows us to consider an electronic configuration space made up by all possible values of the coefficient matrix $\mathbf{c}$. We can also express our problem as equated to finding the values of $\mathbf{c}$ that minimise the electronic energy. In principle ψ can be represented exactly in this way (when the basis is said to be complete). This would require an infinite number of basis functions. Instead a sufficiently large basis such as to be nearly complete is employed in practical calculations. Many choices are available for ϕ, most commonly hydrogen orbitals (which can be derived analytically), Slater orbitals or Gaussian functions[19]. An alternative to the linear combination of atomic orbitals approach is given in section 2.2.1.7.

2.2.1.2. QM Hamiltonian

An expression is needed to describe the interactions between different components of the system. This is given by the QM Hamiltonian, in atomic units[24]:

$$\hat{H} = -\frac{1}{2}\sum_{i=1}^{N}\nabla_i^2 - \frac{1}{2}\sum_{I=1}^{M}\frac{1}{m_I}\nabla_I^2 - \sum_{i=1}^{N}\sum_{I=1}^{M}\frac{q_I}{r_{iI}} + \sum_{i=1}^{N}\sum_{j>i}^{N}\frac{1}{r_{ij}} + \sum_{I=1}^{M}\sum_{J>I}^{M}\frac{q_I q_J}{r_{IJ}} \tag{2.29}$$

here lower case subscripts denote indices of electrons, whilst upper case subscripts denote indices of nuclei. ∇^2 denotes the Laplacian operator, q_I is the atomic number of the Ith nuclei, r is the separation of the two subscripted particles, whilst M and N denote the number of nuclei and electrons in the system respectively. Thus the first two terms give the kinetic energy component of the system, the third gives the energy for the interaction between all electrons and nuclei, the fourth gives an electron-electron repulsion term and the last a nuclei-nuclei repulsion term.

This full Hamiltonian can be simplified by making use of the observation that the mass of the electron is much less than that of the nuclei and hence their movement is much more rapid. This separation of timescales allows us to make the assumption that for any given configuration of the nuclei, the electronic degrees of freedom will instantaneously relax. This is known as the Born-Oppenheimer (BO) approximation and allows formulation of an electronic Hamiltonian that considers the nuclear configuration of the system as a parameter for the electronic system components rather than as variables as in 2.29. This gives the electronic Hamiltonian:

$$\hat{H}_{elec} = -\frac{1}{2}\sum_{i=1}^{N}\nabla_i^2 - \frac{1}{2}\sum_{i=1}^{N}\sum_{I=1}^{M}\frac{q_I}{r_{iI}} + \sum_{i=1}^{N}\sum_{j>i}^{N}\frac{1}{r_{ij}} \tag{2.30}$$

where all terms not including the electrons from equation (2.29) have been dropped. Thus we derive a QM energy function that takes a nuclear configuration of the system as input and uses the electronic Hamiltonian to derive the electronic energy. The constant nuclear energy contributions in the full Hamiltonian can then be included. Use of the electronic Hamiltonian leads to a great simplification versus the full Hamiltonian and underpins modern QM approaches.

2.2.1.3. Schrödinger Equation

The stationary levels of a QM system is derived by finding solutions to the time-independent Schrödinger equation[25]:

$$E\Psi = \hat{H}_{elec}\Psi \tag{2.31}$$

here the electronic energy is given by E and $\hat{H}_{elec}$ is the electronic Hamiltonian expressed as an operator. This represents an eigenvalue problem such that the eigenvalues of $\hat{H}_{elec}$ represent the total electronic energies associated with the ground and various excited states of the system. The ground state energy of the system is given by the lowest eigenvalue of equation (2.31). We shall denote this eigenvalue as E_0 and the corresponding ground state wavefunction as Ψ_0. Finding these eigenvalues corresponds to finding the ground state energy and excited states of the system. A method to find the ground state energy is given by the variational principle. The corresponding electronic energy, $\tilde{E}$, of any trial wavefunction, $\tilde{\Psi}$ can be calculated using:

$$\tilde{E} = \int \tilde{\Psi}^*(\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n)\hat{H}_{elec}\tilde{\Psi}(\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n)d\mathbf{x}_1 d\mathbf{x}_2 ... d\mathbf{x}_N \tag{2.32}$$

the variational principle states that for any given trial wavefunction, $\tilde{E} \geq E_0$. This is in some ways obvious, E_0 being the global minimum of the electronic energy, by definition therefore any other point in the electronic configuration space must have a higher energy. However, we can use this to show that the relative quality of a trial wavefunction can be assessed by its energy, improved wavefunctions having lower energies than their more approximate counterparts.

We have now defined the constrained minimisation problem that constitutes electronic structure theory. Essentially we wish to minimize the electronic energy of a trial wavefunction, with respect to the weights of its basis set functions, **c**. This is subject to the constraint (so long as our basis functions are orthonormal) that the squared values of the columns of **c** sum to 1. In practice, however, this is an extremely difficult problem. Directly finding solutions to the electronic Hamiltonian is complicated by the presence of the electron-electron repulsion term. This constitutes a notoriously difficult many-body problem, making the electronic energy surface a complicated function of **c**. For this reason a large variety of techniques have been developed based around forming approximations to the true electronic Hamiltonian. In the following we shall derive Hartree-Fock theory, one of the first developed approaches, that conveniently demonstrates some of the fundamental properties of *ab initio* calculations. This will be followed by the more modern and popular Density Functional Theory (DFT), that is used extensively throughout this work.

2.2.1.4. Hartree-Fock Theory

As one of the first methods developed for the approximate solution of the Schrödinger equation[26], the Hartree-Fock method is frequently used as an introduction to the field of electronic structure theory. In fact, the Hartree-Fock equations arise naturally through the use of a Slater determinant to represent the many-electron wavefunction of the system. This is due to the fact that the Slater determinant is too simplistic to represent all possible many-body wavefunctions, making its use inherently approximate and leading to simplifications of the electronic Hamiltonian. More exact methods, such as full configuration interaction, make use of a linear combination of Slater determinants in order provide a complete description of the many-body wavefunction[24].

The fundamental effect of employing a Slater determinant is to replace the electron-electron term of equation (2.30) with a mean field term, such that each electron only encounters the average effect of every other electron within the system. This is given by the Fock operator $\hat{f}$:

$$\hat{f} = \hat{h} + \sum_{j}^{N} \hat{v}_j^C - \hat{v}_j^{Ex} \tag{2.33}$$

considering the various terms in more detail:

$$\hat{h} = -\frac{1}{2}\sum_{i}^{N} \nabla_i^2 - \sum_{i}^{N}\sum_{I=1}^{M} \frac{q_I}{r_{iI}} \tag{2.34}$$

this is the core Hamiltonian operator and contains all of the truly single electron terms of equation (2.30). These are preserved unchanged and are simple to solve. Next is the Coulomb operator:

$$v_j^C \chi_i(\mathbf{x}_1) = \int |\chi_j(\mathbf{x}_2)|^2 \frac{1}{r_{12}} d\mathbf{x}_2 \chi_i(\mathbf{x}_1) \tag{2.35}$$

as may be guessed, this accounts for the charge-charge interactions of the electrons. However this is no longer a many-body problem but a series of effective single-electron integrals. Electrons are considered as uncorrelated with one another and as will be discussed this is the major short-coming of the Hartree-Fock approach. The final term is known as the exchange operator:

$$v_j^{EX} \chi_i(\mathbf{x}_1) = \int \chi_j^*(\mathbf{x}_2)\chi_i(\mathbf{x}_2) \frac{1}{r_{12}} d\mathbf{x}_2 \chi_j(\mathbf{x}_1) \tag{2.36}$$

this term is a purely quantum effect, having no analogous classical equivalent, that arises from the electron anti-symmetry requirement. The forms of the Coulombic and exchange operators are noted to be dependent on the value of the wavefunction. For this reason the Hartree-Fock equations pose a non-linear problem and must be solved self-consistently, through an iterative procedure. This is known as a self-consistent field (SCF) procedure.

As mentioned above, the HF approach does not treat correlation effects between electrons. This causes HF to consistently overestimate the ground state energy. The correlation energy of a system is in fact defined as the difference between the best estimate of the ground-state using HF theory and the true ground state energy of the system. Despite its shortcomings however HF is still employed in molecular calculations.

2.2.1.5. Density Functional Theory

The calculations in this work are based around the method of Density Functional Theory (DFT)[27], a more modern technique, popularly used across a range of applications[28]. The fundamental quantity of DFT is the charge density, given by:

$$n(\mathbf{r}_1) = N \int \dots \int |\Psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)|^2 d\omega_1 d\mathbf{x}_2 \dots d\mathbf{x}_N \tag{2.37}$$

here we are carrying out an integration of the many-electron wavefunction of the system over all electronic degrees of freedom but the spatial of the first electron. As we have removed all spin variables the charge density is independent of spin. The charge density is a much simpler description of the electronic component of the system, but importantly it has been shown that the ground state wavefunction of a system can be uniquely mapped to a ground state density and vice-versa. In principle we can now define a series of functionals that act on the charge density to give the energy of the system:

$$E[n(\mathbf{r})] = T[n(\mathbf{r})] + V[n(\mathbf{r})] + U[n(\mathbf{r})] \tag{2.38}$$

the three terms from (2.30) are given by their functional equivalents, T for the kinetic energy of the electrons, V the electron-nucleus interaction and U the electron-electron term. Whilst the form of the V functional is trivially derived, there is no known expression for T or U. How this problem is handled has come to define the various forms of DFT. The most successful, and seemingly ubiquitous in modern computational chemistry approaches, is known as the Kohn-Sham formalism[29], given by:

$$T + U = T_0 + U_0 + U_{xc} \tag{2.39}$$

here T_0 is the functional giving the kinetic energy of a set of N non-interacting electrons:

$$T_0 = \frac{1}{2}\sum_{i}^{N}\int \psi_i \nabla^2 \psi_i d\mathbf{r} \tag{2.40}$$

whilst equivalently U_0 is the classical electron-electron Coulombic repulsion term:

$$U_0[n(\mathbf{r})] = \frac{1}{2}\int\int \frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}\mathrm{d}\mathbf{r}\mathrm{d}\mathbf{r}' \tag{2.41}$$

this can be directly compared with $\hat{v}^C$ and is often known as the Hartree energy. This leaves the final term U_{xc}, the exchange-correlation functional, the correct form of which also remains unknown. This term encompasses all of the non-classical exchange and correlation effects of electron interaction, as well as the correction to the kinetic energy. It may seem that the situation has not improved much. However this formulation has the significant advantage compared to equation (2.38) that the contribution of the unknown, U_{xc} is much smaller than the unknown terms T and U, the majority of the physics of the system being captured by U_0 and T_0. Using equation (2.39) and the variational principle gives the set of N Kohn-Sham equations:

$$\left[-\frac{1}{2}\nabla_i^2 - \sum_{I}^{M}\frac{q_I}{r_1 I} + \int \frac{n(\mathbf{r}_2)}{r_{12}}d\mathbf{r}_2 + V_{XC}(r_1)\right]\psi_i(\mathbf{r}_1) = \varepsilon_i \psi_i(\mathbf{r}_1) \tag{2.42}$$

$$V_{XC}(\mathbf{r}) = \frac{\delta U_{XC}(n(\mathbf{r}))}{\delta n(\mathbf{r})} \tag{2.43}$$

where ε_i is the energy of ith wavefunction, and the use of δ denotes a functional derivative. This is effectively the Schrödinger equation for a set of N non-interacting electrons in a fictional potential that gives the same ground state density as the potential of the real interacting system. A Slater determinant is able to exactly represent the many-electron wavefunction of a non-interacting system and hence is suitable for representation of the charge density under the Kohn-Sham formalism. As previously with Hartree-Fock theory, we note that the operator of equation (2.42) contains the electron density and is hence the wavefunction of the system. Thus solution of the Kohn-Sham equations may be performed iteratively with the SCF procedure.

Whilst QM theory offers greater potential accuracy than classical approaches, the associated computational cost is much larger. Both Hartree-Fock and DFT calculations bear a computational cost that scales as $\mathcal{O}(M^3)$. This has prohibited its wide-spread application in free energy calculations; until recently the required computing power has not been widely available. Recently, linear scaling formulations of DFT have been proposed, making tractable QM calculations of protein-ligand complexes comprising thousands of atoms.

2.2.1.6. Exchange-Correlation Functionals

In practice we are able to partially evaluate the exchange-correlation term. Formulation and improvement of novel approximations to U_{xc} is a major pre-occupation in the study of DFT. There are numerous constraints that can be applied to the form of the functional and the use of ideal reference systems is popular. The first and most basic of the proposed functionals is the Local Density Approximation (LDA)[30]:

$$U_{xc}^{LDA}[n(\mathbf{r})] = \int n(\mathbf{r})\varepsilon_{xc}[n(\mathbf{r})]d\mathbf{r} \tag{2.44}$$

where ε_{xc} is the exchange-correlation functional of the homogeneous electron gas of matching density. Using the exchange-correlation contribution of such an idealised reference system is obviously quite rudimentary and not particularly realistic. However it can provide surprising realism for many solid-state systems and forms an important starting point for the construction of more sophisticated functionals. Considering separately the two parts of ε_{xc}, the exchange contribution is known and can be expressed simply:

$$U_{x}^{LDA}[n(\mathbf{r})] = -\frac{3}{2}\left(\frac{3}{4\pi}\right)^{\frac{1}{3}} \int (n_\alpha(\mathbf{r}))^{\frac{4}{3}} + (n_\beta(\mathbf{r}))^{\frac{4}{3}} d\mathbf{r} \tag{2.45}$$

where n_α and n_β are the up spin and down spin contributions to the electron density. The correlation contribution is more problematic as no correct analytical expression is known, while various suggestions have been made based on parameterisation approaches[31]. This is discussed in slightly more detail in section 2.6.

Although use of the LDA functional provides a functional with the correct properties it is not frequently used for biomolecular simulations. Instead a class of functionals known as generalised gradient approximations (GGA) are more commonly used. These include terms involving the gradient of the electron density to provide a more powerful description of exchange and correlation. There are a great many different GGA functionals, and their forms are somewhat more complex than the LDA. For this reason we shall not consider them in detail, but content ourselves with noting the popular Perdew-Burke-Ernzerhof (PBE)[32] functional:

$$U_{c}^{PBE} = \int \varepsilon_c[n(\mathbf{r})] + H[n(\mathbf{r}), \nabla n(\mathbf{r})]d\mathbf{r} \tag{2.46}$$

$$U_{x}^{PBE} = \int \varepsilon_x[n(\mathbf{r})]F_x[n(\mathbf{r}), \nabla n(\mathbf{r})]d\mathbf{r} \tag{2.47}$$

this is developed from the uniform electron gas, where H and F_x are additional corrections dependent also on the gradient of the density, $\nabla n(\mathbf{r})$. The PBE functional has been shown to offer a good compromise between speed and accuracy in describing biological compounds[33], and is frequently used in this context[34,35,36,37,38].

2.2.1.7. Plane Wave Basis Sets

Whilst the linear combination of atomic orbitals approach provides an intuitive method for the construction of molecular orbitals a popular alternative has arisen through the use of plane waves. Heavily used in the context of solid-state systems, plane waves provide a natural and powerful description of periodic systems[30].

$$\psi(\mathbf{r}) = \sum_{\mathbf{G}} c_{\mathbf{G}} \exp[i\mathbf{G}.\mathbf{r}] \tag{2.48}$$

where $\mathbf{G}$ are defined as reciprocal lattice vectors, defining waves of different frequency and direction, $c_{\mathbf{G}}$ are complex coefficients and $i = \sqrt{-1}$. Unlike atomic centred basis sets, plane waves are evenly spaced throughout the volume of the simulation cell. As previously noted, to truly represent ψ, an infinite number of plane waves would be required so instead the number of plane waves is truncated at a set number. A great advantage of plane waves however, is that the size of the basis set can be easily and systematically varied. Typically the size of the basis is determined through use of the kinetic cut off energy:

$$\nabla^2 \exp[i\mathbf{G}.\mathbf{r}] = \frac{|G|^2}{2} \leq E_{cut} \tag{2.49}$$

where E_{cut} is the cut off value used for the kinetic energy term. The magnitude of $c_{\mathbf{G}}$ is inversely proportional to the magnitude of G and hence high energy plane waves can be excluded without greatly influencing the accuracy of ψ. In practice an appropriate value of E_{cut} must be found for any given system of interest.

As plane waves propagate throughout the entire volume of the system, the number of plane waves associated with a particular value of E_{cut} increases with the size of the simulation cell. For this reason, the computational cost of plane wave DFT implementations increase as a function of not only the number of electrons but also the size of the periodic box.

An additional shortcoming of plane waves is that they are poorly suited to describing electrons in non-valence orbitals[19]. These core electrons tend to require large numbers of plane waves to be accurately described and additionally the wavefunctions of valence electrons in the core region tend to oscillate rapidly also making them difficult to describe. Explicit treatment of the core electrons can be avoided by noting that they are largely invariant with respect to the environment of an atom. Thus rather than explicitly modelling the core electrons their effect is replaced by a modified potential that mimics their inclusion (see figure 2.1). Such pseudopotentials can be formulated on a number of criteria, one of these being norm conservation. This requires that the wavefunction generated with a pseudopotential, matches the density of the true wavefunction. In general the use of pseudopotentials considerably reduces

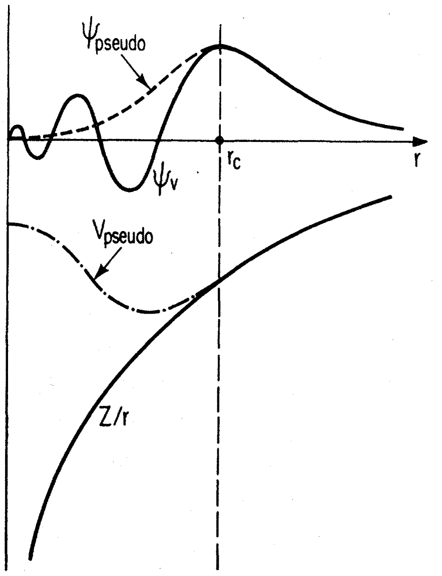

Figure 2.1.: Diagrammatic pseudopotential representation, taken from Payne *et al.*[39] The true nuclear-electron potential Z/r and the true wavefunction, ψ_V as distance from the nucleus increases. The pseudopotential V_{pseudo} replaces Z/r and gives rise to the pseudo wavefunction ψ_{pseudo}. At the distance from the nucleus, r_c, the pseudopotential and pseudo wavefunction become equal to the true potential and wavefunction.

the computational cost of plane wave calculations with only minimal compromises in accuracy.

2.2.2. Classical Theory

The rationale for classical energy models is the approximation of true energy surfaces through the use of potentials with simple functional forms and fitted parameters. The simple form of the potentials used are easy to calculate, hence making them well suited to describing large molecular systems. The dependence on parameters greatly limits their generality. This approach avoids directly considering the electronic degrees of freedom of the system. Replacing the complex nature of the electronic component is the underlying cause of the high degree of parameterisation required for classical models. Excellent efficiency savings in computational effort of several orders of magnitude are produced however. A collection of potentials and parameters is known as a force field model. The functional form of the AMBER force field[40], a popular and well known choice used throughout this work, is given by:

$$V(\mathbf{x}) = \sum_{bonds} K_b(b-b_0)^2 + \sum_{angles} K_\theta(\theta-\theta_0)^2 + \sum_{torsions} \left(\frac{V_n}{2}\right)(1+cos[n\phi-\phi_0]) + \sum_{nonbonded} \left(\frac{A_{ij}}{r_{ij}^{12}}\right) - \left(\frac{B_{ij}}{r_{ij}^{6}}\right) + \left(\frac{q_i q_j}{4\pi\epsilon_0 r_{ij}}\right) \tag{2.50}$$

the energy of the system is given as the sum of a series of terms each designed to approximate an actual physical interaction between atoms. Covalent interactions are described by the first term which makes use of potential form known as a harmonic oscillator. Here, the energy of each covalent bond, is a function of the deviation of the actual bond length, b, from some ideal bond length, b_0, multiplied by the force constant, K_b, where b is the same as b_0 the energy of this potential is zero. However, the energy for squashing or stretching the bond increases quadratically. The angle term is likewise a harmonic oscillator dependent on deviation of the angle θ formed between three bonded atoms, from an ideal angle, θ_0 (figure 2.2a). The torsion term is more complex due to the need to represent multiple minima. The torsion angle ϕ is defined between fours of bonded atoms as shown in 2.2b. In parallel with the bond and angle terms V_n acts as the spring constant, and ϕ_0 gives the position of the first peak. Additionally n is multiplicity of the torsion, defining the number of minima within a full 360^{o} turn.

The final components give the energy contributions from all pairings of atoms that are not covalently bonded. r_{ij} is the interatomic separation between atoms i and j. The first two terms containing A and B comprise a Lennard-Jones potential, which has multiple roles. At very short ranges the r^{-12} term dominates giving large energy penalties as atoms approach one another; this mimics Pauli repulsion effects where atoms repel one another as their electron orbitals overlap. The higher power of the r^{-12} means that it drops off to zero at longer ranges more quickly than the r^{-6} term. This gives favourable negative energies at longer ranges to model the effect of attractive van der Waals interactions[41]. The position of the potentials minima and its depth are controlled by a combination of the parameters A and B.

The final term accounts for electrostatic interactions within the system, where q_i and q_j are the charges of the i^{th} and j^{th} atoms respectively and ϵ_0 is the vacuum permittivity. It is typical for all atoms within classical models to be assigned a non-zero partial point charge. This is to account for the partial charges arising due to differences in electronegativities between covalently bonded atoms. Within a molecule, these energy terms are only calculated between atoms that are separated by at least three bonds. Where separated by exactly three bonds these so called 1-4 terms are scaled by an additional factor depending on the force field employed[42,40,43].

The functional form of a force field alone is not sufficient to describe the potential energy surface of a molecular system. These must be combined with sets of parameters that give the correct ideal bond lengths, angles etc. The fact that these must be determined is the great disadvantage of force field models, as creation of parameter sets is a computational and labour intensive task achieved by the fitting of simulation data to match experimental quantities and high-level QM calculations. Many different parameter sets have been developed for different applications. Necessarily the use of fitted parameters limits the generality of different parameter sets which have been

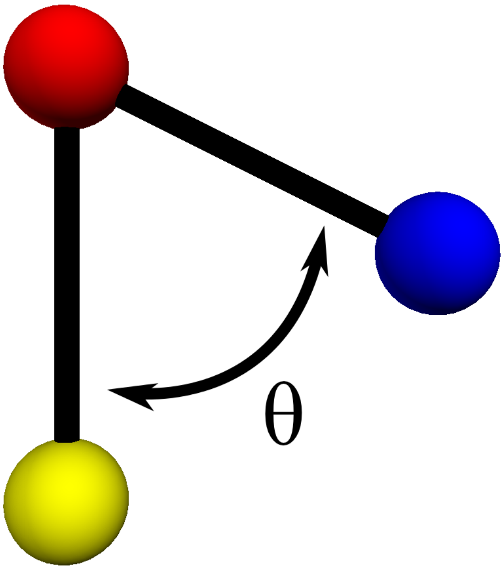

(a) Bond angles are considered between all triplets of bonded atoms. This is considered as the smallest angle formed by the two end atoms with respect to the central atom.

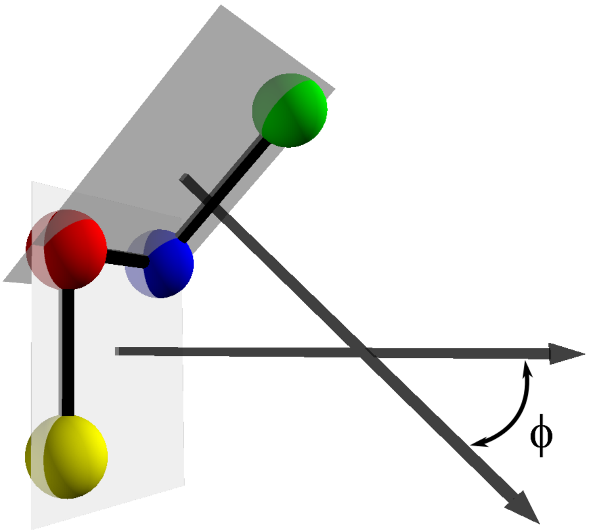

(b) Torsions are formed by groups of four contiguously bonded atoms and represent rotation around bonds. Picturing two planes that are defined by passing through the first and last three atoms in the quartet, the torsion angle is given by the intersection of these planes. This is equivalent to the angle of intersection of the normal vectors of each plane as shown above.

Figure 2.2.

derived under specific chemical and physical conditions. Details of the parameter sets used in this work and their derivation are given below.

Another notable disadvantage of the functional form of (2.50) is that it is too simplistic to account for various physical effects. Covalent bonds are unable to break for instance. As bonds are stretched the energy of their interaction should drop off towards a constant value. However, the form of the harmonic oscillator models the energy as continuously increasing. This prohibits the use of force field models in the study of chemical reactions or enzymes where bonds must necessarily be allowed to break and form. Additionally (2.50) does not include any terms that describe polarisation effects between molecules. Although various extensions have been proposed to develop polarisable force field models, these remain in a tentative and experimental stage[44]; an implementation extending polarisation effects to classical models is given in section 2.2.2.5.

2.2.2.1. AMBER 94 Force Field (FF94)

FF94 was the first force field developed that matched the functional form of equation (2.50). For this reason, an in-depth examination is given here of the process that was used to derive these parameters. The FF94 parameters are based on the preceding force field by Weiner *et al.*[42,45], that was designed primarily for simulation work with implicit solvation models. FF94 was designed with the aim of providing a force field model that would be compatible with explicit solvent simulations using prevailing models of water through a consistent algorithmic approach. The scope of the force field was limited to the fitting of parameters for prevalent biological monomers, namely amino acids and DNA bases[40].

The bond and angle parameters were taken directly from the Weiner *et al.* force field and modified to reproduce the experimental vibrational frequencies of small molecule fragments that make up amino acids and DNA bases. Dihedral parameters were fitted to reproduce the conformational energies of 2nd order Møller-Plesset perturbation theory (a high level QM technique). Lennard-Jones parameters were fitted to reproduce empirical densities and enthalpies of vaporisation in liquid simulations of chemically representative compounds. For instance, parameters for sp^3 hybridised carbon and aliphatic hydrogen were fitted to liquid methane, ethane, propane and butane, whilst aromatic carbon and hydrogen were fitted to liquid benzene. Parameters for other elements - nitrogen, oxygen and sulphur - were taken directly from the OPLS force field. Hydrogen parameters were given further post hoc modifications to account for the particular sensitivity of hydrogen to its chemical environment. Finally partial charges for all atoms were determined using the Restrained Electrostatic Potential (RESP) fitting[46], to multiple conformations, at the Hartree-Fock/6-31G* level. A scaling factor of $\frac{1.0}{1.2}$ for 1-4 interactions is used for this force field and its derivatives.

2.2.2.2. AMBER 99SB Force Field (FF99SB)

The development of the AMBER 99 force field was derived directly from FF94, and focussed on the improvement of torsional terms through the inclusion of higher order Fourier terms[43]. Considerable improvements versus FF94 were observed in recreating conformational energies at the GVB/LMP2 level. Further improvements to protein backbone torsion angles were made and released as FF99SB[47]. FF99SB is used extensively throughout this work to model an Adenosine-Thymidine DNA base pair and full protein-ligand systems.

2.2.2.3. General Amber Force Field (GAFF)

The aim of GAFF[48] was to provide a generic set of force field parameters suitable for a wide variety of small molecules relevant to drug design. This provides a particular challenge, as it is necessary to derive parameters that cover a sizeable portion of chemical space. Combined however with effective algorithms for the determination of partial charges, such as RESP and AM1-bcc, GAFF greatly improved the ease of applying force field models to novel compounds.

GAFF parameters were designed to cover small organic compounds comprised of carbon, oxygen, nitrogen, phosphorus, hydrogen, sulphur and halogens. The chemical range of all of these compounds was covered by 35 unique atom types. In principle GAFF can be applied out of the box to any compound comprised of the above elements.

The majority of bond and angle parameters were derived as averaged values taken from experimental X-ray and neutron diffraction data[49,50]. Additional parameters where also derived based on MP2/6-31G* calculations. Owing to the large number of atom types, the basic GAFF parameter set is not complete, but included are empirical rules for extrapolating missing parameters from those provided. Torsional parameters were derived by fitting potentials to recreate torsional scans at the MP4/6-311G(d,p) level for a set of 200 model molecules. Lennard-Jones parameters were adapted directly from the FF99 force field.

2.2.2.4. Water Models

As the ubiquitous biological solvent, considerable effort has been invested in developing an appropriate classical model for water. This has proved challenging however, due to the unusual properties of water and the diversity of roles it fulfil[20]. Very many different potentials have been proposed[51]. We shall here discuss only the two employed in this work. These are the TIP3P and TIP4P models[52].

A diagrammatic representation of the two models is given in figure 2.3 along with the

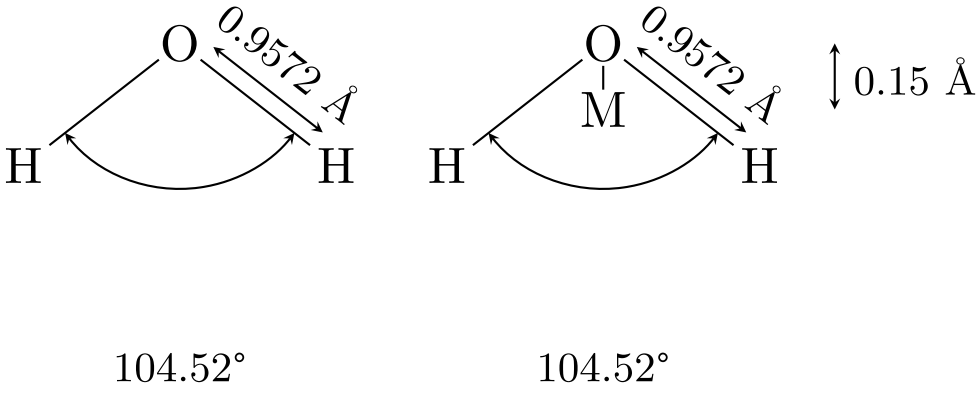

(a) TIP3P

(b) TIP4P

Atom	TIP3P	TIP4P
O	-0.834	0.00
H	0.417	0.52
M	N/A	-1.04

(c) Atomic charges for each model

Figure 2.3.: TIP3P and TIP4P water models.

charges associated with the relevant atoms. The internal geometry of each water is kept rigid, allowing only three rotational and three translational degrees of freedom per molecule. The key difference between the models is the inclusion of a fourth dummy atom on the TIP4P model. This dummy atom is assigned a partial charge. However, it does not have any Lennard-Jones interactions. Inclusion of the additional site greatly improves the distribution of charge and gives more accurate densities[53]. Full parameters for both models are given by Jorgensen *et al.*[52]. Although the TIP4P model offers benefits in terms of accuracy, the inclusion of an additional point charge increases the computational cost.

2.2.2.5. Drude Oscillators

Multiple extensions to equation (2.50) have been proposed to account for polarisation effects within classical systems. Here we shall consider the use of Drude oscillators[54]. Firstly each atom, i, in the model system is assigned a point charge n_i and a mobile Drude particle. The Drude particle carries a point charge d_i that does not interact electrostatically with n_i, instead each Drude particle is anchored to its corresponding nucleus by a harmonic potential with a sufficiently large spring constant, k_i^d, to prevent significant displacement. The charge d_i is able to interact electrostatically with all other charges within the system, except n_i. We also define the net charge on each atom $q_i = n_i + d_i$. Additionally, The energetic contribution of the Drude particles is given by:

$$U^{Drude}(\mathbf{r}^d;\mathbf{x}) = \sum_{i=1} k_i^d |\mathbf{r}_i^d|^2 + \sum_{i=1}\sum_{j>i} \frac{k_e d_i d_j}{r_{ij}} + \sum_{I=1}\sum_{j\neq I} \frac{k_e n_I d_j}{r_{Ij}} + \sum_{i=1}\sum_{j>i} \frac{k_e n_I n_J}{r_{IJ}} \tag{2.51}$$

where k_e is Coulomb's constant, $\mathbf{r}_i^d$ denotes the displacement vector between the i^{th} Drude particle and its paired nucleus. Values of r with multiple subscripts denote Cartesian distances between system components, capital letters being nuclei and lower case letters Drude particles. All of the distances can be written in terms of $\mathbf{r}^d$ and $\mathbf{x}$, the nuclear configuration of the system, allowing U^{Drude} to be expressed as a function of these. The system is treated using the Born-Oppenheimer approximation, with the Drude particles considered as electronic degrees of freedom that relax instantaneously to their minimum for any given configuration of nuclei. This is denoted by the only parametric dependence of equation (2.51) on $\mathbf{x}$. Thus we write the polar contribution to the potential energy:

$$U^{polar}(\mathbf{x}) = \min\{U^{Drude}(\mathbf{x})\} \tag{2.52}$$

this represents the miminum energy configuration of the Drude degrees of freedom, $\mathbf{r}^d$, for a given configuration of nuclei. For each Drude particle, we have introduced two free parameters, d_i and k^d that control the polarisability of each atom. Smaller values of the spring constant increase the polarisability of an atom by allowing greater displacement of the drude particle from the nucleus. On the other hand, increasing the magnitude of d_i, typically by making it more negative to simulate a polarised electron density, will also increase the polarisability. Both of these parameters can be related to a single term, the isotropic polarisability, α[55]:

$$\alpha_i = \frac{d_i^2}{k_i^d} \tag{2.53}$$

typically values of k_i^d are chosen so as to minimize potential numerical problems that can arise due to large displacements of Drude particles from their respective nuclei. Specific values of α_i can then be achieved by fitting the appropriate value of d_i.

2.2.3. Ewald Summation

The treatment of electrostatics in molecular simulation is a well established problem, especially compounded by the common use of periodic boundary conditions. The $\frac{1}{r}$ form of the electrostatic potential leads it to form very long-ranged interactions. This potentially allows interactions between charges to span periodic images and requires that such potentials are modelled appropriately. Ewald summation[56] provides an efficient and frequently used approach to calculate electrostatic energies in infinite periodic systems. We will here deal with only a very simple qualitative description. Starting by slightly formalising the contribution of the Coulomb term from equation (2.50):

$$U^{Coul} = \sum_{i} \sum_{j>i} \frac{k_e q_i q_j}{|\mathbf{r}_{ij}|} \tag{2.54}$$

where $\mathbf{r}_{ij}$ is the displacement vector between the i^{th} and j^{th} atom. We now expand this to account for interactions arising through periodic boundary conditions:

$$U^{Coul} = \sum_{\mathbf{n}}^{\infty} \sum_{i} \sum_{j>i} \frac{k_e q_i q_j}{(\mathbf{r}_{ij} + \mathbf{n})} \tag{2.55}$$

where $\mathbf{n}$ is the displacement vector between the central image and the periodic image under consideration and we are summing over all images. It can be shown that this series is only conditionally convergent, i.e. the order in which the terms are considered changes the answer. Naively attempting to solve equation (2.55) is therefore a very difficult process and has lead to the development of alternative approaches such as the Ewald summation and its extension Particle Mesh Ewald (PME)[57]. The basic insight is that equation (2.55) can be separated into short range terms that converge quickly, and longer range terms that converge slowly, through the introduction of Gaussian charge distributions that screen the interactions of the charges. The final result for the Ewald energy, U^{Ewald}:

$$\begin{aligned} U^{Ewald} =& \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \left(\sum_{\mathbf{n}}^{\infty} q_i q_j \frac{\operatorname{erfc}[\kappa(\mathbf{r}_{ij} + \mathbf{n})]}{\mathbf{r}_{ij} + \mathbf{n}} + \right. \qquad (2.56) \\ & \left. \frac{1}{\pi L^3} \sum_{\mathbf{k} \neq 0} \frac{4\pi^2 q_i q_j}{|\mathbf{k}|^2} \exp\left[\frac{-|\mathbf{k}|^2}{4\kappa^2} \right] \cos[\mathbf{k} \cdot \mathbf{r}_{ij}] \right) - \\ & \frac{\kappa}{\pi^{\frac{1}{2}}} \sum_{i=1}^{N} q_i^2 + \frac{2\pi}{3L^3} \left| \sum_{i=1}^{N} q_i \mathbf{r}_i \right|^2 \qquad (2.57) \end{aligned}$$

where we assume a cubic system of length L and carry out a summation over the set of reciprocal lattice vectors $\mathbf{k} = \frac{2\pi\mathbf{n}}{L^2}$. κ is a free parameter representing the width of the Gaussian distribution used in screening the charges. The erfc is the complentary error function, that has a steep transition between 0 and 1. Thus by choosing an appropriately large value of κ the first sum can be made such that only $\mathbf{n} = \mathbf{0}$, i.e. the sum over charges in the central image, contributes significantly and the rest can be truncated. The larger the value of κ used the greater the number of terms are needed in the second sum over the reciprocal lattice vectors, but this provides a straight forward convergent series.

Particle Mesh Ewald extensions are based around the introduction of the fast Fourier transformation (FFT) to handle the sums in reciprocal space. As the FFT is a discrete operation it requires the potentials under consideration to be evaluated at points on a grid. The use of the FFT provides a computational cost that scales as $\mathcal{O}(N \log N)$ compared with $\mathcal{O}(N^2)$ for conventional Ewald summation.

2.3. Sampling Methods

Although the theory of statistical mechanics provides powerful tools for the treatment of thermodynamic systems, large practical difficulties remain. Typical molecular systems have a very high (3N) dimensionality but only small, intricately shaped, regions of the phase space contribute meaningfully to ensemble average values. This problem precludes the use of direct integration techniques to derive exact solutions for systems of interest. Instead computer simulation techniques have arisen as an alternative approach to collect statistically meaningful information. These techniques are employed to generate a collection of configurations that have the correct Boltzmann distribution of states and can be used in the evaluation of ensemble average properties.

2.3.1. Molecular Dynamics (MD)

MD makes use of Newton's Second Law of Motion[22]:

$$\mathbf{a}_i(t) = \frac{\mathbf{F}_i(t)}{m_i} \tag{2.58}$$

for a given configuration of a molecular system, at time t, the forces acting upon the i^{th} atom, $\mathbf{F}_i$ are accessible, and can be related to the associated set of accelerations, $\mathbf{a}_i$, through the mass of each individual atom m_i. The formal definition of the forces acting within an isolated system is given by:

$$\mathbf{F}_i(t) = -\frac{dU(\mathbf{x}(t))}{d\mathbf{x}_i} \tag{2.59}$$

such that the forces, $\mathbf{F}_i(t)$, are given by the negative gradient of the potential energy surface of the atomic coordinates at time t. This provides sufficient information to advance the system through time taking into account the current velocities and the accelerations to give a new set of atomic velocities and positions. This can be written as the second order differential equation problem:

$$\frac{d^2\mathbf{x}_i(t)}{dt^2} = -\frac{dU(\mathbf{x}(t))}{d\mathbf{x}_i m_i} \tag{2.60}$$

solving an equation of this form requires the boundary conditions $\mathbf{x}_i(0)$ i.e. the starting configuration of the system and $\frac{d\mathbf{x}_i(0)}{dt} = \mathbf{v}_i(0)$, i.e. the starting velocities. Ideally a closed form solution of equation (2.60) would be possible. However, we are faced with a notorious many body problem and are forced to fall back on numerical integration techniques. This is a very broad field and will not be covered here. However a summary of some notable techniques including those used throughout this work are included in appendix B.

At present we have a formulation of MD that conserves energy, corresponding to the NVE ensemble. Energy exchanges back and forth between the potential and kinetic components of the system but the total energy remains constant. In order to model different ensembles we need to extend the formulation to produce a Boltzmann distribution of states. This is achieved through the use of thermostats, algorithms that modify the velocities or forces to enforce a desired distribution. There are various forms of thermostat with different advantages and use cases. We shall here consider only the Langevin thermostat[58]. This is implemented through modifying the force term at every time step:

$$\mathbf{F}_i^L(t) = \mathbf{F}_i(t) - \gamma m_i \mathbf{v}_i(t) + (2\gamma \frac{1}{\beta} m_i)^{\frac{1}{2}} \mathbf{R}(t) \tag{2.61}$$

here we define the Langevin forces, $\mathbf{F}_i^L$, in terms of the deterministic forces (equation (2.59)), a friction term and a 'kick' term. $\mathbf{R}(t)$ is a Gaussian process, in practice implemented as a pseudo-random number generator choosing values from a Gaussian distribution of mean 0 and variance 1. The constant γ is a free parameter, known as the friction constant. The contrasting roles of the friction and kick terms act together to give the overall correct Boltzmann distribution, the friction term removing energy from the system whilst the kick term adds energy. Each particle is thermostated individually, making this form of thermostat unsuitable for measuring certain dynamic properties such as diffusion rates. The friction constant controls the behaviour of the thermostat by modifying the average kick experienced by each atom and the strength of the friction term.

2.3.2. Monte Carlo (MC)

MC methods have found widespread use across many fields, defined by the use of stochastic sampling approaches to evaluate system properties. The application of MC in the calculation of numerical integrals is common. However, we will here consider the application of Markov chain MC to the exploration of the configuration space of molecular systems. A Markov chain is formally defined as a system that undergoes stochastic transitions between states[19]. Such a system is said to be memoryless, i.e. the transition probabilities depend only on the current state of the system and not its previous states.

The hallmark of Markov chain MC is stochastic sampling of a system's configuration state space. As noted in section 2.1 the state space for molecular systems is comprised of 3N dimensions with only small regions contributing significantly to molecular integrals. Stochastically sampling such a phase space naïvely therefore is very inefficient and the chance of randomly generating a point in the phase space of the system that is relevant to the equilibrium distribution is vanishly small. This problem is resolved through the use of importance sampling[59]; this generates new configurations using an alternative distribution and an acceptance test to recover the desired Boltzmann distribution. States are added to the Markov chain with the appropriate probability and hence are weighted equally.

A standard process for the construction of a Markov chain with importance sampling is given by the Metropolis-Hastings algorithm[60]:

1. From some starting configuration $\mathbf{x}_1$, generate a trial configuration $\mathbf{x}_2$ by small stochastic displacement of $\mathbf{x}_1$.
2. Calculate the acceptance probability, P^{acpt}, from $\mathbf{x}_1$ and $\mathbf{x}_2$.
3. Draw a random number, r, from a uniform distribution in the range 0 to 1.
4. If r is less than P^{acpt}, add $\mathbf{x}_2$ to the Markov chain, $\mathbf{x}_2$ becomes $\mathbf{x}_1$.
5. If r is greater than P^{acpt}, add another copy of $\mathbf{x}_1$ to the Markov chain.
6. Repeat from 1. as desired.

As stated above, a Markov chain requires that the chance of accepting a new configuration should depend only on the current state and the proposed state. This is reflected in the acceptance probability, known as the Metropolis test, given by:

$$P^{acpt}(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \min\left(1, \frac{P(\mathbf{x}_2)P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P(\mathbf{x}_1)P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)}\right) \tag{2.62}$$

as defined previously, P denotes the Boltzmann factor, such that:

$$\frac{P(\mathbf{x}_2)}{P(\mathbf{x}_1)} = \exp\left[-\beta(U(\mathbf{x}_2) - U(\mathbf{x}_1))\right] \tag{2.63}$$

additionally, $P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)$ represents the probability of generating configuration $\mathbf{x}_2$ from $\mathbf{x}_1$. The value of this term is dependent on the procedure used in generating trial configurations. In the case outlined above the P_g terms will cancel so long as the displacements are drawn from a distribution symmetrical about zero. The difference between $\mathbf{x}_1$ and $\mathbf{x}_2$ can be arbitrarily large. However, using stochastic displacement, the chance of generating a viable configuration is closely tied to the size of the move. In practice, the size of the average displacement is tuned such that a reasonable acceptance rate is achieved. This leaves the acceptance test for the Metropolis-Hastings algorithm as:

$$P^{acpt}(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \min\left(1, \exp\left[-\beta(U(\mathbf{x}_2) - U(\mathbf{x}_1))\right]\right) \tag{2.64}$$

the appropriate acceptance tests for more sophisticated structure generation procedures and their application to hybrid free energy calculations are considered in section 2.5. The acceptance test is chosen such that the structures accepted into the Markov chain follow the desired Boltzmann distribution. How this works can be demonstrated through the concept of a limiting (stationary) distribution. We start by defining the transition operator, $\hat{\mathbf{T}}$, that acts on probability distributions[19]:

$$\hat{\mathbf{T}} p_i(\mathbf{x}) = p_{i+1}(\mathbf{x}) \tag{2.65}$$

where $p_i(\mathbf{x})$ is some arbitrary distribution over configuration space. The limiting distribution of $\hat{\mathbf{T}}$ is defined as:

$$p(\mathbf{x})_{lim} = \lim_{i \to \infty} p_i(\mathbf{x}) \tag{2.66}$$

thus meeting the criteria:

$$\hat{\mathbf{T}} p_{lim}(\mathbf{x}) = p_{lim}(\mathbf{x}) \tag{2.67}$$

rationalising the alternative name of stationary distribution as it is unaltered by the application of $\hat{\mathbf{T}}$. As developed previously p_{lim} is an eigenfunction of $\hat{\mathbf{T}}$ with an eigenvalue of 1. In practice we do not deal directly with probability distributions but instead with a series of configurations. However, the process is analogous. By successive applications of a transition operator to each configuration, the limiting distribution is approached for the states within the Markov chain. We wish to define $\hat{\mathbf{T}}$ therefore such that the limiting distribution corresponds to the equilibrium Boltzmann distribution. This is given by the condition[19]:

$$\sum_j P(\mathbf{x}_i) T(\mathbf{x}_i \to \mathbf{x}_j) = \sum_j P(\mathbf{x}_j) T(\mathbf{x}_j \to \mathbf{x}_i) \tag{2.68}$$

where $T(\mathbf{x}_1 \to \mathbf{x}_2)$ is the probability of the Markov chain transitioning from $\mathbf{x}_1$ to $\mathbf{x}_2$. In practice a more restrictive condition is used, known as detailed balance. Here equation (2.68) is satisfied by ensuring the condition is met for every j:

$$P(\mathbf{x}_1) T(\mathbf{x}_1 \to \mathbf{x}_2) = P(\mathbf{x}_2) T(\mathbf{x}_2 \to \mathbf{x}_1) \tag{2.69}$$

the transition probability can be decomposed into the product of the probability of generating a configuration and the probability of accepting a configuration:

$$P(\mathbf{x}_1) P_g(\mathbf{x}_1 \to \mathbf{x}_2) P^{acpt}(\mathbf{x}_1 \to \mathbf{x}_2) = P(\mathbf{x}_2) P_g(\mathbf{x}_2 \to \mathbf{x}_1) P^{acpt}(\mathbf{x}_2 \to \mathbf{x}_1) \tag{2.70}$$

for here, it can be shown that the Metropolis acceptance test meets the detailed balance criterion:

$$P(\mathbf{x}_1) P_g(\mathbf{x}_1 \to \mathbf{x}_2) \min \left[1, \frac{P(\mathbf{x}_2) P_g(\mathbf{x}_2 \to \mathbf{x}_1)}{P(\mathbf{x}_1) P_g(\mathbf{x}_1 \to \mathbf{x}_2)} \right]$$

$$= P(\mathbf{x}_2)P_g(\mathbf{x}_2 \to \mathbf{x}_1) \min\left[1, \frac{P(\mathbf{x}_1)P_g(\mathbf{x}_1 \to \mathbf{x}_2)}{P(\mathbf{x}_2)P_g(\mathbf{x}_2 \to \mathbf{x}_1)}\right] \tag{2.71}$$

$$\begin{aligned} &\min[P(\mathbf{x}_1)P_g(\mathbf{x}_1 \to \mathbf{x}_2), P(\mathbf{x}_2)P_g(\mathbf{x}_2 \to \mathbf{x}_1)] \\ &\qquad = \min\left[P(\mathbf{x}_2)P_g(\mathbf{x}_2 \to \mathbf{x}_1), P(\mathbf{x}_1)P_g(\mathbf{x}_1 \to \mathbf{x}_2)\right] \end{aligned} \tag{2.72}$$

where we have multiplied into the brackets and cancelled the resulting fractions.

2.4. Free Energy Difference Estimation

Whilst it is in principle possible to find the free energy of the thermodynamic system through direct numerical solution of the partition function, for biologically relevant molecular systems this is too complex to attempt. Instead free energy calculations focus on calculating the free energy difference between two thermodynamic states without directly calculating the partition function of either end point. There are many different approaches to this form of calculation. However in this work we shall make use primarily of a number of equilibrium techniques collectively referred to as free energy perturbation (FEP). These constitute three different free energy difference estimators, each of which shall be derived and motivated in the following section. The underlying idea behind all these approaches is that a free energy difference can be expressed in terms of one or more ensemble averages[1].

When calculating a free energy difference between two states it is convenient to define function that describes a smooth transition between the end-points. The most frequently used form is given by:

$$U_\lambda = \lambda U_1 + (1 - \lambda) U_0 \tag{2.73}$$

where U_0 and U_1 are the potential energy function of the end-points, denoted as states 0 and 1 respectively, and λ is a scaling factor taking a value between zero and one. The functional dependence of U on $\mathbf{x}$ has been skipped here for clarity. This is referred to as the λ-coupling approach and will be used, at least implicitly, throughout this section.

2.4.1. The Zwanzig Equation

This is also known as exponential averaging, or the perturbative method. This is the first derived of the estimators we shall consider being given by Zwanzig in 1954[2]. The derivation is given as follows:

$$\begin{aligned}
\Delta A_{0\rightarrow 1} &= A_1 - A_0 && (2.74)\\
&= -\frac{1}{\beta} \ln \frac{Z_1}{Z_0} && (2.75)\\
&= -\frac{1}{\beta} \ln \frac{\int \exp\left[-\beta U_1\right] \mathrm{d}\mathbf{r}}{Z_0} && (2.76)\\
&= -\frac{1}{\beta} \ln \frac{\int \exp\left[-\beta U_1\right] \exp\left[-\beta (U_0 - U_0)\right] \mathrm{d}\mathbf{r}}{Z_0} && (2.77)\\
&= -\frac{1}{\beta} \ln \frac{\int \exp\left[-\beta U_0\right] \exp\left[-\beta (U_1 - U_0)\right] \mathrm{d}\mathbf{r}}{Z_0} && (2.78)\\
&= -\frac{1}{\beta} \ln \left\langle \exp\left[-\beta (U_1 - U_0)\right] \right\rangle_0 && (2.79)
\end{aligned}$$

$$= -\frac{1}{\beta} \ln \langle \exp\left[-\beta \Delta U\right] \rangle_0 \tag{2.80}$$

where $\Delta U = U_1 - U_0$ and the ensemble average is over state 0. The formulation of the Zwanzig equation provides the convenient property that it is based on a single ensemble average, hence only requiring MD or MC sampling of one state. This fact is used profitably to calculate MM to QM free energy differences efficiently, as will be discussed in section 2.5. However, the Zwanzig equation has some notable drawbacks. The exponential term within the ensemble average leads to poor numerical behaviour as snapshots are weighted unevenly. Thus individual configurations can have a disproportionately strong influence on the ensemble average, making it difficult to assess whether a calculation has converged[23]. There is no way to rule out that there is not some, as yet unsampled, configuration that will contribute strongly to the free energy difference. This effect lowers the efficiency of the operator requiring the end-states of the calculation to be more similar to one another than other estimators. This increases the number of intermediate windows that must be used for a perturbation and leads to the Zwanzig equation not being as commonly used in free energy calculations[19].

2.4.2. Thermodynamic Integration (TI)

Also known as the integration method. This estimator, given by Kirkwood[4], is derived explicitly using the λ coupling approach:

$$\begin{aligned}
\Delta A_{0\to 1} &= A_1 - A_0 && (2.81)\\
&= \int_0^1 \frac{\partial A_\lambda}{\partial \lambda} \mathrm{d}\lambda && (2.82)\\
&= \int_0^1 \frac{\partial[-\frac{1}{\beta}\ln Z_\lambda]}{\partial \lambda} \mathrm{d}\lambda && (2.83)\\
&= -\frac{1}{\beta}\int_0^1 \frac{\partial[\ln Z_\lambda]}{\partial Z_\lambda}\frac{\partial Z_\lambda}{\partial \lambda} \mathrm{d}\lambda && (2.84)\\
&= -\frac{1}{\beta}\int_0^1 \frac{1}{Z_\lambda}\frac{\partial Z_\lambda}{\partial \lambda} \mathrm{d}\lambda && (2.85)\\
&= -\frac{1}{\beta}\int_0^1 \frac{1}{Z_\lambda}\frac{\partial[\int \exp\left[-\beta U_\lambda\right] \mathrm{d}\mathbf{r}]}{\partial \lambda} \mathrm{d}\lambda && (2.86)\\
&= -\frac{1}{\beta}\int_0^1 \frac{1}{Z_\lambda}\int \frac{\partial[\exp\left[-\beta U_\lambda\right]]}{\partial \lambda} \mathrm{d}\mathbf{r}\mathrm{d}\lambda && (2.87)\\
&= -\frac{1}{\beta}\int_0^1 \frac{1}{Z_\lambda}\int \frac{\partial[\exp\left[-\beta U_\lambda\right]]}{\partial[-\beta U_\lambda]}\frac{\partial[-\beta U_\lambda]}{\partial \lambda} \mathrm{d}\mathbf{r}\mathrm{d}\lambda && (2.88)\\
&= \int_0^1 \frac{1}{Z_\lambda}\int \exp\left[-\beta U_\lambda\right]\frac{\partial[U_\lambda]}{\partial \lambda} \mathrm{d}\mathbf{r}\mathrm{d}\lambda && (2.89)\\
&= \int_0^1 \int \frac{\exp\left[-\beta U_\lambda\right]\frac{\partial[U_\lambda]}{\partial \lambda}}{Z_\lambda} \mathrm{d}\mathbf{r}\mathrm{d}\lambda && (2.90)
\end{aligned}$$

$$= \int_0^1 \left\langle \frac{\partial U_\lambda}{\partial \lambda} \right\rangle_\lambda \mathrm{d}\lambda \tag{2.91}$$

$$= \int_0^1 \left\langle \frac{\partial[\lambda U_1 + (1-\lambda)U_0]}{\partial \lambda} \right\rangle_\lambda \mathrm{d}\lambda \tag{2.92}$$

$$= \int_0^1 \langle U_1 - U_0 \rangle_\lambda \, \mathrm{d}\lambda \tag{2.93}$$

the resulting integral can be evaluated using numerical quadrature approaches as outlined in appendix A. In contrast to the Zwanzig equation this expression requires the evaluation of at least two ensemble averages for solution of the integral. Typically however, the integral is further subdivided at multiple intermediate values of λ. The great benefit of TI is that it avoids the exponential weighting of the Zwanzig equation, making it considerably less unstable numerically and giving smoother convergence behaviour[1]. TI also has the notable advantage that the ensemble average at each λ-window is dependent only on that value of λ. This allows one retrospectively to add additional λ points to a calculation without the need to re-evaluate averages at previously completed λ-windows.

2.4.2.1. Finite-Difference Thermodynamic Integration (FDTI)

For practical reasons the final steps of the previous derivation i.e.:

$$\left\langle \frac{\partial U_\lambda}{\partial \lambda} \right\rangle_\lambda = \langle U_1 - U_0 \rangle_\lambda \tag{2.94}$$

can be difficult to implement. This lead to the development of FDTI as a workaround for this problem and an alternative to regular TI[61]. Starting as previously:

$$\Delta A_{0\to 1} = A_1 - A_0 \tag{2.95}$$

$$= \int_0^1 \frac{\partial A_\lambda}{\partial \lambda} \mathrm{d}\lambda \tag{2.96}$$

we then note the following limit:

$$\frac{\partial A_\lambda}{\partial \lambda} = \lim_{h \to 0} \frac{A_{\lambda+h} - A_{\lambda-h}}{2h} \tag{2.97}$$

and introduce its finite difference approximation:

$$\lim_{h \to 0} \frac{A_{\lambda+h} - A_{\lambda-h}}{2h} \approx \frac{A_{\lambda+h} - A_{\lambda-h}}{2h} \tag{2.98}$$

$$\approx \frac{A_{\lambda\to\lambda+h} - A_{\lambda\to\lambda-h}}{2h} \tag{2.99}$$

the finite difference approximates the limit well so long as h is sufficiently small. The

Zwanzig equation is then used to evaluate the free energy differences in this term giving:

$$\Delta A_{0\to1} \approx \int_0^1 \frac{\langle \exp\left[-\beta(U_{\lambda+h} - U_\lambda)\right]\rangle_\lambda - \langle \exp\left[-\beta(U_{\lambda-h} - U_\lambda)\right]\rangle_\lambda}{2h} \mathrm{d}\lambda \tag{2.100}$$

The two ensemble averages in this expression are for the same value of λ and hence can be evaluated with a single simulation. Although this formulation makes use of the Zwanzig equation the small values of h used ensure excellent overlap of states and rapid convergence of the perturbations. The use of FDTI has been reported to offer superior convergence characteristics to TI[61].

2.4.2.2. Replica-Exchange Thermodynamic Integration (RETI)

A further augmentation suggested for TI is the use of Hamiltonian Replica Exchange (HRE)[62]. In practice this method is compatible with any of the estimators presented in this section, however, it is most commonly used in the form of RETI and hence is considered here. HRE takes the form of a specialised MC move undertaken at regular intervals through a TI calculation. Such a calculation consists of a number of simultaneous MC or MD simulations carried out at a range of λ values. A HRE move attempts to swap configurations between simulations with neighbouring λ values, subject to the acceptance test[63]:

$$\begin{aligned} &P_A^{RETI}(\mathbf{x}_{\lambda m}, \mathbf{x}_{\lambda n}) = \\ &\qquad \min(1, \exp\left[-\beta((U_{\lambda n}(\mathbf{x}_{\lambda m}) - U_{\lambda n}(\mathbf{x}_{\lambda n})) - (U_{\lambda m}(\mathbf{x}_{\lambda m}) - U_{\lambda m}(\mathbf{x}_{\lambda n})))\right]) \end{aligned} \tag{2.101}$$

where λm and λn represent neighbouring λ windows within a calculation. Correspondingly $\mathbf{x}_{\lambda m}$ and $\mathbf{x}_{\lambda n}$ are configurations generated by sampling at each λ window. If the move is accepted the configurations are swapped between λ values, $\mathbf{x}_{\lambda m}$ becoming $\mathbf{x}_{\lambda n}$ and vice versa.

This process is designed to enhance sampling at each λ window by providing access to a larger volume of configuration space provided by sampling under similar, but distinct, Hamiltonians. The real advantage of this process is that it comes at almost no additional computational cost, the only restriction being that computations must be carried out concurrently.

2.4.3. Bennett's Acceptance Ratio (BAR)

The derivation of this estimator is somewhat a more complex than either of the two proceeding and is not given here in full, but is developed along the following lines[3]:

$$\begin{aligned}
\Delta A_{0\to 1} &= -\frac{1}{\beta}\ln\frac{Z_1}{Z_0} && (2.102)\\
&= -\frac{1}{\beta}\ln\frac{Z_1 g(\mathbf{x})\exp\left[-\beta(U_0+U_1)\right]}{Z_0 g(\mathbf{x})\exp\left[-\beta(U_0+U_1)\right]} && (2.103)\\
&= -\frac{1}{\beta}\ln\frac{\langle g(\mathbf{x})\exp\left[-\beta U_1\right]\rangle_0}{\langle g(\mathbf{x})\exp\left[-\beta U_0\right]\rangle_1} && (2.104)
\end{aligned}$$

here $g(\mathbf{x})$ is an arbitrary weighting function, that modifies the potential energy surface. As it appears in both the numerator and the denominator, the value of the weighting function does not influence the free energy difference, but its value will influence the statistical properties of equation (2.104). As $g(\mathbf{x})$ can be freely chosen, the best choice is the value of the function that minimises the statistical variance. The details of this are not given here, but it can be shown that the optimal value is given by:

$$g(\mathbf{x}) = \left(\frac{Z_0}{n_0}\exp\left[-\beta U_1\right] + \frac{Z_1}{n_1}\exp\left[-\beta U_0\right]\right)^{-1} \tag{2.105}$$

substituting equation (2.105) into equation (2.104) gives the final result:

$$\begin{aligned}
\Delta A_{0\to 1} &= \frac{1}{\beta}\ln\frac{\langle F(U_0-U_1+C)\rangle_1}{\langle F(U_1-U_0-C)\rangle_0} + C - \frac{1}{\beta}\ln\frac{n_1}{n_0} && (2.106)\\
C &= \Delta A_{0\to 1} + \frac{1}{\beta}\ln\frac{n_1}{n_0} && (2.107)
\end{aligned}$$

here $F(x)$ is the Fermi function $F(x) = (1+\exp[\beta x])^{-1}$, whilst n_0 and n_1 are the number of snapshots sampled from each state. From an arbitrary initial guess equations (2.106) and (2.107) are solved self-consistently until a converged estimate of $\Delta A_{0\to 1}$ is obtained. Here C is a constant that adjusts the relative heights of the potential energy surfaces, U_0 and U_1. Self-consistent solution gives the near-optimal value of C that minimises the variance of $\Delta A_{0\to 1}$. At this value of C, denoted here as C^{opt}, the condition $\left\langle f(U_0-U_1+C^{opt})\right\rangle_1 = \left\langle f(U_1-U_0-C^{opt})\right\rangle_0$ is met. BAR is often considered the most efficient estimator of those presented, as its derivation is performed with the minimisation of the variance of calculated free energy differences in mind.

2.4.4. Free Energy Cycles

Whilst in principle the methods outlined above allow the calculation of free energy differences between any two thermodynamic states, this does not mean that such calculations are inherently practical to attempt. The greater the difference between the end points of the calculation the greater the number of λ windows will be required to evaluate it. Long simulation times at each λ are also required to allow the convergence of ensemble averages as large scale conformational differences may be observed between windows.

The direct calculation of free energies of binding or solvation can be circumvented through the use of thermodynamic cycles as shown in figures 2.4a and 2.4b[19]. Attempting to evaluate (in the case of hydration) either of the terms ΔA_P^{solv} or ΔA_Q^{solv} can be computationally intensive. As the ligand is inserted into the water, large scale changes in the structure of the solvent to form hydration shells can occur on relatively long timescales. Calculation of relative free energy differences (again in the case of hydration defined as $\Delta\Delta A_{P\to Q}^{solv} = \Delta A_Q^{solv} - \Delta A_P^{solv}$) can be achieved much more efficiently, through use of the relation:

$$\Delta\Delta A_{P\to Q}^{solv} = \Delta A_Q^{solv} - \Delta A_P^{solv} = \Delta A_{P\to Q}^{(aq)} - \Delta A_{P\to Q}^{(g)} \tag{2.108}$$

calculation of the terms $\Delta A_{P\to Q}^{(aq)}$ and $\Delta A_{P\to Q}^{(g)}$ is comparatively easy due to the inherent state space overlap between thermodynamic states in the same phase. Assuming that P and Q are similar enough to one another no large shifts within the solvent structure or protein binding site are likely to occur providing ideal conditions for free energy calculations to converge quickly. Other practical advantages are offered through the cancellation of systematic errors in both computational free energy calculations and experimental determination of free energies. Whilst relative free energies cannot be used to assess the absolute potency of a potential drug molecule they are valuable in optimisation procedures. Starting from a lead molecule of demonstrated efficacy, systematic modifications can be made to explore the nearby chemical space for increases in binding affinity using relative free energy calculations[1].

2.4.5. Single vs Dual Topology Calculations

Calculation of terms such as $\Delta A_{P\to Q}^{(g)}$, $\Delta A_{P\to Q}^{(aq)}$ and $\Delta A_{P\to Q}^{(b)}$ relies on establishing a non-physical pathway between different molecules, often referred to as an alchemical perturbation. This can pose practical challenges as the choice of path is arbitrary and, to provide optimal convergence, should be chosen to minimise fluctuations in the free energy as a function of λ.

Use of the single topology approach makes use of λ-coupling as given by equation (2.73) and requires the simulation of a single ligand molecule. Atoms shared between the

$$\begin{tikzcd}
\mathrm{P}_{(g)} \arrow[r, "\Delta A^{(g)}_{P\to Q}"] \arrow[d, "\Delta A_P^{solv}"'] & \mathrm{Q}_{(g)} \arrow[d, "\Delta A_Q^{solv}"] \\
\mathrm{P}_{(aq)} \arrow[r, "\Delta A^{(aq)}_{P\to Q}"'] & \mathrm{Q}_{(aq)}
\end{tikzcd}$$

(a) Relative free energy of hydration cycle

$$\begin{tikzcd}
\mathrm{P}_{(aq)} \arrow[r, "\Delta A^{(aq)}_{P\to Q}"] \arrow[d, "\Delta A_P^{bind}"'] & \mathrm{Q}_{(aq)} \arrow[d, "\Delta A_Q^{bind}"] \\
\mathrm{P}_{(b)} \arrow[r, "\Delta A^{(b)}_{P\to Q}"'] & \mathrm{Q}_{(b)}
\end{tikzcd}$$

(b) Relative free energy of binding cycle

Figure 2.4.: Thermodynamic cycles for the calculation of relative free energy differences of the two ligands P and Q.

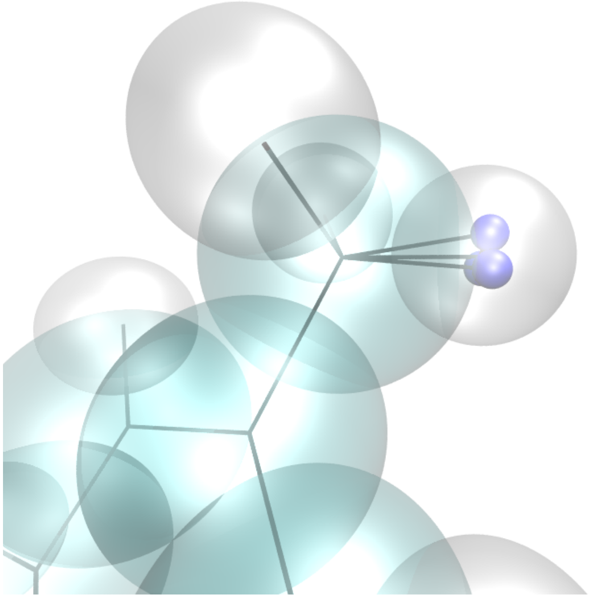

Figure 2.5.: Single topology ligand. The small blue spheres indicate dummy atoms drawn within the van der Waals radius of a hydrogen atom.

end points of the perturbation are mapped onto one another and the value of λ is used to scale potential energy functions. Some difficulty is introduced where the end states have differing numbers of atoms. This requires the introduction of dummy atoms that make no contribution to the energy but which are gradually switched on to interact with the rest of the system. Such dummy atoms are normally withdrawn within the van der Waals radii of their bonding partner (see figure 2.5).

An alternative to single topology approaches is offered through the use of dual topology, allowing one to side step the issues of mapping atoms between ligands and dummy atoms. This makes use of a distinct form of λ-coupling:

$$U_{\lambda}^{dual} = U_0^{intra} + U_1^{intra} + \lambda U_1^{inter} + (1-\lambda) U_0^{inter} + U^{rest} \tag{2.109}$$

where U^{inter} and U^{intra} are the interaction and intra-molecular energies associated with the ligand, and U^{rest} is the energy of the system not involving either ligand. When λ is 0, ligand 0 is fully interacting with the rest of the system and ligand 1 is in the gas phase, and vice versa when λ is 1. The interaction between the two ligands is not included at any time. Using U_{λ}^{dual}, equivalent relative free energy differences can be calculated using the modified thermodynamic cycles in figure 2.6. In the case of the solvation free energy, the cycle can be reduced to a single calculation as the contributions of the gas phase free energy difference are implicitly accounted for. For the binding free energy, it is simple to see that the gas phase contributions within the cycle cancel out making figure 2.6b directly equivalent to figure 2.4b.

$$\mathrm{P}_{(aq)} + \mathrm{Q}_{(g)} \xrightarrow{\Delta\Delta A^{solv}_{P\rightarrow Q}} \mathrm{Q}_{(aq)} + \mathrm{P}_{(g)}$$

(a) Relative free energy of hydration calculation

$$\begin{tikzcd}
\mathrm{P}_{(aq)} + \mathrm{Q}_{(g)} \arrow[r, "\Delta\Delta A^{solv}_{P\rightarrow Q}"] \arrow[d, "\Delta A^{bind}_{P}"'] & \mathrm{Q}_{(aq)} + \mathrm{P}_{(g)} \arrow[d, "\Delta A^{bind}_{Q}"] \\
\mathrm{P}_{(b)} + \mathrm{Q}_{(g)} \arrow[r, "\Delta\Delta A^{(b)}_{P\rightarrow Q}"'] & \mathrm{Q}_{(b)} + \mathrm{P}_{(g)}
\end{tikzcd}$$

(b) Relative free energy of binding cycle

Figure 2.6.: Dual topology thermodynamic cycles for the calculation of relative free energy differences of the two ligands P and Q.

2.4.6. Potential of Mean Force (PMF) Calculations

In parallel with the concept of a potential energy surface we may develop the idea of a free energy landscape, or PMF, that can be defined formally as[64]:

$$A(\lambda) = -\frac{1}{\beta} \ln \int \delta[\lambda - \gamma(\mathbf{x})] \exp\left[-\beta U(\mathbf{x})\right] d\mathbf{x} \tag{2.110}$$

where by analogy we have written the free energy as a function of the variable λ, known in this context as the reaction coordinate, and the corresponding γ is an arbitrary function of the spatial coordinates. The role of γ is to map λ to a chosen degree of freedom within the system. A simple example would be $\gamma(\mathbf{x}) = \mathbf{x}_j^z$, denoting the z-axis coordinate of the jth atom within a system. This definition can be useful to define say the PMF of an ion within an inhomogenous medium such as a lipid bilayer. As γ is arbitrary however λ can be mapped to any degree of freedom desired, for instance the value of a torsion angle, or the distance between the centres of mass of two molecules. This definition can easily be generalised to consider PMFs over multiple degrees of freedom.

The PMF can be thought of as showing how the free energy of a thermodynamic state is partitioned over a particular degree of freedom. As there are many arbitrary ways of quantifying the degrees of freedom of any given system so too there are corresponding PMFs. The physical meaning of $A(\lambda)$ can be made apparent by noting:

$$P(\lambda) = \exp[-\beta A(\lambda)] = \langle \delta[\lambda - \gamma(\mathbf{x})] \rangle \tag{2.111}$$

where $P(\lambda)$ is the probability density of the unconstrained system over λ. Thus, where an appropriate degree of freedom is chosen for the reaction coordinate, the PMF offers considerable information about the behaviour properties of the system under study e.g. prediction of the relative occupancies of different minima.

In theory the PMF could be calculated from an unconstrained simulation by histogramming occurrences of different λ values. This can be highly inefficient where large free energy barriers exist within the PMF however. Several different approaches have been developed therefore to speed convergence of such calculations. The best known of these is Umbrella sampling[65], however we shall consider use of the potential of mean constraint force.

Using the formalism of TI and equation (2.110) the free energy change associated with different values of λ is given by:[66]

$$A(\lambda_B) - A(\lambda_A) = \int_{\lambda_A}^{\lambda_B} \left\langle \frac{\partial(\exp\left[-\beta[U(\mathbf{x};\lambda_B) - U(\mathbf{x};\lambda_A)]\right])}{\partial\lambda} \right\rangle_{\lambda} d\lambda \tag{2.112}$$

where we are considering the free energy difference between the system at λ_A and λ_B

and the potential energy has been written as parametrically dependent on λ. The ensemble average can be interpreted as the mean force experienced by a system constrained to a particular λ value. By differentiating equation (2.112) with respect to λ it is easy to see that the slope of the PMF is directly related to the mean force experienced by a system constrained at the associated λ value. By running a series of MD simulations with different constrained values of λ, and calculating the magnitude of the force required to maintain the constraints, the value of $\frac{dA(\lambda)}{d\lambda}$ at a number of points can be determined. These may then be used to determine $A(\lambda)$ through integration with, for example, the Euler method (appendix B). Algorithms to enforce constraints and provide the corresponding forces are widely implemented, for instance the SHAKE and RATTLE algorithms.[67,68]

This result can be more intuitively rationalised by noting that any degree of freedom within a molecular system may be considered as sampling under a potential that constitutes the averaged effects of the rest of the degrees of freedom of the system. Equation (2.110), the PMF, can in fact be thought of as the definition of this. If we accept this interpretation, then we can use equation (2.59) to demonstrate the relationship between the forces experienced at each λ window and the gradient of the PMF.

2.5. Hybrid Quantum and Classical Free Energy Techniques

Owing to the above stated advantages and disadvantages of classical and QM theory various attempts have been made to combine them in simulation work such as to get the most out of both[5,7,8,69,10]. A range of different strategies have been explored and are developed in technical detail. This is followed by a discussion and comparison of the collection of techniques that comprise the state of the art and the focus of this work.

2.5.1. MM/QM

The first, and perhaps most straightforward, combination of MM and QM theory was developed by Warshel, Levitt and Karplus[5] as the so-called QM/MM method, for which they were awarded 2013 Nobel Prize in Chemistry. They proposed the idea of embedding a QM region within a classical simulation as a method of studying enzyme reaction mechanisms. MM methods are inherently unable to deal with chemical reactions, as they are cannot capture bond breaking or charge transfer polarisation effects. Assuming such effects are only relevant within the reaction site of the enzyme, suitably accurate results can be achieved using this method. By limiting the size of the QM region, only a fraction of the computational expense is required compared to treating the whole system at the QM level. This technique has since been widely employed in the study of biological systems and many others (the reviews Lin & Truhlar [70] and Senn & Thiel [71] give an impression of the extent of its use and details of its implementation).

The main technical difficulty in the implementation of QM/MM is dealing with the interface between the two regions of the system. Covalent bonds that cross the QM/MM boundary can be handled in a number of ways, by introducing additional moeities, or special cap atoms that satisfy the electronic requirements of the QM region[72]. The treatment of the non-covalent interactions across the boundary may also be handled with varying degrees of complexity[73]. The simplest form known as mechanical embedding, treats the van der Waals forces purely classically and does not describe Coulomb interactions. More sophisticated electrostatic embedding maintains the classical Lennard-Jones interactions, but includes the classical point charges within the QM Hamiltonian. This allows the MM region to polarise the QM region, placing it within a more realistic electrostatic context.

The flexibility of QM/MM methods allows them to be used to study a great range of problems. Although considerably cheaper than pure QM methods, use of a QM/MM Hamiltonian in free energy calculations is still comparatively expensive. For this reason specialised approaches for the efficient calculation of free energy differences have been developed. These can be freely combined with the QM/MM approach, providing an additional level of improvement.

2.5.2. QM Corrections with the Zwanzig Equation

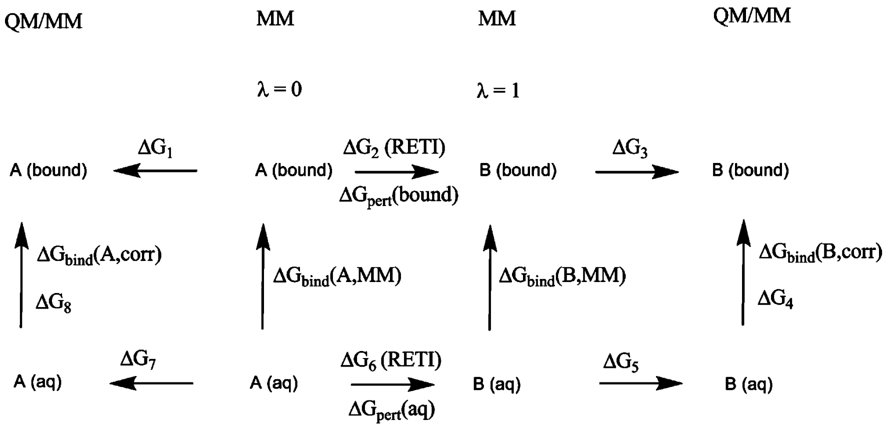

Figure 2.7.: Extended MM/QM free energy cycle, diagram taken from ref Beierlein *et al.* [6].

As developed fully in section 2.4, the Zwanzig equation is able to express the free energy difference between two systems in terms of an ensemble average over only a single state. Noting that the same chemical state, described using different Hamiltonians, constitutes separate thermodynamic states we can apply the same process to calculate the free energy difference between the MM and QM descriptions of the system:

$$\Delta A_{MM \rightarrow QM} = -\frac{1}{\beta} \ln \langle \exp\left[-\beta(U_{QM} - U_{MM})\right] \rangle_{MM} \qquad (2.113)$$

calculations of this form can be combined with a purely classical thermodynamic cycle[10] to give QM corrections to MM free energy differences (see figure 2.7). How this is more efficient than a purely quantum calculation is not immediately obvious, as each MM energy used within the estimator is paired with a QM energy. However, due to the nature of MD and MC sampling, the consecutive configurations produced are highly similar. As such there is no value to include all configurations within the estimator. Instead snapshots are chosen from the trajectory at a constant interval. In this sense the MM state is being used as a reference state, to produce a series of uncorrelated snapshots that can be post processed to the QM level. This is far more efficient than attempting to complete the alchemical perturbation purely at the QM level, which would requiring sampling of the QM state and would generate many correlated configurations at great cost. This approach has also been used by a number of different groups for a range of applications[11,12,13,14,6].

Whilst it is common practice to subdivide a perturbation between states using a λ-coupling approach, in this case any intermediate λ window carries the same computational cost as sampling under the full QM Hamiltonian. To maintain the efficiency of the method therefore it is necessary to carry out all MM to QM perturbations in a single step. This method will be referred to throughout this work as the single-step free energy perturbation (SSFEP) approach.

An appealing improvement on this method is offered by the use of enveloping distribution sampling (EDS).[74] This approach defines a reference Hamiltonian constructed from a combination of the MM and QM Hamiltonians, as given by:

$$U_{EDS} = -\frac{1}{\beta} \ln \left(\exp[-\beta U_{MM}] + \exp[-\beta U_{QM}] \right) \tag{2.114}$$

where U_{ESD} is the reference state potential energy function. An MM to QM free energy difference can then be ascertained by simulation of the reference Hamiltonian and use of the Zwanzig equation to perform two individual perturbs to both the MM and QM states from a single simulation trajectory. As the reference Hamiltonian is constructed of both Hamiltonians it provides good phase space overlap with both the MM and QM states. This ensures good convergence of the perturbations from the reference state. Unfortunately a potential energy function of the form given by equation (2.114), contains QM contributions and hence sampling of the reference state will be as computationally intensive as direct sampling of the QM level.

2.5.3. Molecular Mechanics Based Importance Function (MMBIF)

This collection of methods aims to generate an appropriately distributed QM ensemble from sampling only on an MM potential. The term MMBIF was coined by Iftimie[69,75] from work deriving QM free energy surfaces and reaction rates from MM potentials. This was adopted by Woods and Mulholland[7] and applied to the relative hydration free energy of water and methane. Using the same free energy cycle as the SSFEP methodology but using MMBIF to generate ensembles cheaply at the QM level it is feasible to employ a λ-coupling approach to subdivide the MM to QM perturbation. This allows the use of (more efficient) estimators than the Zwanzig equation. In the case of Woods and Mulholland, TI was used. What follows is a technical development of the technique as used by Iftimie and Mulholland using Multi-'Timestepping' Monte Carlo (MTMC).

MTMC has been used for various applications in computational physics and chemistry[76,77,7]. The basic idea of the technique is to generate a Markov chain at the QM level. Structures for this Markov chain are, however, not generated through the Hastings method (stochastic atomic displacements) but instead by sampling on an auxiliary Markov chain at the MM level. This necessitates the use of a specialised acceptance test given by:

$$P_{MMBIF}(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \min \left(1, \exp\left[-\beta \Delta\Delta U\right]\right) \tag{2.115}$$

$$\Delta\Delta U = \left(U_{QM}(\mathbf{x}_2) - U_{MM}(\mathbf{x}_2)\right) - \left(U_{QM}(\mathbf{x}_1) - U_{MM}(\mathbf{x}_1)\right) \tag{2.116}$$

the algorithm is given in full as follows[7]:

1. From some starting configuration $\mathbf{x}_1$, carry out a series of MC moves at the MM level with the Metropolis-Hastings algorithm (see section 2.4) to generate $\mathbf{x}_2$.
2. Get the QM energy of $\mathbf{x}_2$ i.e. calculate the value of $U_{QM}(\mathbf{x}_2)$
3. Apply the acceptance test equation (2.115)
4. If test passed $\mathbf{x}_2$ is added to the Markov chain becomes $\mathbf{x}_1$
5. If test failed discard $\mathbf{x}_2$ and $\mathbf{x}_1$ is again added to the Markov chain
6. Repeat from 1. as desired

The derivation of the acceptance test equation (2.115) is given as follows. Starting from the general form of the acceptance test at the QM level:

$$P_{QM}^{Acpt}(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \min\left(1, \frac{P_{QM}(\mathbf{x}_2)P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P_{QM}(\mathbf{x}_1)P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)}\right) \tag{2.117}$$

using the Boltzmann distribution for the QM state:

$$\frac{P_{QM}(\mathbf{x}_2)}{P_{QM}(\mathbf{x}_1)} = \exp\left[-\beta(U_{QM}(\mathbf{x}_2) - U_{QM}(\mathbf{x}_1))\right] \tag{2.118}$$

we are left with the need to determine, $\frac{P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)}$. The basis of the technique is that we use a classical Markov chain to carry out the structure generation process, the chance of a structure being generated is therefore the outcome of a series of standard classical MC moves[75]:

$$P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \prod_{i=1}^{n} T_{MM}(\mathbf{x}_1^{i-1} \rightarrow \mathbf{x}_1^{i}) \tag{2.119}$$

here $\mathbf{x}_1^i$ are the intermediate, classically generated, configurations between $\mathbf{x}_1$ and $\mathbf{x}_2$ as we define $\mathbf{x}_1 = \mathbf{x}_1^0$ and $\mathbf{x}_2 = \mathbf{x}_1^n$. T_{MM} is the classical transition probability. This is readily extended to give the quantity we desire:

$$\frac{P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)} = \frac{\prod_{i=1}^{n} T_{MM}(\mathbf{x}_1^{i} \rightarrow \mathbf{x}_1^{i-1})}{\prod_{i=1}^{n} T_{MM}(\mathbf{x}_1^{i-1} \rightarrow \mathbf{x}_1^{i})} \tag{2.120}$$

we use the classical Markov chain detailed balance condition, $\frac{T_{MM}(\mathbf{x}_1^{i} \rightarrow \mathbf{x}_1^{i-1})}{T_{MM}(\mathbf{x}_1^{i-1} \rightarrow \mathbf{x}_1^{i})} = \frac{P_{MM}(\mathbf{x}_1^{i-1})}{P_{MM}(\mathbf{x}_1^{i})}$, to give:

$$\frac{P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)} = \frac{\prod_{i=1}^{n} P_{MM}(\mathbf{x}_1^{i-1})}{\prod_{i=1}^{n} P_{MM}(\mathbf{x}_1^{i})} \tag{2.121}$$

$$= \frac{P_{MM}(\mathbf{x}_1^0)}{P_{MM}(\mathbf{x}_1^n)} \tag{2.122}$$

$$= \frac{P_{MM}(\mathbf{x}_1)}{P_{MM}(\mathbf{x}_2)} \tag{2.123}$$

$$= \exp\left[-\beta(U_{MM}(\mathbf{x}_1) - U_{MM}(\mathbf{x}_2))\right] \tag{2.124}$$

inserting equations (2.118) and (2.124) into equation (2.117) gives the full acceptance test:

$$P_{QM}^{Acpt}(\mathbf{x}_1 \to \mathbf{x}_2) \tag{2.125}$$

$$= \min\left(1, \frac{P_{QM}(\mathbf{x}_2)P_g(\mathbf{x}_2 \to \mathbf{x}_1)}{P_{QM}(\mathbf{x}_1)P_g(\mathbf{x}_1 \to \mathbf{x}_2)}\right) \tag{2.126}$$

$$= \min\left(1, \exp[-\beta(U_{QM}(\mathbf{x}_2) - U_{QM}(\mathbf{x}_1))] \exp[-\beta(U_{MM}(\mathbf{x}_1) - U_{MM}(\mathbf{x}_2))]\right) \tag{2.127}$$

$$= \min\left(1, \exp\left[-\beta\Delta\Delta U\right]\right) \tag{2.128}$$

$$= P_{MMBIF}(2.115) \tag{2.129}$$

2.5.4. Non-Boltzmann Reweighting

The most recently proposed of this class of techniques was put forward by König and Boresch[8]. Borrowing from the theoretical basis of the PMF technique Umbrella Sampling, they use its core relation[65]:

$$\langle X \rangle_u = \frac{\langle X \exp[\beta V] \rangle_b}{\langle \exp[\beta V] \rangle_b} \tag{2.130}$$

here u represents an ensemble average on an unbiased system, and b represents a biased system. V is the bias potential by which the two states differ and X is some property of interest.Most commonly this is chosen as a Harmonic potential and X is chosen as the PMF. More generally however this result can be seen as recovering an ensemble average of property X from sampling under a different Hamiltonian. Sampling of the biased state is sometimes referred to as non-Boltzmann sampling of the unbiased state, giving rise to the name of this technique. Maintaining the formalism of biases for the moment, V can be defined as:

$$V = U_{MM} - U_{QM} \tag{2.131}$$

this definition will allow the recovery of QM ensemble average properties from sampling under an MM Hamiltonian i.e.:

$$\langle X \rangle_{QM} = \frac{\langle X \exp[\beta V] \rangle_{MM}}{\langle \exp[\beta V] \rangle_{MM}} \tag{2.132}$$

this can be profitably applied to the calculation of free energies, starting with the Zwanzig equation at the QM level:

$$\Delta A_{QM,0\to 1} = -\frac{1}{\beta}\ln\left\langle \exp[-\beta(U_{QM,1} - U_{QM,0})]\right\rangle_{QM,0} \tag{2.133}$$

here and in the following subscripts denote both the level of theory and an alchemical change. So whilst $U_{QM,0}$, $U_{QM,1}$, $U_{MM,0}$ and $U_{MM,1}$ all correspond to thermodynamically distinct states, only differences in numeral denote changes to chemically distinct states. We then choose $X = \exp\left[-\beta(U_{QM,1} - U_{QM,0})\right]$:

$$\begin{aligned}
&\exp\left[-\beta\Delta A_{QM,0\to 1}\right] \\
&= \left\langle \exp\left[-\beta(U_{QM,1} - U_{QM,0})\right]\right\rangle_{QM,0} && (2.134)\\
&= \frac{\int \exp\left[-\beta U_{QM,0}\right]\exp\left[-\beta(U_{QM,1} - U_{QM,0})\right] d\mathbf{x}}{\int \exp\left[-\beta U_{QM,0}\right] d\mathbf{x}} && (2.135)\\
&= \frac{\int \exp\left[-\beta(U_{MM,0} - U_{MM,0})\right]\exp\left[-\beta U_{QM,0}\right]\exp\left[-\beta(U_{QM,1} - U_{QM,0})\right] d\mathbf{x}}{\int \exp\left[-\beta(U_{MM,0} - U_{MM,0})\right]\exp\left[-\beta U_{QM,0}\right] d\mathbf{x}} && (2.136)\\
&= \frac{\dfrac{\int \exp\left[-\beta U_{MM,0}\right]\exp\left[-\beta(U_{QM,1} - U_{QM,0})\right]\exp\left[-\beta(U_{QM,0} - U_{MM,0})\right] d\mathbf{x}}{\int \exp\left[-\beta U_{MM,0}\right] d\mathbf{x}}}{\dfrac{\int \exp\left[-\beta U_{MM,0}\right]\exp\left[-\beta(U_{QM,0} - U_{MM,0})\right] d\mathbf{x}}{\int \exp\left[-\beta U_{MM,0}\right] d\mathbf{x}}} && (2.137)\\
&= \frac{\left\langle \exp\left[-\beta(U_{QM,1} - U_{QM,0})\right]\exp[\beta V_0]\right\rangle_{MM,0}}{\left\langle \exp[\beta V_0]\right\rangle_{MM,0}} && (2.138)
\end{aligned}$$

giving the so called non-Boltzmann Zwanzig result:

$$\Delta A_{QM,0\to 1} = -\frac{1}{\beta}\ln\frac{\left\langle \exp\left[-\beta(U_{QM,1} - U_{QM,0})\right]\exp[\beta V_0]\right\rangle_{MM,0}}{\left\langle \exp[\beta V_0]\right\rangle_{MM,0}} \tag{2.139}$$

where $V_0 = U_{MM,0} - U_{QM,0}$. This result in principle gives the correct free energy difference between two QM states, from sampling of only a single MM state and can me applied as a post-processing technique. Note that for each snapshot both $U_{QM,0}$ and $U_{QM,1}$ must be calculated requiring two different QM computations. This does not offer any apparent computational savings over the use of SSFEP therefore. As the difference between the two QM states is likely to be greater then that between each QM and its respective MM state, use of this result is unlikely to be practical. However the above derivation is very flexible and through a completely analogous series of steps can be used to derive non-Boltzmann BAR:

$$\Delta A_{QM,0\to 1} = -\frac{1}{\beta}\ln\left(\frac{\left\langle F(U_{QM,0} - U_{QM,1} + C)\exp[\beta V_1]\right\rangle_{MM,1}\left\langle \exp[\beta V_0]\right\rangle_{MM,0}}{\left\langle F(U_{QM,1} - U_{QM,0} + C)\exp[\beta V_0]\right\rangle_{MM,0}\left\langle \exp[\beta V_1]\right\rangle_{MM,1}}\right) + C \tag{2.140}$$

$$C = \Delta A_{QM,0\to 1} + \frac{1}{\beta}\ln\frac{n_1}{n_0} \tag{2.141}$$

Figure 2.8.: Use of non-Boltzmann BAR (NBB) to calculation QM correction terms to an MM perturbation. Circles in gray represent sampled states across the λ coordinate, states in recovered using reweighting through NBB. Diagram taken from König *et. al.*[8]

where $V_1 = U_{MM,1} - U_{QM,1}$. By allowing use of BAR as an estimator therefore the numerical limitations of the Zwanzig equation are circumvented, at the cost of twice the number of snapshots being processed at the QM level. This does not alleviate the problem of poor overlap between $U_{QM,0}$ and $U_{QM,1}$ so this result remains unlikely to have any practical usage. Instead the non-Boltzmann results can also be used to perturb between MM and QM states. For the non-Boltzmann Zwanzig equation:

$$\Delta A_{QM \to MM,0 \to 1} = -\frac{1}{\beta} \ln \frac{\langle \exp\left[-\beta(U_{MM,1} - U_{QM,0})\right] \exp[\beta V_0] \rangle_{MM,0}}{\langle \exp[\beta V_0] \rangle_{MM,0}} \tag{2.142}$$

similarly, an equivalent result using BAR can be derived. This can be used to efficiently calculate free energy differences by using a typical λ-coupling window at the MM level, but using equation (2.142) to correct the final perturbation at each end-point (see figure 2.8). This allows the majority of the alchemical change to be carried out at the MM level, but the use of BAR to calculate corrections to the QM level.

2.5.5. Discussion

The above techniques for hybrid free energy calculation have mostly been presented in the literature as proof of concept demonstrations in the context of free energies. As the first proposed, SSFEP is the most extensively used[6,10,11,12,13,14,7,15,16,17,9], and has been applied to the most ambitious collection of test systems, including protein-ligand calculations[6]. The use of MMBIF approaches has been limited to simple water to methane perturbations[7] and model fluids[78]. The non-Boltzmann approach of Boresch has been applied to a set of free energy of hydration calculations[79] with both implicit and explicit solvent.

Despite being the most widely used, SSFEP suffers from the drawback of depending on the Zwanzig equation as an estimator. Its formulation allows for efficient one-sided sampling. However, this in turn also leads to the poor numerical behaviour of estimator. Determining the convergence of such calculations is challenging particularly given the restriction that the perturbation must be carried out in a single step. For this reason both MMBIF and non-Boltzmann approaches represent appealing alternatives, having been used with TI or BAR, respectively. Both of these methods come at a greater degree of computational expense however. Use of TI or BAR necessitates that at least two MM ensembles must be post-processed to the QM level. Some questions remain as to whether or not this increased cost is amply compensated by more sophisticated estimators. The use of and validation of SSFEP is the focus of the first chapter.

The use of MMBIF to derive free energy differences has been demonstrated using a water to methane perturbation by Woods *et al.*[7]. A range of different water models was used to carry out this process and in doing so demonstrate the drawback of this approach. It was found that when attempting to use TIP5P as the classical water model very poor acceptance rates for the QM Markov chain are observed, i.e. the system gets stuck and is unable to find configurations suitable for the QM ensemble. As they note this is likely due to poor phase space overlap of the two Hamiltonians. Given the simplicity of the test system they consider, this is perhaps surprising and suggests a potential shortcoming of this technique. Although MMBIF reduces the requirement of overlap between states by allowing the use of TI as an estimator this is replaced by the need for sufficient phase space overlap to carry out construction of the ensemble. Whether this represents a true gain in the efficiency of MM to QM calculations is therefore an open question. On the other hand the acceptance rates of MMBIF moves provide valuable information about the quality of the simulation. In contrast to SSFEP calculations where it is difficult to tell if a perturbation has converged due to poor overlap, MMBIF calculations provide a key indicator that can highlight problematic calculations.

Augmentations to the MMBIF methodology have been proposed that aim to alleviate the degree of overlap required between states[78]. Coe *et al.* note that the temperature and pressure of the classical auxiliary Markov chain can be considered as free parameters, as these properties are corrected to the appropriate values at the QM level by the acceptance test. They then develop the argument that for any simulation there is some optimum value for the classical temperature/pressure that maximises acceptance to the QM ensemble, and propose an approach to determine these values.

Some points of similarity can be remarked between the MMBIF and non-Boltzmann approaches. Both make use of an MM potential as reference state to derive QM properties. In the case of MMBIF the actual ensemble is recovered, whereas non-Boltzmann approaches recover the value of any given ensemble average property by appropriate reweighting of the MM. In principle it is possible to use both MMBIF and non-Boltzmann approaches to carry out perturbations completely at the QM-level, with intermediate alchemical QM states matched by the corresponding alchemical MM reference potential. The QM calculations required for the intermediate states, however, may simply obscure any benefits of this approach.

A practical benefit of the non-Boltzmann and SSFEP approaches are that the QM calculations are not sequentially dependent on one another. In both of these cases, an MM simulation is run, and a series of snapshots are selected and post-processed to the QM level. Such calculations are therefore able to exploit the massively parallel architecture of modern super computers, leading to savings in real time if not in CPU time. In contrast, for MMBIF calculations each new configuration for a QM calculation is dependent on completion of the preceding. As the scalability of individual calculations is inherently limited by system size this can provide a bottleneck in the time taken for such calculations. This drawback is compensated by the power of MMBIF to bias sampling towards the desired ensemble, whereas other approaches are able only to evaluate those configurations that are generated by the MM potential. This introduction of biasing is argued to improve the efficiency of sampling offered by MMBIF by increasing the probability of generating configurations compatible with the QM ensemble[69] (this is considered in chapter 5).

Some reservations have been expressed within the literature regarding the use of hybrid free energy techniques. Many studies have focussed on demonstrating that converged calculations can be obtained through these methodologies however the work of Beierlein[6], Warshel[15] and Woods[7] have all noted a decrease in the accuracy of hybrid free energy calculations relative to purely MM calculations. All three note a decrease in accuracy in the calculation of the free energy of hydration of methane. Woods suggests that this discrepancy is due to MM water models not being well optimised for use in QM/MM work. However, the same explanation cannot be extended to the results of Warshel. Ultimately, the different QM and sampling approaches in these studies make it very difficult to form a solid conclusion; that each seems to encounter the same problem is interesting. It may be that the increased expense of the QM calculations prevents adequate convergence in the examples considered, or simply that the particular test cases considered are already well described by the parameterised MM potentials. That said, Beierlein does achieve improvement in the results of a protein-ligand free energy calculation employing COX-2.

2.6. Software Packages

The aim of this section is to briefly outline the technical details of the software packages used throughout this work for clarity and brevity.

2.6.1. AMBER 12

The AMBER[80] software suite comprises a collection of programs focussed on MD sampling of MM energy models. The core applications that carry out MD are SANDER and PMEMD, dynamics being propagated using the leapfrog integrator. A number of algorithms are available for pressure and temperature regulation of the system. Long range electrostatics are handled through the use of PME or conventional Ewald summation in periodic systems. A CUDA implementation of PMEMD for use on GPUs was also employed. AMBER 12 is used extensively throughout the molecular simulation community.

The MM representation of molecules throughout this work have been produced through use of the ANTECHAMBER application. This executable takes a molecular structure as input and determines GAFF atom types and assigns partial charges for all atoms. A number of methods can be used for partial charge derivation. However, the AM1-bcc method was used exclusively in this work.

2.6.2. ProtoMS 2.6 and 3.0

A code developed in-house, ProtoMS[76], is designed to undertake MC sampling of MM systems using the Metropolis-Hastings algorithm. ProtoMS implements the FF99 and GAFF force fields. As such ANTECHAMBER is used to determine application of the GAFF force field to small molecules. Long range electrostatics are not treated specifically and a distance cut off is used for both LJ and charge-charge interactions.

In addition to periodic boundary conditions, ProtoMS is frequently used to simulate systems with a solvent cap. Here a sphere of waters is placed around the contents of the simulation cell. Any waters that stray outside off the defined cap radius are penalised energetically by a Harmonic term. This is often combined with the use of a protein scoop, whereby residues beyond a specified distance of a user-defined point are removed from the protein structure. The scoop is usually placed around the binding site of protein and used to limit the system size.

2.6.3. CASTEP 5.5

CASTEP[81] is a plane-wave DFT package using Born-Oppenheimer MD with the Velocity Verlet algorithm. Long range electrostatics are treated using Ewald summation. It is designed for massively parallel calculations with strong scalability. It supports the use of pseudopotentials and offers 'On-the-fly' pseudopotential generation. A wide range of functionals are available including hybrid functionals such B3LYP as well PBE and LDA. The employed implementation of the LDA functional is given by Perdew and Zunger[82].

Forces for molecular dynamics and structure minimisation are calculated analytically with the expression:

$$\mathbf{F} = \sum_{i,j} c_i^* c_j \int \psi_j^*(\mathbf{r}) \frac{\partial \hat{H}}{\partial \mathbf{x}} \psi_i(\mathbf{r}') \mathrm{d}\mathbf{r}\mathrm{d}\mathbf{r}' \qquad (2.143)$$

Where ψ_i and ψ_j are the plane-wave functions of the basis, and $\mathbf{x}$ are the nuclear spatial coordinates. This makes the evaluation of forces relatively cheap once a converged electron density has been calculated.[83]

3. Single-Step Free Energy Perturbation

3.1. Introduction

Of the presented approaches to hybrid free energy techniques by far the most commonly applied of these is the single step free energy perturbation (SS-FEP)[6,10,11,12,13,14,7,15,16,17,9]. The Zwanzig equation[2] is used for the calculation of additional QM correction terms to MM based free energy cycles using a single-step free energy perturbation[10] (figure 3.1) as given by:

$$\Delta A_{MM\rightarrow QM} = -\frac{1}{\beta}\ln\left\langle \exp\left[-\beta(U_{QM} - U_{MM})\right]\right\rangle_{MM} \tag{3.1}$$

The one-sided sampling of the Zwanzig equation allows the technique to avoid extremely costly sampling with the QM Hamiltonian. This advantage is countered however by a more stringent requirement for overlap between perturbation end states than other free energy difference estimators (see section 2.5).

The unstable numerical formulation of the Zwanzig equation and its inherent directionality make it difficult to determine whether the condition of sufficient overlap has been met[1] - there may be rare, as yet unsampled configurations, that will heavily influence the calculated free energy difference. To our knowledge, it has yet to be rigorously demonstrated that in general the overlap of QM and MM free energy surfaces is sufficient to allow its use. Previous work from this group has developed an alternative approach based around charge perturbation to test for convergence of hybrid MM and QM calculations[6]. This provides a necessary but not sufficient condition for convergence. We directly address the convergence of single step perturbations by considering the calculation of the reverse QM to MM process. As free energy is a state property the free energy difference between the MM and QM states is invariant based on the direction of the calculation. This provides a rigorous test for convergence based on the condition:

$$\Delta A_{MM\rightarrow QM} + \Delta A_{QM\rightarrow MM} = 0 \tag{3.2}$$

in addition to the previously defined $\Delta A_{MM\rightarrow QM}$ the reverse perturbation, from the QM to the MM state is denoted by $\Delta A_{QM\rightarrow MM}$. In this chapter we present results

$$\begin{tikzcd}
L^{solv}_{MM} \arrow[r, "\Delta A^{solv}_{MM\rightarrow QM}"] \arrow[d, "\Delta A^{bind}_{MM}"'] & L^{solv}_{QM} \arrow[d, "\Delta A^{bind}_{QM}"] \\
L^{bound}_{MM} \arrow[r, "\Delta A^{bound}_{MM\rightarrow QM}"'] & L^{bound}_{QM}
\end{tikzcd}$$

Figure 3.1.: Free energy cycle for the calculation of QM correction terms to an MM free energy difference for the ligand L.

considering direct assessment of the quality of phase space overlap between MM and QM states through extensive generation of QM MD trajectories to allow calculation of the reverse, QM to MM, perturbation. Throughout this work the deviation from zero of equation 3.2 will be referred to as the *discrepancy* of a perturbation.

An adenosine-thymidine DNA base pair is used as a model system, chosen to represent a compromise between biological realism and computational tractability. Previous work has considered the suitability of different MM water models in hybrid calculations[14]. However, the base pair system we consider here provides a far more ambitious and biologically relevant system. The size of the system is sufficient to allow extensive sampling of the QM phase space, whilst also representing a ubiquitous biological dimer.

We generate DFT QM ensembles using MD with the PBE[32] and LDA[29] functionals. The PBE functional has been shown to offer a good compromise between speed and accuracy in describing biological compounds[33], and is frequently used in this context[34,35,36,37,38]. The LDA functional provides a less realistic description of the system's dynamics but usefully demonstrates the behaviour of the single-step perturbation where the MM and QM phase-spaces differ more markedly. Classical trajectories are generated using the AMBER ff99SB[43] and GAFF[48] force fields.

It is common practice when employing hybrid free energy techniques to make use of interaction energies in the place of total energies within the free energy difference estimator[6,10,11,12,13,14,7,15,16,17,9]. The interaction energy of a system, U^{inter}_{AB}, is given by:

$$U^{inter}_{AB} = U_{AB} - U_A - U_B \tag{3.3}$$

where A and B denote two different components of the system. The interaction energy of the two is given by the energy of the complex, U_{AB}, minus the energy of the two components in isolation, U_A and U_B. In the case of a typical MM model the interaction energy can simply be derived by summing the appropriate terms of the force field. For QM models additional calculations for the individual system components are required as the electronic structure of each in isolation will differ from that in the complex due to polarization effects between them. Interaction energies are then simply substituted in the place of total energies within the estimator. In the case of the Zwanzig equation:

$$\Delta A_{0\rightarrow 1} = -\frac{1}{\beta} \ln \left\langle \exp[-\beta \Delta U^{inter}] \right\rangle_0 \tag{3.4}$$

this substitution is not without theoretical difficulties as the derivation of the Zwanzig equation is carried out using total energies. As such the consequences of this approximation are unclear and form the basis of chapter 4. As defined above, the interaction energy of a system includes the energy of polarization and hence free energy calculations using interaction energies are still able to capture these effects.

3.2. Measuring Phase Space Overlap

Given the requirement for phase space overlap between perturbation end states, it is desirable to be able to estimate overlaps of probability distributions within phase space. We here consider two expressions as developed by Bennett[3]. Firstly we note the expression given for direct calculation of phase space overlap in the form of the following integral:

$$O = 2\int \frac{P_0(\mathbf{x})P_1(\mathbf{x})}{P_0(\mathbf{x}) + P_1(\mathbf{x})}\mathrm{d}\mathbf{x} \tag{3.5}$$

where $P_0(\mathbf{x})$ and $P_1(\mathbf{x})$ give the probability of a configuration $\mathbf{x}$ under different states. Owing to the unfeasibility of evaluating integrals of more than a few dimensions, we make use of this expression in only a few single dimensional cases to estimate phase space overlap of particular degrees of freedom. A more practical method for evaluating overlap within higher dimensionality systems is also given by Bennett, directly related to the acceptance ratio method. BAR is given by equations (2.106) and (2.107). As described previously in section 2.4.3, these are solved self-consistently until a converged estimate for $\Delta A_{0\to 1}$ is obtained. This provides the value, C^{opt} that minimises the variance of free energy difference estimate. Bennett notes that the value to which the ensemble averages converge (we shall refer to this as O_{BAR}) provides information about the phase space overlap of the end states of a perturbation, and is given by:

$$O_{BAR} = \left\langle f(U_0 - U_1 + C^{opt})\right\rangle_1 = \left\langle f(U_1 - U_0 - C^{opt})\right\rangle_0 \tag{3.6}$$

a value of 1 would indicate perfect overlap of the Boltzmann distributions of U_0 and U_1 within phase space. Where the converged value of O_{BAR} is small (compared to unity) the overlap of end states is poor. Although this property is discussed directly by Bennett[3], to our knowledge it has not been profitably used to calculate overlaps between perturbation end states. Evaluation of O_{BAR} requires sampling of both states 0 and 1, and hence cannot be used to assess quality of phase space overlaps for use of the Zwanzig equation, without additional simulation work. One might be curious regarding the relationship between equations (3.5) and (3.6). Figure 3.2 demonstrates their precise equivalence in the case of a single dimensional harmonic spring test case, this is more fully explained in section 4.3.

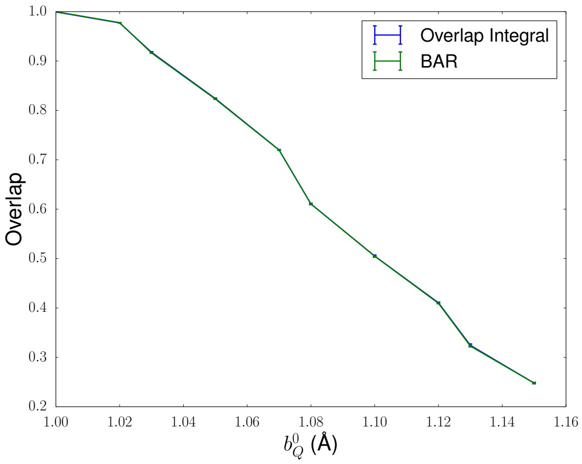

Figure 3.2.: Comparison of O_{BAR} and the overlap integral (equation (3.5)).

3.3. Methods and Calculation Setup

3.3.1. QM Calculations

All QM calculations were carried out using the plane-wave DFT package CASTEP 5.5[81]. Calculations using the LDA[29] functional were carried out with a kinetic energy cutoff of 900 eV with norm-conserving pseudopotentials[84]. PBE[32] calculations used a kinetic energy cutoff of 500 eV and ultra-soft pseudopotentials automatically generated by CASTEP. Kinetic energy cutoffs in each case were tested and chosen based on the requirement of converged energies. Electronic energies were converged to a tolerance of 10^{-5} eV per atom between SCF cycles, using a maximum g-vector of 0.1 Å^{-1} for charge mixing and a grid spacing factor of 2.0 relative to the diameter of the cutoff sphere. A cubic periodic box with sides of 20 Å was used for LDA calculations and 25 Å for PBE; both box sizes are more than sufficient to accommodate the A-T dimer. Long range electrostatics were treated through Ewald summation.

3.3.2. MM Calculations

All MM calculations were carried out using the AMBER 12 software suite[80,85]. Calculations were carried out using both the GAFF[48] and ff99SB[43] force fields. Partial charges for use with the GAFF force field were produced with ANTECHAMBER using the AM1-bcc charge method[86,87]. A cut-off of 8 Å, was used in the calculation of non-bonded interactions and the Particle Mesh Ewald (PME) method was used for long range electrostatics. The PME was validated against the conventional Ewald approach for electrostatics to confirm the equivalent treatment between the MM and QM Hamiltonians (see Supporting Information). A cubic periodic box with sides of 20 Å was used for ff99SB calculations and 25 Å for GAFF.

3.3.3. Molecular Dynamics and SSFEP

The same MD protocol was used for both the MM and QM systems. Initial structures for production MD were generated by the NAB module of AMBER, and subsequently minimised for 50 iterations with the appropriate potential energy function. Bases were modelled with the associated deoxyribose component but without phosphate present. For each Hamiltonian, 5 independent repeats with the same starting configuration were run. All MD calculations were carried out with a time-step of 0.25 fs, as determined by the requirement for constant energy dynamics under the NVE ensemble. Production MD runs were carried out in the NVT ensemble with periodic boundary conditions. Temperature control was achieved using the Langevin thermostat with a collision constant of 0.1 ps^{-1} to regulate the system at 300 K.

The only differences between MD calculations for the MM and QM systems, besides the choice of Hamiltonian, lies in the different algorithms used by CASTEP and AMBER. AMBER employs the leap-frog integrator to solve equations of motion, whilst minimisations employed the conjugate gradient algorithm. Born-Oppenheimer *ab initio* MD calculations in CASTEP employed the velocity Verlet algorithm, whilst minimisations were based on the Broyden-Fletcher-Goldfarb-Shannon (BFGS) algorithm[88].

Generation of QM and MM trajectories was carried out simultaneously and were continued until the discrepancy of all perturbations with interaction energies was close to zero. Five repeat NVT trajectories were calculated under each Hamiltonian. In total, for the GAFF and PBE trajectories, 80,000 configurations were generated per repeat, for the LDA and ff99SB, 5000 configurations were generated per repeat. For calculation of free energy differences the first 500 configurations from each trajectory were discarded as equilibration and every one in ten frames was post-processed with all Hamiltonians for the calculation of total and interaction energies. Calculation of free energy differences was carried out with a custom python script implementation.

3.3.4. Potential of Mean Force Calculations

Potential of Mean Force calculations were carried out using MD with linear constraints with CASTEP.[89] An additional 25 short (1500 time-steps) MD runs were carried out with linear constraints placed on the N-H–N hydrogen bond between thymidine and adenosine. The N-H and H–N bonds were considered as separate degrees of freedom constrained at 0.2 Å intervals, from 1.0 - 1.8 Å. This gives a 5 by 5 grid of points, corresponding to the 25 runs. Constraints were enforced using the RATTLE[67] algorithm. The mean force required to maintain each constraint is equal to the gradient of the free energy surface at that point. The surface itself is then generated through use of the Euler method[90], taking the lowest point of the PMF to be zero.

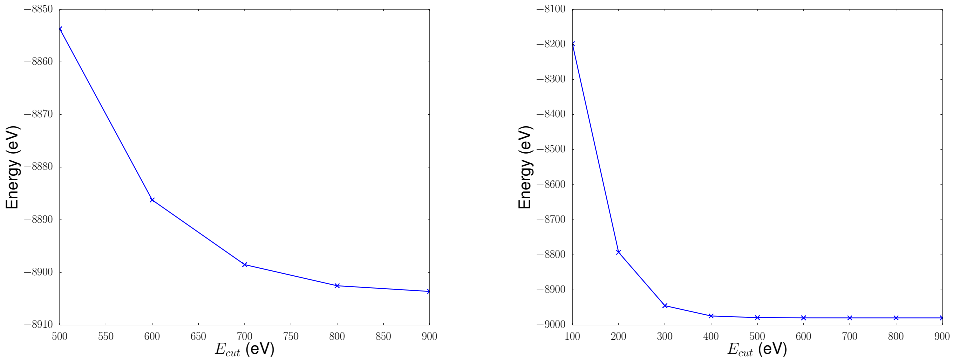

(a) LDA

(b) PBE

Figure 3.3.: Convergence of QM energies with increasing E_{cut}.

3.4. Results and Discussion

3.4.1. Validation of Kinetic Cutoff Energies

As considered more fully in section 2.2.1.7, a necessary prerequisite for the use of a plane wave DFT package such as CASTEP, is the selection of an appropriate kinetic cutoff energy, E_{cut}, for the basis set. This should be sufficiently large to provide converged energies, but is associated with increased computational costs and hence the minimum value that offers convergence is desirable. Calculated energies for the LDA and PBE functionals over a range of E_{cut} are shown in figure 3.3 support the use of the select cutoff values of 500 eV and 900 eV for the PBE and LDA functionals respectively.

3.4.2. Validation of MD Procotol

As an extensive amount of MD work was to be undertaken, it was considered desirable to validate the MD protocol under use. A simple test for the stability of an MD trajectory is its ability to conserve energy under the NVE ensemble. This allows selection of an appropriate timestep for the integration algorithms used by the different software packages (see appendix B and section 2.6). The results shown for short NVE trajectories for both classical MD with ff99SB in AMBER and QM MD with the LDA functional in CASTEP. These results were generated using the same protocol as given by section 3.3 but without the use of the Langevin thermostat. The classical results show good conservation of average energy across all timesteps considered. Although the averages are well maintained considerable differences in the variability of the energy are seen.

Molecular dynamics at the QM level is noted to require shorter time-steps due to the presence of the iterative SCF procedure at every time step[39]. Regardless of the

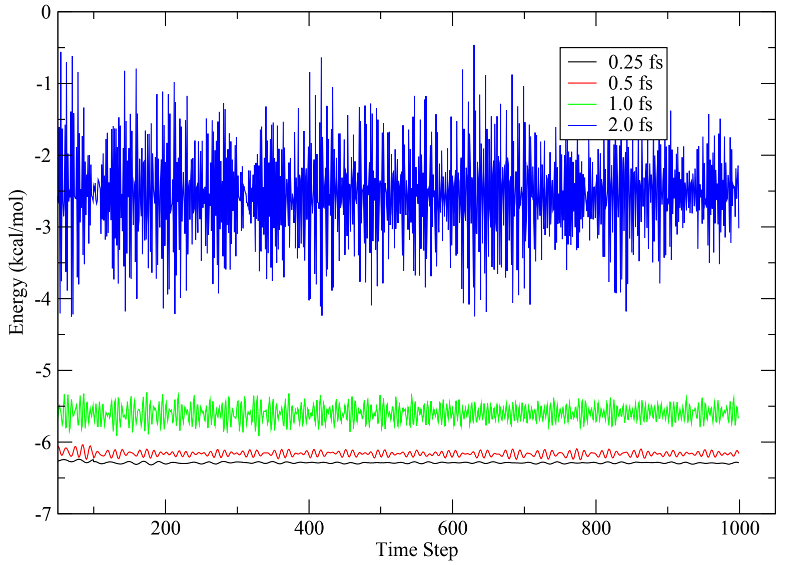

(a) Classical MD

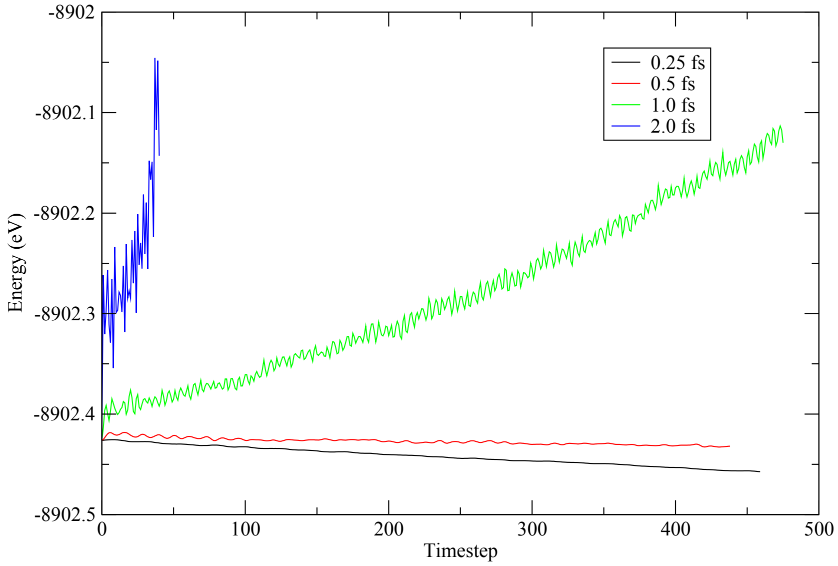

(b) Quantum MD

Figure 3.4.: System total (potential and kinetic) energies under constant energy MD conditions with different timesteps.

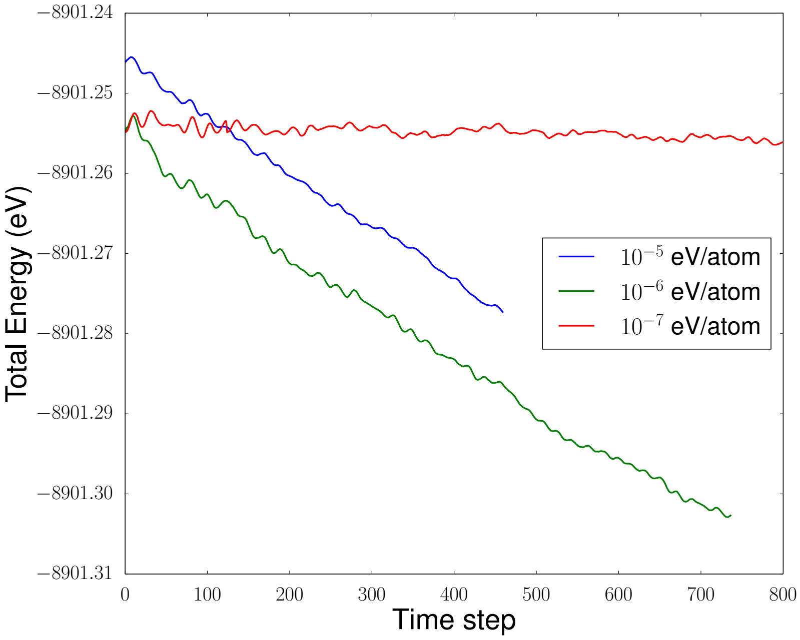

Figure 3.5.: System total (potential and kinetic) energies under constant energy MD conditions with different tolerence values used for convergence of the electronic energy calculation.

timestep used for the QM MD, net drift in the total energy of the system is seen, although this is markedly improved for the 0.5 and 0.25 fs trajectories. This effect arises from the extrapolation of wavefunctions carried out between timesteps. The number of SCF cycles required for convergence of the system energy at each timestep can be greatly reduced if a good estimate of the final density is used to start the procedure. It is typical for QM MD packages therefore to extrapolate an estimate of the system density at each timestep by using the final density of the previous timestep. This extrapolation is then subject to the standard SCF procedure. This practice however has the unfortunate property of breaking the time reversibility of the integration scheme i.e. the energy of any configuration becomes dependent on the configuration before leading to the observed drift in energies[91]. Drift arising for this reason can be reduced by altering the electronic energy convergence tolerance used with the SCF procedure. Results showing the change in total energy drift as a function of the tolerance are given in figure 3.5 for trajectories generated with a timestep of 0.25 fs. The use of 10^{-5} eV/atom and 10^{-6} eV/atom are comparable. However 10^{-7} eV/atom offers a marked improvement, confirming the source of energy drift as being due to the wavefunction extrapolation. The similarity in drift magnitude between the tolerances of 10^{-5} eV/atom and 10^{-6} eV/atom arises due to the length of the convergence window used. CASTEP requires that the electronic energy remains converged to the specified tolerance for a specified number of SCF cycles (in this case 3). This allows the electronic energy to converge beyond the required tolerance. MD runs using 10^{-5}

eV/atom as the required tolerance typically actually converge to 10^{-6} eV/atom, leading to little practical difference between these values. Additional cycles are required however to converge the electronic energy more tightly to 10^{-7} eV/atom, leading to a decrease in the energy drift of the system. Owing to the increased computational cost, however, it was decided to use a 0.25 fs timestep and the default electron energy convergence tolerance of 10^{-5} eV/atom.

3.4.3. Minimised Configurations

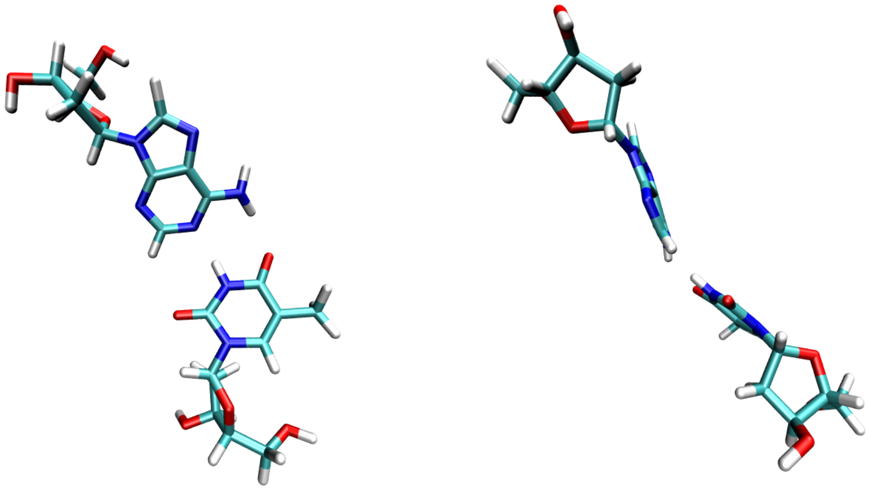

(a) LDA

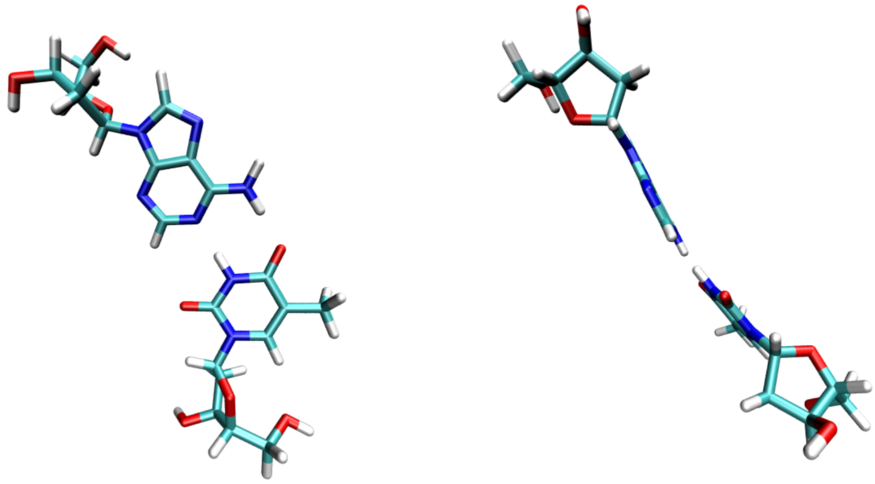

(b) PBE

Figure 3.6.: Initial minimised configurations used for QM MD runs.

The minimised initial configurations used for MD with the four different Hamiltonians are shown in figures 3.6 and 3.7. FF99SB is the only Hamiltonian to give a planar orientation of the bases. This may be due to parameterisation of this force field to reproduce stacking within DNA duplexes. Heavy atom RMSD values be-

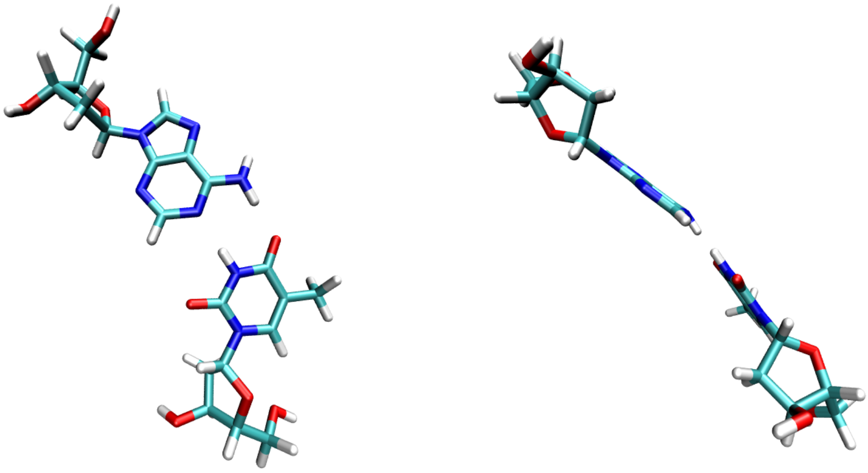

(a) gaff

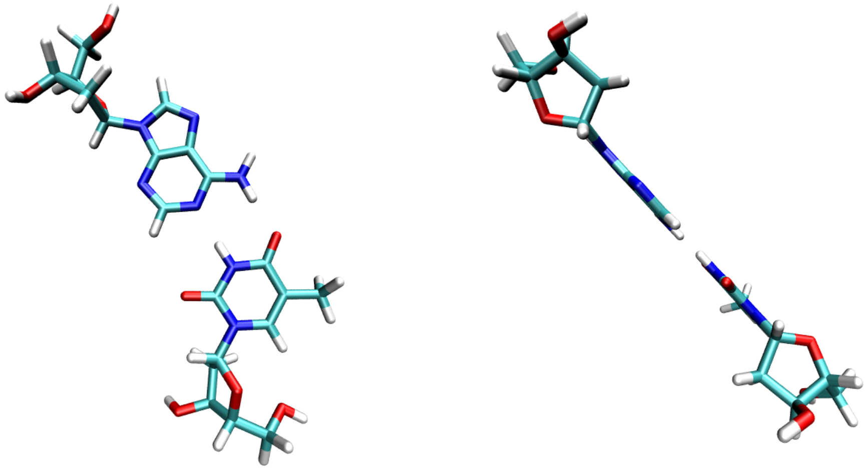

(b) ff99SB

Figure 3.7.: Initial minimised configurations used for classical MD runs.

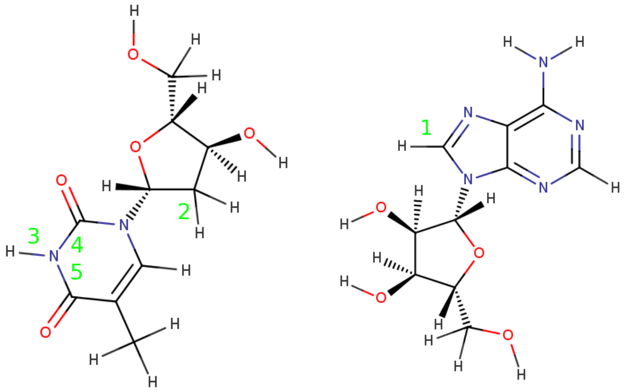

(a)

Hamiltonian	1	2	3	4	5
PBE	1.180	1.154	0.993	1.389	1.382
LDA	1.086	1.099	1.066	1.382	1.391
GAFF	1.080	1.092	1.006	1.320	1.319
FF99SB	1.083	1.091	1.011	1.373	1.374

(b)

Configuration	RMSD vs. PBE	RMSD vs. FF99SB
PBE	0.00	3.85
LDA	0.35	3.93
GAFF	3.98	0.49
FF99SB	3.84	0.00

(c)

Figure 3.8.: 5 bond lengths (in Å) with the greatest difference in minimised configurations between Hamiltonians. Bond positions are highlighted on the structures in (a). Corresponding lengths under different Hamiltonians are given in (b). (c) Heavy atom RMSDs (in Å) of minimized configurations under different Hamiltonians. Structures were aligned to the heavy atoms of the bases, excluding the sugars.

tween different structures are given in figure 3.8c and show the closest accordance between structures at the same level of theory. Agreement of the LDA structure with the more sophisticated PBE functional suggests an, at least, reasonable description of the system is achievable with the former. Although RMSDs provide information about similarities in Cartesian space it is informative to consider explicitly details of the internal geometry. Figure 3.8b gives the five bond lengths within the base pair that show the largest difference between Hamiltonians. It is interesting to note that in each case there tends to be a specific outlier that does agree well with the other structures. Although we might expect the supposedly poorer LDA and GAFF Hamiltonians to show greater variability, no particular trend can be observed within the outliers. For bonds 1 and 2 PBE is the outlier, for bond 3, LDA, for bonds 4 and 5, GAFF. The consequences of marked differences in the bonded degrees of freedom are discussed in more detail in section 3.4.6.

3.4.4. Single-Step FEP

Results of all perturbations between the four considered Hamiltonians are given in figure 3.9. Although our stated aim is to examine perturbations between MM and QM states, it was considered trivial additional work to complete the calculations for all possible perturbations. Using the protocol described in section 3.3.3, gave 8000 data points for each repeat of the perturbations between GAFF and PBE and 450 data points for all other perturbations.

Completion of the larger cycle allows for a more rigorous test of convergence through the computation of cycle closures. Unfortunately cycle closures are non-trivial to calculate as each leg has two separately calculated free energy differences associated with it. Different forward and reverse calculations can be used in any permutation to provide a value for the cycle closure. We compromise by calculating all possible permutations for each cycle and reporting the minimum, maximum and mean unsigned closures. It is immediately apparent from these results that interaction energies provide much tighter cycle closures than using total energies. Although the reported minimum closures using total energies are close to zero, the large associated standard errors suggest this is simply spurious, through a fortunate combination of different components of the cycle. The mean and maximum closures are exceedingly poor however, and suggest the unsuitability of total energies in hybrid free energy work. Although a recent paper has presented results that give successful convergence with total energy calculations, the general applicability of this approach has yet to be demonstrated in a system as complex as that considered here[8].

The convergence of each leg of the cycle can be assessed by calculating the discrepancy between the forward and reverse perturbation as given by equation 3.2. This is shown in figure 3.10. For total energies no particular leg in the cycle can be highlighted as responsible for the poor convergence; even the best converged leg (the PBE to LDA perturbation) has a discrepancy of greater than 10 kcal.mol^{-1}. The use of interaction energies however, is much more compelling. All perturbations fall close to or less than one standard error from zero, with the exception of the PBE to ff99SB calculations (p-value < 0.05 from an unpaired Student's t-test that the free energy differences in either direction are drawn from different distributions).

The magnitude of the free energy differences are considerable when using total energies. Interpretation of these values should be taken with care as all energies calculated are given with respect to an arbitrary reference value, determined by the Hamiltonian. The difference in this reference value between Hamiltonians gives very large apparent free energy differences. The use of the free energy cycles such as in figure 3.1 accounts for this reference state effect and gives meaningful relative free energy changes. Interpretation of the values associated with individual legs of the cycle should be carried out with care.

It might be argued that the convergence of calculations using total energies fails simply due to the large differences in the size of the energy values associated with each Hamiltonian. To test for the possibility of numerical instability caused by difference in reference state, arbitrary constants were used to adjust energy values within individual perturbations. This allows for the exponential terms in figure 3.1 to be scaled to numerically tractable regions and the unadjusted free energy difference can then be recovered by removing the arbitrary constant used. In practice this procedure was found to have no effect on the discrepancy of the each perturbation. Moreover, it can be shown analytically that the discrepancy is invariant with respect to the difference in reference state between Hamiltonians (see appendix D). As long as care is taken to avoid numerical overflows in the exponential terms, the difference in scale of the energy values has no effect on the convergence properties of a calculation.

These results indicate that, in practical terms, the use of single step perturbation techniques is restricted to interaction energies. In addition to the significantly superior convergence properties of interaction energies, they provide a more intuitive interpretation for the resulting free energy differences, as differences in the strength of interaction under different Hamiltonians. For interaction energies, all Hamiltonians share a naturally defined common reference state, namely the two bases at infinite separation. In practice, the use of interaction energies is commonplace with hybrid MM and QM work.[6,10,11,12,13,14,7,15,16,17,9] Despite this prevalence however interaction energies are not formally correct in the context of free energy calculations based on the Zwanzig equation which is derived for total energies. A rigorous theoretical and practical examination of the consequences of using interaction energies are presented in chapter 4. In practice however, the poor convergence of total energy calculations leaves little choice but to use interaction energies.

The failure of calculations using total energies is suggestive of poor overlap between the potential energy surfaces of the different Hamiltonians. That only total energies are affected suggests the problem pertains to the intra-molecular degrees of freedom of the system. This is considered in more detail in a section 3.4.6.

3.4.5. QM MD Trajectories

Within the QM trajectories some examples of proton exchange were observed between the N-H of the thymidine and the N hydrogen bonding partner of adenosine (see figure 3.11). Marked exchange events were observed within two of the five trajectories with the LDA functional; this is particularly significant given their short duration. In contrast, the PBE functional demonstrated comparatively little exchange, only two events occurring within one the five repeats of considerably greater length. Characterisation of the free energy barrier of proton exchange under the LDA functional was carried out through potential of mean force of constraint (PMFC) calculations, using CASTEP. This reveals a free energy barrier of around 1.0 kcal.mol^{-1}, well within the range expected to be crossed due to thermal fluctuations at 300 K. This value is perhaps underestimated due to the coarse resolution of the PMF and the short, constrained trajectories used to generate it. The key features of the landscape appear to be recreated however, and transitions between the minima occur across the saddle point. The observation of hydrogen exchange within this system may also be attributed to the propensity of DFT functionals to underestimate proton exchange barriers.[92]

The comparative rarity of proton exchange events under the more accurate PBE functional suggests that exchange is due to the shortcomings of the LDA functional, leading to unphysically low barriers within the MD runs. Production of an LDA ensemble is still of considerable value as it is noted that a converged free energy difference can still be calculated even where the QM Hamiltonian includes non-classical effects, such as charge transfer or polarisation. Owing to the formulation of the Zwanzig equation, configurations with very high energies in the classical Hamiltonian (such as a highly stretched covalent bond in the case of the proton exchange) are negligibly likely to occur under classical dynamics and hence do not contribute to the free energy difference. Conversely, whilst sampling under the QM Hamiltonian, configurations stabilised by non-classical effects have large negative values of ΔU and hence small contributions to the overall free energy difference.

3.4.6. Phase Space Overlap

The failure of calculations to converge with the use of total energies is indicative of a violation of the requirement for sufficient phase space overlap of not only the MM and QM potential energy surfaces, but of all the energy models. That this problem can be ameliorated with the use of interaction energies suggests the practical reason for the widespread use of this approximation.

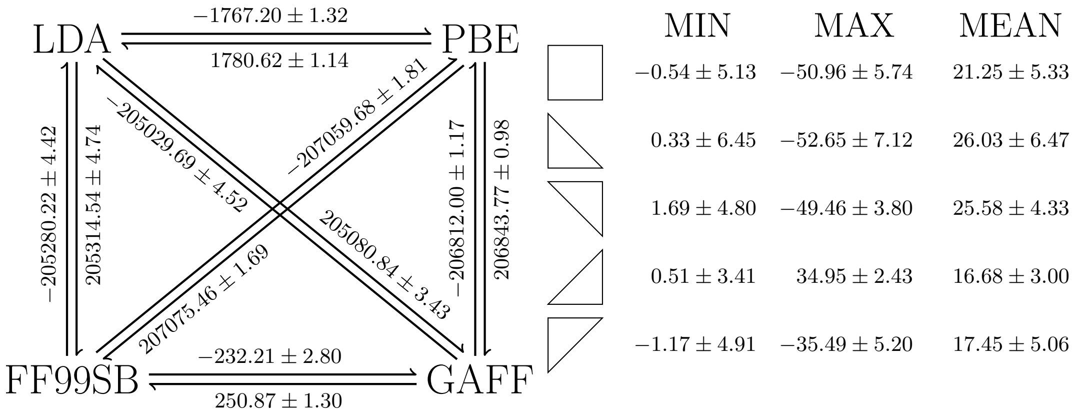

(a) Total energies

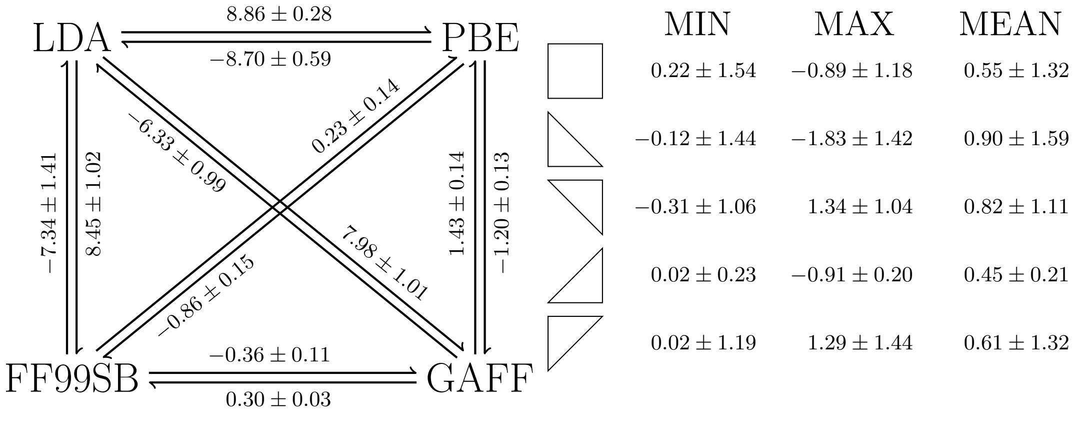

(b) Interaction Energies

Figure 3.9.: Free energy cycles constructed between all Hamiltonians. A single standard error for each perturbation is shown, derived from the standard deviation of the 5 repeats of each calculation. As each leg of the cycle has two associated free energy differences there are multiple different combinations that can be used to construct a closed cycle. We calculated all combinations and on the right of each diagram present the minimum, maximum and mean closures of the illustrated cycle. Standard errors for closures are calculated by summing the variance of each leg of the cycle involved. Standard errors for mean closures were calculated by taking the average variance of all possible leg permutations for each cycle. All free energy differences are in kcal.mol^{-1}.

Perturbation	Total Energies	Interaction Energies
GAFF ↔ ff99SB	$3.71 \times 10^{-3} \pm 3.72 \times 10^{-3}$	0.98 ± 0.00
LDA ↔ ff99SB	$4.38 \times 10^{-10} \pm 8.75 \times 10^{-10}$	0.12 ± 0.04
LDA ↔ GAFF	$1.46 \times 10^{-16} \pm 2.91 \times 10^{-16}$	0.08 ± 0.04
PBE ↔ ff99SB	$5.46 \times 10^{-5} \pm 1.20 \times 10^{-4}$	0.56 ± 0.02
PBE ↔ GAFF	$9.77 \times 10^{-12} \pm 1.38 \times 10^{-11}$	0.40 ± 0.03
PBE ↔ LDA	$2.19 \times 10^{-5} \pm 1.76 \times 10^{-5}$	0.56 ± 0.14

Table 3.1.: O_{BAR} values for each perturbation using total and interaction energies, calculated as described in section 3.2.

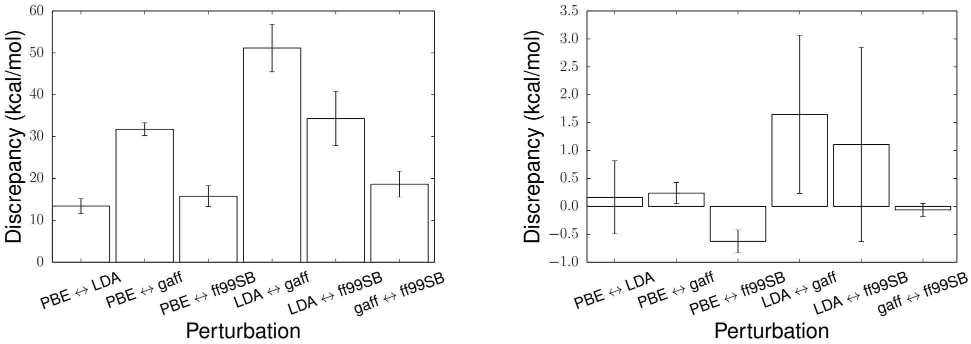

(a) Total energies

(b) Interaction energies

Figure 3.10.: Discrepancies for forward and reverse perturbations within the free cycles. One standard error is shown for all results, calculated by summing the variance of the forward and reverse calculation.

To examine the extent to which using interaction energies improves phase space overlap, the value of O_{BAR} was calculated for all perturbations using total and interaction energies (table 3.1). Although values of O_{BAR} cannot profitably be compared between perturbations due to differing simulation lengths, values for total and interaction energies within perturbations can be compared directly as they are produced from the same data. The use of interaction energies provides between 5 and 16 orders of magnitude improvement in the value O_{BAR}. The smaller overlap values for interaction energies between the LDA functional and classical potentials can be rationalised in terms of the proton transfer events seen in the LDA trajectories. It is comforting to note that the calculated overlap is superior between the PBE functional and classical potentials, than to the LDA functional. Perhaps unsurprisingly, the specialised parameters of the ff99SB force field are noted to offer enhanced overlap with the PBE functional compared to the GAFF force field. Regardless, the values of O_{BAR} presented for perturbations involving GAFF are still more than sufficient to suggest the feasibility of the single-step perturbation.

The striking improvement in phase space overlap provided by interaction energies suggests that the poor total energy results are likely due to the failures in the overlap of intra-molecular degrees of freedom. Using interaction energies reduces the number of degrees of freedom that are considered within the perturbation to exclude intra-molecular terms. Additionally, it is noted generally that intra-molecular potentials tend to be less ‘soft’ then their inter-molecular counterparts. This suggests that in general it is easier to satisfy the required phase space overlap for inter-molecular interactions that have broader probability distributions.

To pinpoint the particular intra-molecular degrees of freedom that give rise to poor total energy overlaps, a simple analysis restricted to the system bond lengths was used. From the trajectory data, distributions for all bond lengths under the PBE and GAFF Hamiltonians were generated, as these are the longest and hence best sampled

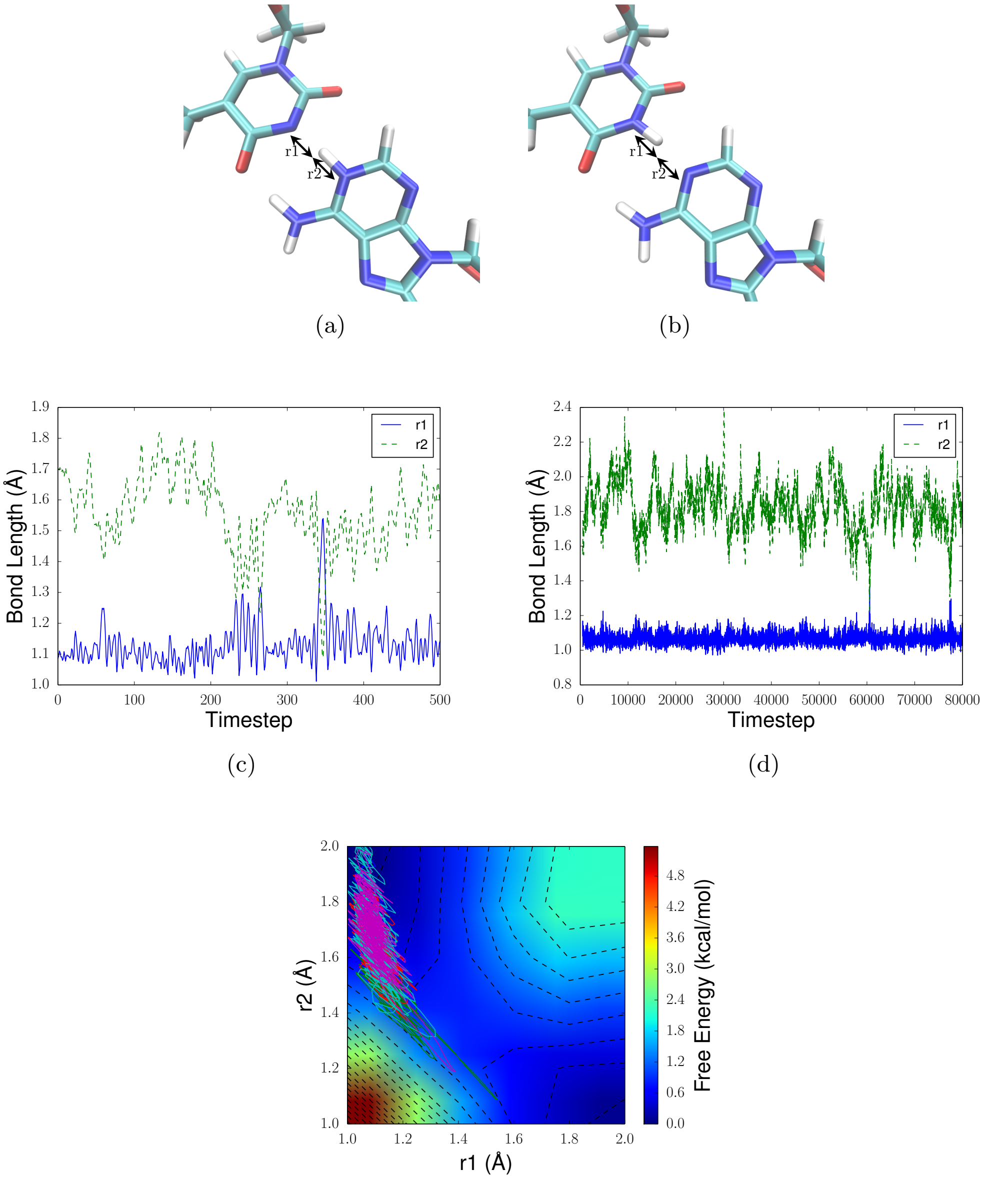

(e)

Figure 3.11.: (a) and (b) show example configurations from an LDA MD run, with the proton exchanged (a) and not exchanged (b). (c) and (d) give time series of r1 and r2 from an example LDA (c) and PBE (d) MD run. (e) Free energy surface of proton exchange between bases using the LDA functional. The solid lines indicate the paths taken by the 5 LDA MD trajectories. Dashed contour lines are plotted every 0.25 kcal.mol^{-1}.

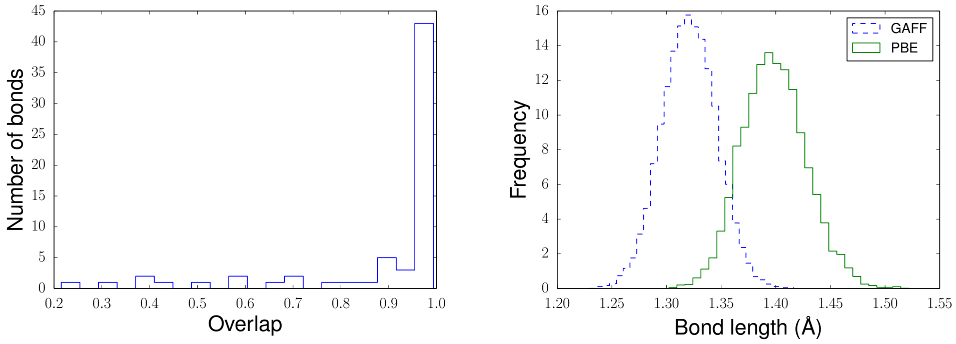

(a) (b)

Figure 3.12.: (a) Distribution of overlaps between GAFF and PBE ensembles for covalent bonds in the base pair calculated using equation 3.5. (b) Distribution of bond C4-O4 bond of thymidine under different Hamiltonians.

trajectories. Overlaps between the distributions of corresponding bonds under different Hamiltonians were then calculated using equation 3.5. The 64 covalent bonds in the base pair give rise to a distribution of overlaps as shown in figure 3.12a. The majority of bonds display excellent overlap between the MM and QM ensembles, but a number demonstrate considerably reduced overlap caused by an offset in equilibrium lengths. The worst example of this is given in figure 3.12b, showing the C4-O4 bond within thymidine (using the Amber force field atom naming conventions[43]).

Each bonded degree of freedom can be approximated as varying independently with respect to the other bonds of the system (figure 3.13 shows a correlation analysis between all bonded degrees of freedom). An estimate of the combined overlap of the PBE and GAFF Hamiltonians can therefore be obtained by taking the product of the overlaps for each individual bond. This overlap estimate is limited to a sub-region of the configuration space of the system as defined by those covalently bonded degrees of freedom and gives a value of 2.132×10^{-5}. This represents a generous upper-bound on the overlap of the two states as the inclusion of additional degrees of freedom can only serve to lower the combined overlap of the system. Although the majority of bonds within the system present an overlap of greater than 0.95, the comparatively small number with poor overlap values can combine to give a globally poor overlap between states. This estimate of the overlap falls short of that required for the convergence of calculations using BAR[93]. As a less efficient estimator, the Zwanzig equation requires even better phase space overlap between states. The values of O_{BAR} for the different perturbations presented in table 3.1 support the use of the Zwanzig equation as they suggest significant overlap is achieved between inter-molecular degrees of freedom.

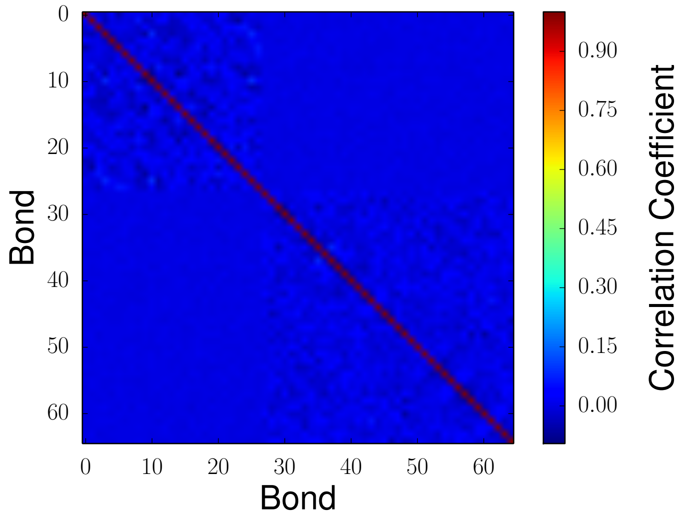

(a) GAFF

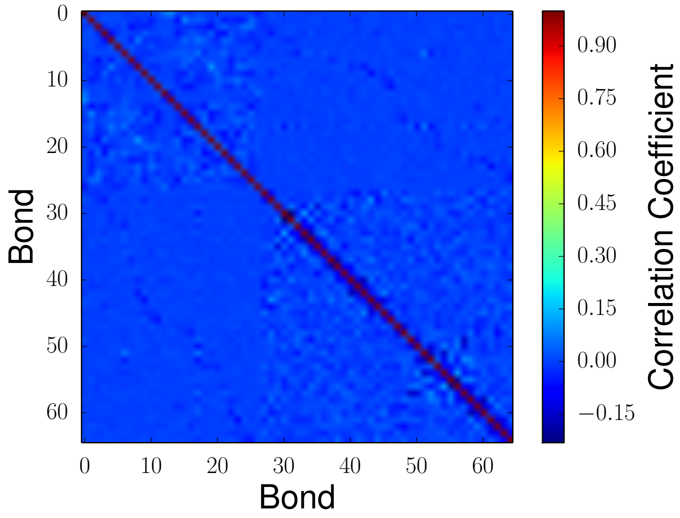

(b) PBE

Figure 3.13.: Correlation matrices for all covalent bonds under the GAFF and PBE Hamiltonians. Correlation coefficients are calculated from simulation data time series.

3.5. Conclusions

The data presented in this chapter constitute a direct validation of the MM to QM single step free energy perturbation procedure, through completion of the reverse QM to MM perturbation. This required the generation of extensive *ab initio* MD trajectories within a model biological system. The A-T DNA dimer chosen for these calculations represents a compromise between biological complexity and expense of calculations. In total over 100 ps of *ab initio* MD was generated using plane-wave DFT.

Importantly the practical restriction that perturbations must be carried out with interaction energies instead of total potential energies is established. Discrepancies between forward and reverse perturbations are shown to be on the order of 10s of kcal.mol^{-1}, for total energies, but nearing zero for interaction energies. Although single step perturbation techniques have been used for some time, the requirement to use interaction energies is often glossed-over or not explicitly stated.

The failure of total energy calculations with the Zwanzig equation is explained in terms of poor phase space overlap between MM and QM Hamiltonians. Marked differences between the phase-space distributions of intra-molecular degrees of freedom are highlighted as problematic. Although limited to only the covalently bonded degrees of freedom, our analysis gives very low upper-bound estimates for total energy phase-space overlap. This analysis also suggests caution in hybrid free energy work around the common practice of enforcing bond length constraints. Constraints may improve overlap between MM and QM ensembles, by removing problematic degrees of freedom from being sampled, but run the risk of constraining ensembles outside their global minimum, distorting calculated free energy differences. The extent to which this problem may be avoided through the use of interaction energies is unclear. König *et al.* has examined the effect of bond length constraints in a simple hybrid free energy perturbation of ethane to methanol[8].

Interaction energy calculations are demonstrated to exhibit markedly better overlap between ensembles, and improved convergence of single step free energy calculations. The presence of non-classical proton exchange interactions between the bases does not prevent stable convergence of the calculated free energy differences.

In addition to the single step methodology considered here, there have been other notable suggestions for hybrid free energy work based around more elaborate sampling or reweighting techniques[7,8]. The generated QM ensembles to provide a valuable data set for the analysis of other methodologies, and this is considered chapter 5. The contents of this paper have been published under the title “Direct Validation of the Single Step Classical to Quantum Free Energy Perturbation”[94].

4. Interaction Energies

4.1. Introduction

An aspect of hybrid free energy calculations that is not widely discussed in the literature is the use of interaction energies. The interaction energy of a system is somewhat subjectively defined. Generally, where the study of the interaction of two components of a thermodynamic system are considered[34]:

$$U_{AB}^{inter} = U_{AB} - U_A - U_B \tag{4.1}$$

$$U_{AB}^{inter} = U_{AB} - U_{AB}^{intra} \tag{4.2}$$

Where the interaction energy, U_{AB}^{inter}, of the two system components A and B is given by the total energy of the complex, U_{AB}, minus the total energy of each component in isolation, U_A and U_B. We here refer to the combined energies of the isolated system components as the system's intramolecular energy, U_{AB}^{intra}. This definition can be intuitively applied to systems of common interest for free energy calculations. When calculating free energies of hydration for instance, the solvent can be considered one system component and the solute another. In the case of protein-ligand free energy calculations the protein can be grouped with the solvent as one component and the ligand treated as the second. Despite the name, the intramolecular energy term may contain energetic interactions between molecules (eg. solvent molecules), although never between the two different components of the system. The terminology is somewhat inconsistent therefore and might be better thought of as the intra-component energy, however we shall proceed with the use of intra-molecular.

As demonstrated in the results of chapter 3 the use of interaction energies is largely necessitated in the hybrid free energy calculations in order to provide sufficient overlap between perturbation end-points . Some benefit is also gained in certain situations by simplifying free energy cycles, as when using interaction energies, some free energy differences are necessarily zero, eg. the alchemical transformation of a ligand *in vacuo*. This produces free energy calculations analogous to those used in dual topology calculations (figure 2.6).

Interaction energies are used within free energy calculations by simply substituting them in the place of the appropriate total energy values. For instance in the case of the Zwanzig equation, equation (2.80) becomes:

$$\Delta A_{0\to1}^{inter} = -\frac{1}{\beta}\ln\left\langle \exp\left[-\beta\Delta U^{inter}\right]\right\rangle_0 \tag{4.3}$$

Where $\Delta U^{inter} = U_1^{inter} - U_0^{inter}$. The Zwanzig equation is derived with the use of total energies and we are not aware of any attempt to justify the use of interaction energies on a statistical mechanics basis. Although it is only explicitly stated in a few papers[6,10,11,12,13,14,7,15,16] it is understood that the use of interaction energies is implicit in the majority of hybrid free energy work.

Similarly, the effect of interaction energies on calculations using MMBIF and non-Boltzmann approaches are unknown. In their application to water and methane, Woods and Mulholland go so far as to keep the solutes entirely rigid, thereby removing the effect of intramolecular degrees of freedom[7]. The extent to which this may be expected to achieve realistic results when applied to more complex systems of biological interest is unknown. Their approaches are addressed in chapter 5.

A great deal of theoretical uncertainty is introduced through the use of interaction energies as free energy techniques are formally derived using total energies[2,4,3] and the validity of carrying out a straightforward substitution with interaction energies is not readily apparent. This chapter shall consider the use of interaction energies as a proxy for total energies in free energy calculations through a thorough first-principles statistical mechanics analysis. It is shown that under certain conditions interaction energies fail as an appropriate substitute for total energies. To our knowledge this forms a novel theoretical critique of common practice in hybrid free energy calculations and further theoretical observations are made that lay the ground work for design of appropriate diagnostic tools for use on a calculation by calculation basis. The derived theoretical results are applied and verified in a series of test systems of increasing realism. The wider implications in the context of single and dual topology calculations are also considered.

4.2. Theoretical Analysis

Here we work from first-principles statistical mechanics to provide some insight into the effects of using interaction energies in free energy calculations. The use of interaction energies is most common with the single step free energy perturbation (SSFEP) approach. For this reason interaction energies within SSFEP are used as a starting point for this analysis. To recap briefly, SSFEP is carried out through application of the Zwanzig equation[2]:

$$\Delta A_{0\to 1} = -\frac{1}{\beta}\ln\langle \exp[-\beta\Delta U]\rangle_0 \tag{4.4}$$

for clarity we reiterate that $\Delta A_{0\to 1}$ is a configurational free energy difference between the thermodynamic states 0 and 1, where we only consider the potential energy terms from the system Hamiltonian. All references to total energies in this chapter refer to total potential energies, as the kinetic energy component can be factorised out and treated analytically where necessary (see section 2.4). Equation (4.4) is an ensemble average and can therefore be rewritten as an integral over configuration space[2]:

$$\Delta A_{0\to 1} = -\frac{1}{\beta}\ln\int P_0 \exp\left[-\beta\Delta U(\mathbf{x})\right] d\mathbf{x} \tag{4.5}$$

Where $\mathbf{x}$ are the configurational degrees of freedom of the system and P_0 is the Boltzmann probability term:

$$\begin{aligned} P_0(\mathbf{x}) &= \frac{\exp\left[-\beta U_0(\mathbf{x})\right]}{Z_0} && (4.6)\\ &= \frac{\exp\left[-\beta U_0(\mathbf{x})\right]}{\int \exp\left[-\beta U_0(\mathbf{x})\right] d\mathbf{x}} && (4.7)\end{aligned}$$

the use of interaction energies with (4.4) in practice is carried out by simple substitution of interaction energies in the place of total energies, to give the interaction free energy difference:

$$\begin{aligned} \Delta A^{inter}_{0\to 1} &= -\frac{1}{\beta}\ln\left\langle \exp\left[-\beta\Delta U^{inter}\right]\right\rangle_0 && (4.8)\\ \Delta A^{inter}_{0\to 1} &= -\frac{1}{\beta}\ln\int P_0(\mathbf{x})\exp\left[-\beta\Delta U^{inter}(\mathbf{x})\right] d\mathbf{x} && (4.9)\end{aligned}$$

where $\Delta U^{inter} = U^{inter}_1 - U^{inter}_0$. The Boltzmann probability term P_0 is unchanged between equations (4.5) and (4.9) as the Hamiltonian under which sampling is carried out is the same. Although they are used interchangeably in hybrid free energy work, it is not self evident however, what relation $\Delta A^{inter}_{0\to 1}$ bears to $\Delta A_{0\to 1}$. The Zwanzig equation is derived with the use of total energies and the physical interpretation of the weighting of interaction energies by the total energy Boltzmann factor is

$$\begin{tikzcd}
1{:}\mathrm{P}_{(g)} \arrow[r, "\Delta \mathrm{A}_{1\to 2}"'] \arrow[d, "\Delta \mathrm{A}_{hyd}\mathrm{P}"] & 2{:}\mathrm{Q}_{(g)} \arrow[d, "\Delta \mathrm{A}_{hyd}\mathrm{Q}"'] \\
3{:}\mathrm{P}_{(aq)} \arrow[r, "\Delta \mathrm{A}_{3\to 4}"] & 4{:}\mathrm{Q}_{(aq)}
\end{tikzcd}$$

Figure 4.1.: Example thermodynamic cycle for the calculation of relative free energies of hydration of P and Q. The shown numbers are used to label the different states of the cycle. The free energy of hydration for P and Q are denoted by ΔA_{hyd}P and ΔA_{hyd}Q respectively. Whilst $\Delta A_{1\to 2}$ and $\Delta A_{3\to 4}$ give the free energy differences between states 1 and 2 and states 3 and 4 respectively.

not readily apparent; certainly its treatment as being equivalent to equation (4.4) is difficult to justify without further analysis.

It should be noted that whilst the values given by equations (4.5) and (4.9) are not directly equivalent, interaction energies are not employed to calculate individual free energy differences between two thermodynamic states. Instead they are employed within thermodynamic cycles to calculate relative free energies of hydration or binding such that an additional free energy calculation is always performed. For this reason, we shall proceed by considering the application of equations (4.5) and (4.9) within the context of a relative free energy of hydration calculation. As given in figure 4.1 we shall consider the notional relative free energy of hydration of the compounds P and Q. Using equations (4.5) and (4.9) we can give expressions for the relative free energy of hydration, $\Delta\Delta A_{hyd}$, and its equivalent using interaction energies $\Delta\Delta A_{hyd}^{inter}$:

$$\Delta\Delta A_{hyd} = \Delta A_{3\to 4} - \Delta A_{1\to 2} \tag{4.10}$$

$$= -\frac{1}{\beta}\ln\int P_3 \exp\left[-\beta(U_4 - U_3)\right]\mathrm{d}\mathbf{x} + \frac{1}{\beta}\ln\int P_1 \exp\left[-\beta(U_2 - U_1)\right]\mathrm{d}\mathbf{x} \tag{4.11}$$

Combining logarithms:

$$= -\frac{1}{\beta}\ln\frac{\int P_3 \exp\left[-\beta(U_4 - U_3)\right]\mathrm{d}\mathbf{x}}{\int P_1 \exp\left[-\beta(U_2 - U_1)\right]\mathrm{d}\mathbf{x}} \tag{4.12}$$

Expanding P_3 and P_1:

$$= -\frac{1}{\beta}\ln\frac{\dfrac{\int \exp\left[-\beta U_3\right]\exp\left[-\beta(U_4 - U_3)\right]\mathrm{d}\mathbf{x}}{Z_3}}{\dfrac{\int \exp\left[-\beta U_1\right]\exp\left[-\beta(U_2 - U_1)\right]\mathrm{d}\mathbf{x}}{Z_1}} \tag{4.13}$$

Expanding Z_3 and Z_1, and cancelling U_3 and U_1:

$$= -\frac{1}{\beta} \ln \frac{\dfrac{\int \exp\left[-\beta U_4\right] \mathrm{d}\mathbf{x}}{\int \exp\left[-\beta U_3\right] \mathrm{d}\mathbf{x}}}{\dfrac{\int \exp\left[-\beta U_2\right] \mathrm{d}\mathbf{x}}{\int \exp\left[-\beta U_1\right] \mathrm{d}\mathbf{x}}} \tag{4.14}$$

and:

$$\Delta\Delta A^{inter}_{hyd} = \Delta A^{inter}_{3\to 4} \tag{4.15}$$

$$= -\frac{1}{\beta} \ln \int P_3 \exp\left[-\beta(U^{inter}_4 - U^{inter}_3)\right] \mathrm{d}\mathbf{x} \tag{4.16}$$

Expanding P_3:

$$= -\frac{1}{\beta} \ln \frac{\int \exp\left[-\beta(U^{intra}_3 + U^{inter}_3)\right] \exp\left[-\beta(U^{inter}_4 - U^{inter}_3)\right] \mathrm{d}\mathbf{x}}{Z_3} \tag{4.17}$$

Expanding Z_3 and cancelling U^{inter}_3:

$$= -\frac{1}{\beta} \ln \frac{\int \exp\left[-\beta(U^{intra}_3 + U^{inter}_4)\right] \mathrm{d}\mathbf{x}}{\int \exp\left[-\beta U_3\right] \mathrm{d}\mathbf{x}} \tag{4.18}$$

here the subscripted numbers are used to denote the different thermodynamic states corresponding to the cycle of figure 4.1. We have made use of the equality noted in section 2.4.4 to relate the relative hydration free energy to the two alchemical perturbations $\Delta A_{3\to 4}$ and $\Delta A_{1\to 2}$. Although the total energy expression is the standard ratio of partition functions, the interaction energy term contains the partition function of state 3, along with a hybrid partition function containing the intramolecular terms of state 3 and and the intermolecular terms of state 4. The interaction free energy difference can therefore be thought of as the perturbation between state 3 and the hybrid state. Equations (4.14) and (4.18) are general expressions of the relative free energy difference. We can tailor these for the calculation at hand by noting the following decompositions for the potential energy at each point in the cycle:

$$U_1(\mathbf{x}) = U^l_{\mathrm{P}}(\mathbf{x}) \tag{4.19}$$

$$U_2(\mathbf{x}) = U^l_{\mathrm{Q}}(\mathbf{x}) \tag{4.20}$$

$$U_3(\mathbf{x}) = U^{intra}_3(\mathbf{x}) + U^{inter}_3(\mathbf{x}) = U^l_{\mathrm{P}}(\mathbf{x}) + U^s(\mathbf{x}) + U^{inter}_3(\mathbf{x}) \tag{4.21}$$

$$U_4(\mathbf{x}) = U^{intra}_4(\mathbf{x}) + U^{inter}_4(\mathbf{x}) = U^l_{\mathrm{Q}}(\mathbf{x}) + U^s(\mathbf{x}) + U^{inter}_4(\mathbf{x}) \tag{4.22}$$

states 1 and 2 correspond to the molecule *in vacuo* and hence only contain the ligand energy terms, $U^l_{\mathrm{P}}(\mathbf{x})$ and $U^l_{\mathrm{Q}}(\mathbf{x})$. The potential energy function used for the ligands is the same for states 1 and 3 and states 2 and 4 respectively. Additionally states 3 and 4 contain contain solvent molecules as an environment for the ligand and as such we have introduced the observation that the intramolecular energies can be decomposed into the contributions of the ligand, U^l, and the solvent U^s. The solvent molecules intramolecular energy term, U^s, is unchanged between states 3 and 4 and hence not subscripted. We note that this assumption is true in the case of a chemical perturbation as considered here. In the case of an MM to QM perturbation where states 3 and 4, are described at a different level of theory, the energy model used for the solvent is different between the states and hence separate $U_3^s(\mathbf{x})$ and $U_4^s(\mathbf{x})$ potential energy functions would be required. This distinction does not qualitatively influence the results presented here and the MM to QM perturbation case is not explicitly considered and we proceed with the more general decompositions as given. The interaction potential energy terms always vary depending on the ligand present and hence are distinct between states 3 and 4. From equation (4.14):

$$\Delta\Delta A_{hyd} = -\frac{1}{\beta}\ln\frac{\dfrac{\int \exp\left[-\beta(U_{\mathrm{Q}}^l + U^s + U_4^{inter})\right]\mathrm{d}\mathbf{x}}{\int \exp\left[-\beta(U_{\mathrm{P}}^l + U^s + U_3^{inter})\right]\mathrm{d}\mathbf{x}}}{\dfrac{\int \exp\left[-\beta U_{\mathrm{Q}}^l\right]\mathrm{d}\mathbf{x}}{\int \exp\left[-\beta U_{\mathrm{P}}^l\right]\mathrm{d}\mathbf{x}}} \tag{4.23}$$

and from equation (4.18):

$$\Delta\Delta A_{hyd}^{inter} = -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_{\mathrm{P}}^l + U^s + U_4^{inter})\right]\mathrm{d}\mathbf{x}}{\int \exp\left[-\beta(U_{\mathrm{P}}^l + U^s + U_3^{inter})\right]\mathrm{d}\mathbf{x}} \tag{4.24}$$

this now looks a little more promising and it is apparent that if cancellations of the solvent and ligand contributions to the free energy difference were allowed these expressions would be equivalent. However as currently written in equations (4.19) to (4.22) all of the potential energy terms are a function of the same set of Cartesian coordinates, $\mathbf{x}$. We can proceed further by introducing the concept of internal coordinates, which provide an alternate description of the degrees of freedom of molecular system,[95] as applied to equations (4.19) to (4.22):

$$U_1(\mathbf{r}^l) = U_{\mathrm{P}}^l(\mathbf{r}^l) \tag{4.25}$$

$$U_2(\mathbf{r}^l) = U_{\mathrm{Q}}^l(\mathbf{r}^l) \tag{4.26}$$

$$U_3(\mathbf{r},\mathbf{q}) = U_3^{intra}(\mathbf{r}) + U_3^{inter}(\mathbf{r},\mathbf{q}) = U_{\mathrm{P}}^l(\mathbf{r}^l) + U^s(\mathbf{r}^s) + U_3^{inter}(\mathbf{r},\mathbf{q}) \tag{4.27}$$

$$U_4(\mathbf{r},\mathbf{q}) = U_4^{intra}(\mathbf{r}) + U_4^{inter}(\mathbf{r},\mathbf{q}) = U_{\mathrm{Q}}^l(\mathbf{r}^l) + U^s(\mathbf{r}^s) + U_4^{inter}(\mathbf{r},\mathbf{q}) \tag{4.28}$$

we are choosing here to represent U as a function of an internal coordinate set such that N degrees of freedom $\mathbf{x}$ have been subdivided into three different categories. In general there will be a set of G global degrees of freedom that correspond to rigid body translations and rotations of the system. These do not contribute to the potential energy of the system and hence do not appear in the above. Of the remaining N - G degrees of freedom, a subset $\mathbf{r}$ of size L degrees of freedom can be defined such that the intramolecular energy is a function of only $\mathbf{r}$. The final N - L - G degrees of freedom are the interaction degrees of freedom and are denoted $\mathbf{q}$. The interaction energy of the system can then be written as a function of both $\mathbf{r}$ and $\mathbf{q}$. The ability to write the potential energy in this form is independent of the energy model used. Thus these results can be readily applied in the context of both MM and QM calculations. We also note that the intramolecular degrees of freedom can be decomposed into those of the ligand being mutated, and the rest of the system, denoted by the superscripts l and s. These terms are directly comparable to equations (4.19) to (4.22), but written using the internal coordinate system. Whilst U^l and U^s can be written as functions of their individual degrees of freedom, the interaction energy terms remain a function of all of non-global degrees of freedom of the system. Updating equations (4.23) and (4.24) using the internal coordinate system given by equations (4.25) to (4.28) gives:

$$\Delta\Delta A_{hyd} = -\frac{1}{\beta}\ln\frac{\dfrac{\int \exp\left[-\beta(U^l_{\mathrm{Q}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_4(\mathbf{r},\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}{\int \exp\left[-\beta(U^l_{\mathrm{P}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_3(\mathbf{r},\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U^l_{\mathrm{Q}}(\mathbf{r}^l)\right]d\mathbf{r}^l}{\int \exp\left[-\beta U^l_{\mathrm{P}}(\mathbf{r}^l)\right]d\mathbf{r}^l}} \tag{4.29}$$

and:

$$\Delta\Delta A^{inter}_{hyd} = -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U^l_{\mathrm{P}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_4(\mathbf{r},\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}{\int \exp\left[-\beta(U^l_{\mathrm{P}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_3(\mathbf{r},\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}} \tag{4.30}$$

as we have undertaken a change in coordinate system, the contribution of the Jacobian should be considered[96], as each change of the coordinate system is paired within a fraction, however these contributions to the free energy difference cancel exactly. Adopting internal coordinates gives expressions with a reduced degree of dependence between the terms of the potential energy function. However separation of the integral into a product of simpler terms is still precluded by the interaction energy term that remains a function of all the degrees of freedom of the system. However, a separated approximation of the equations (4.29) and (4.30) can be given by:

$$\begin{aligned}&\Delta\Delta A_{hyd}\\ \approx&\Delta\Delta A^{(sep)}_{hyd}\end{aligned} \tag{4.31}$$

Introducing $\langle\mathbf{r}\rangle$ in the place of $\mathbf{r}$:

$$= -\frac{1}{\beta}\ln\frac{\dfrac{\int \exp\left[-\beta(U^l_{\mathrm{Q}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_4(\langle\mathbf{r}\rangle_{\mathrm{Q}},\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}{\int \exp\left[-\beta(U^l_{\mathrm{P}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_3(\langle\mathbf{r}\rangle_{\mathrm{P}},\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U^l_{\mathrm{Q}}(\mathbf{r}^l)\right]d\mathbf{r}^l}{\int \exp\left[-\beta U^l_{\mathrm{P}}(\mathbf{r}^l)\right]d\mathbf{r}^l}} \tag{4.32}$$

Separating the integrals:

$$= -\frac{1}{\beta}\ln\frac{\dfrac{\int\exp\left[-\beta U^{l}_{\mathrm{Q}}(\mathbf{r}^{l})\right]d\mathbf{r}^{l}\int\exp\left[-\beta U^{s}(\mathbf{r}^{s})\right]d\mathbf{r}^{s}\int\exp\left[-\beta U^{inter}_{4}(\langle\mathbf{r}\rangle_{\mathrm{Q}}\,,\mathbf{q}))\right]d\mathbf{q}}{\int\exp\left[-\beta U^{l}_{\mathrm{P}}(\mathbf{r}^{l})\right]d\mathbf{r}^{l}\int\exp\left[-\beta U^{s}(\mathbf{r}^{s})\right]d\mathbf{r}^{s}\int\exp\left[-\beta U^{inter}_{3}(\langle\mathbf{r}\rangle_{\mathrm{P}}\,,\mathbf{q}))\right]d\mathbf{q}}}{\dfrac{\int\exp\left[-\beta U^{l}_{\mathrm{Q}}(\mathbf{r}^{l})\right]d\mathbf{r}^{l}}{\int\exp\left[-\beta U^{l}_{\mathrm{P}}(\mathbf{r}^{l})\right]d\mathbf{r}^{l}}} \tag{4.33}$$

Cancelling all intra terms:

$$= -\frac{1}{\beta}\ln\frac{\int\exp\left[-\beta U^{inter}_{4}(\langle\mathbf{r}\rangle_{\mathrm{Q}}\,,\mathbf{q}))\right]d\mathbf{q}}{\int\exp\left[-\beta U^{inter}_{3}(\langle\mathbf{r}\rangle_{\mathrm{P}}\,,\mathbf{q}))\right]d\mathbf{q}} \tag{4.34}$$

and:

$$\begin{aligned}&\Delta\Delta A^{inter}_{hyd}\\ \approx&\Delta\Delta A^{inter(sep)}_{hyd}\end{aligned} \tag{4.35}$$

Introducing $\langle\mathbf{r}\rangle$ in the place of $\mathbf{r}$:

$$= -\frac{1}{\beta}\ln\frac{\int\exp\left[-\beta(U^{l}_{\mathrm{P}}(\mathbf{r}^{l})+U^{s}(\mathbf{r}^{s})+U^{inter}_{4}(\langle\mathbf{r}\rangle_{\mathrm{P}}\,,\mathbf{q}))\right]d\mathbf{r}^{l}d\mathbf{r}^{s}d\mathbf{q}}{\int\exp\left[-\beta(U^{l}_{\mathrm{P}}(\mathbf{r}^{l})+U^{s}(\mathbf{r}^{s})+U^{inter}_{3}(\langle\mathbf{r}\rangle_{\mathrm{P}}\,,\mathbf{q}))\right]d\mathbf{r}^{l}d\mathbf{r}^{s}d\mathbf{q}} \tag{4.36}$$

Separating the integrals:

$$= -\frac{1}{\beta}\ln\frac{\int\exp\left[-\beta U^{l}_{\mathrm{P}}(\mathbf{r}^{l})\right]d\mathbf{r}^{l}\int\exp\left[-\beta U^{s}(\mathbf{r}^{s})\right]d\mathbf{r}^{s}\int\exp\left[-\beta U^{inter}_{4}(\langle\mathbf{r}\rangle_{\mathrm{P}}\,,\mathbf{q})\right]d\mathbf{q}}{\int\exp\left[-\beta U^{l}_{\mathrm{P}}(\mathbf{r}^{l})\right]d\mathbf{r}^{l}\int\exp\left[-\beta U^{s}(\mathbf{r}^{s})\right]d\mathbf{r}^{s}\int\exp\left[-\beta U^{inter}_{3}(\langle\mathbf{r}\rangle_{\mathrm{P}}\,,\mathbf{q})\right]d\mathbf{q}} \tag{4.37}$$

Cancelling all intra terms:

$$= -\frac{1}{\beta}\ln\frac{\int\exp\left[-\beta U^{inter}_{4}(\langle\mathbf{r}\rangle_{\mathrm{P}},\mathbf{q}))\right]d\mathbf{q}}{\int\exp\left[-\beta U^{inter}_{3}(\langle\mathbf{r}\rangle_{\mathrm{P}},\mathbf{q}))\right]d\mathbf{q}} \tag{4.38}$$

where the notation $\langle\mathbf{r}\rangle$ is taken to denote the equilibrium values of the intramolecular degrees of freedom. Each of these has a subscript denoting the thermodynamic state from which the averages are taken and dictated by the intramolecular terms within each integral. The above represents an approach for writing separated approximations to integrals over multiple variables[95]. By including the averaged value of $\mathbf{r}$ it is removed as an explicit degree of freedom that must be integrated over, from here separation of the integral into its components follows naturally.

The obvious question is under what conditions do equations (4.34) and (4.38) well approximate their counterparts. It is above noted that each of the integrals within equations (4.29) and (4.30) corresponds to a partition function which can be related to the probability distribution of a thermodynamic state by equation (2.6). The full expressions are well approximate by their separations where the probability distributions of the system degrees of freedom are unchanged in isolation. This is tantamount to the requirement that the probability distributions of the intramolecular and interaction degrees of freedom are independent of one another i.e. uncorrelated. For example, consider the separation:

$$\begin{aligned}\int \exp\left[-\beta(U^l_{\mathrm{Q}}(\mathbf{r}^l) + U^s(\mathbf{r}^s) + U_4^{inter}(\langle\mathbf{r}\rangle_{\mathrm{Q}}, \mathbf{q}))\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} \\ \approx \int \exp\left[-\beta U^l_{\mathrm{Q}}(\mathbf{r}^l)\right] d\mathbf{r}^l \int \exp\left[-\beta U^s(\mathbf{r}^s)\right] d\mathbf{r}^s \int \exp\left[-\beta U_4^{inter}(\langle\mathbf{r}\rangle_{\mathrm{Q}}, \mathbf{q}))\right] d\mathbf{q}\end{aligned} \tag{4.39}$$

within the full integral, the probability distribution of the intramolecular degrees of freedom $\mathbf{r}$ is altered by the presence of the U_4^{inter}. The separated expression treats the intramolecular terms in isolation however, allowing cancellation with the gas phase term. The more closely the probability distribution of $\mathbf{r}$ matches between the isolated terms and the full interacting system the better the approximation given by the separated integral expression. This same condition applies in reverse, considering the effect of the intramolecular energies on the interaction degree of freedom, with the exception that this time we are simplifying the integral by freezing the intramolecular terms at their averaged values.

In the limit of $\Delta\Delta A^{(sep)}_{hyd}$ and $\Delta\Delta A^{inter(sep)}_{hyd}$, providing sufficiently accurate approximations to equations (4.29) and (4.30), the use of interaction energies in free energy calculations may be considered comparable to total energies, with a discrepancy given by:

$$\Delta\Delta A^{(sep)}_{diff} = \Delta\Delta A^{(sep)}_{hyd} - \Delta\Delta A^{inter(sep)}_{hyd} = -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta U_4^{inter}(\langle\mathbf{r}\rangle_{\mathrm{Q}}, \mathbf{q}))\right] d\mathbf{q}}{\int \exp\left[-\beta U_4^{inter}(\langle\mathbf{r}\rangle_{\mathrm{P}}, \mathbf{q}))\right] d\mathbf{q}} \tag{4.40}$$

interpretation of equation (4.40) requires some care. The difference lies in the equilibrium values used for the intramolecular degrees of freedom in the different states of the cycle. Under what conditions then might this term be expected to be negligible? The degrees of freedom subsumed in $\mathbf{r}$ are $\mathbf{r}^l$ and $\mathbf{r}^s$. The majority of values in $\langle\mathbf{r}\rangle_{\mathrm{Q}}$ will be the same as in $\langle\mathbf{r}\rangle_{\mathrm{P}}$ excepting those ligand degrees of freedom that are mutated. If the mutated terms lead to significant changes in internal geometry however this may be expected to generate contributions to equation (4.40). The inclusion of the solvent intramolecular degrees of freedom meanwhile suggest that changes in the solvent structure caused by the perturbations in the ligand may also contribute to this term.

Having noted the lack of rigour inherited through the use of interaction energies, questions are raised about properties of interaction free energy differences. In particular, are the free energy differences calculated with interaction free energies pathway independent? We address this issue through the introduction of an additional λ window into the above derivations. For total energies:

$$\Delta\Delta A_{hyd}^{2\lambda} = (\Delta A_{3\rightarrow 3.5} + \Delta A_{3.5\rightarrow 4}) - (\Delta A_{1\rightarrow 1.5} + \Delta A_{1.5\rightarrow 2}) \tag{4.41}$$

$$= -\frac{1}{\beta}\ln\frac{\frac{Z_{3.5}}{Z_3}}{\frac{Z_{1.5}}{Z_1}} - \frac{1}{\beta}\ln\frac{\frac{Z_4}{Z_{3.5}}}{\frac{Z_2}{Z_{1.5}}} \tag{4.42}$$

$$= -\frac{1}{\beta}\ln\frac{\frac{Z_4}{Z_3}}{\frac{Z_2}{Z_1}} \tag{4.43}$$

$$= \Delta\Delta A_{hyd} \tag{4.44}$$

for interaction energies:

$$\Delta\Delta A_{hyd}^{inter,2\lambda} \tag{4.45}$$

$$=\Delta A_{3\rightarrow 3.5}^{inter} + \Delta A_{3.5\rightarrow 4.0}^{inter} \tag{4.46}$$

$$=-\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_3^{intra} + U_{3.5}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta U_3\right] d\mathbf{r}d\mathbf{q}} - \frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_{3.5}^{intra} + U_4^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta U_{3.5}\right] d\mathbf{r}d\mathbf{q}} \tag{4.47}$$

$$\neq\Delta\Delta A_{hyd}^{inter} \tag{4.48}$$

where $\Delta\Delta A_{hyd}^{2\lambda}$ and $\Delta\Delta A_{hyd}^{inter,2\lambda}$ are the total and interaction free energies of hydration calculated using 2 λ windows. As expected $\Delta\Delta A_{hyd}$ is invariant with the number of λ windows, as it creates a smooth succession of states i.e. a perturbation from 3 to 3.5 and then from 3.5 to 4. However this is not the case with the use of interaction energies. As discussed above, $\Delta\Delta A_{hyd}^{inter}$ can be understood as the free energy difference between state 3 and a hybrid of states 3 and 4. The introduction of a λ-window however does not create a connection of states. Rather we have a perturbation from state 3 to a hybrid intermediate of states 3 and 3.5, followed by a perturbation from 3.5 to a hybrid of states 3.5 and 4.

4.2.1. Discussion

In the preceding analysis we have shown that the use of interaction energies, although common place, lacks the theoretical correctness associated with total energies. This is a prominent disadvantage as one of the noted strengths of free energy techniques is their rigorous theoretical derivation. The use of interaction energies adds an additional degree of approximation, although the magnitudes of such errors as will be introduced in calculations with systems of biological complexity are not readily apparent. In this context we have reached the limits of the purely theory-based analysis as presented here and must move on to consider what can be learned from calculations performed in practice. This is carried out in the form of a model thermodynamic system in section 4.3, a full protein-ligand system in section 4.4 and an extension of the model system to consider polarisation in section 4.5. First, however, some additional observations are made on the basis of theory alone.

Despite the loss of rigour it engenders, the interaction energy expression 4.18, has advantages with regard to the use of total energies. As previously noted the resultant simplification of the free energy cycles to exclude a vacuum calculation provides some computational savings. It might be argued that this is minimal, as simulations in vacuum are comparatively cheap, but this may not always be the case. Consider the free energy of hydration of a MM/QM system where the QM ligand is embedded in MM solvent. Here the MM component of the system contributes negligibly to the computational requirements in the aqueous phase of the calculation. This entails that the vacuum simulation of the ligand at the QM level will therefore require almost exactly as much CPU time to complete. Differing computational requirements of the gas and aqueous phase will be more prominent as the solvent is modelled more intensively. However in this case the use of interaction energies reduces the cost of the computation by approximately half.

A second benefit is also suggested by the separated equations (4.34) and (4.38). The integral in these terms is reduced to cover only the interaction degrees of freedom. The number of interaction degrees of freedom are far fewer than the number of intramolecular degrees of freedom. Reducing the dimensionality of the integrals suggests far superior convergence characteristics for the interaction energy expressions and can be viewed in terms of reducing the requirement for phase space overlap between states of the perturbation. As only the interaction degrees of freedom contribute directly to the free energy difference, only overlap of the relevant sub-volume of the total phase space is required. This aspect of the use of interaction energies is demonstrated in detail in chapter 3.

A valuable distinction can be made between the use of interaction energies in the context of single and dual topology calculations. The alternative λ-coupling approach used in dual topology calculations is given by equation (2.109). Use of this with any of the free energy difference estimators leads to the cancellation of intramolecular energy terms, leaving only the interaction energies. The work of this section has implicitely assumed the use of single-topology calculations, through the form of the potential energy functions specified. Although superficially single topology with interaction energies and dual topology approaches appear similar, it is shown in appendix C that interaction energies remain rigorously correct in the context of dual topology. Despite being unsuitable for single step MM to QM perturbations with the Zwanzig equation, when combined with some of the alternative approaches considered in section 2.5, dual topology calculations may provide a way to circumvent difficulties associated with the use of interaction energies. This is discussed in more detail in chapter 5.

There are notable points of comparison warranted between our theoretical analysis and the work of Boresch[95,97,98] related to the treatment of bonded terms within free energy calculations. Boresch's treatment is limited to considering the decomposition of the free energy difference of the ligand degrees of freedom into three different contributions[95] and uses this to conclude the equivalency of the single and dual topology free energy calculations. This does not consider the effect of interaction potentials on the intramolecular degrees of freedom however as their independence is simply assumed. Our analysis therefore builds on the work of Boresch and considers situations in which the influence of interaction degrees of freedom violates the assumed independence.

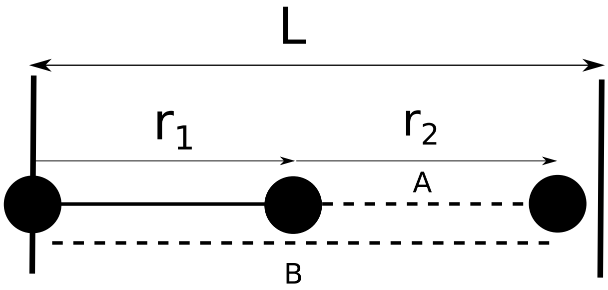

Figure 4.2.: Interaction energy model system. The covalent bond is shown as a solid line whilst the dashed lines denotes Lennard-Jones potentials. Marked are the two bond lengths r_1 and r_2 that constitute a complete internal coordinate system.

4.3. Model System

Having proposed the theoretical results of the previous section we proceed by exploring the implications of the derived terms in the context of a model thermodynamic system. The system under consideration is shown in figure 4.2. This meets the criteria of having a well defined interaction energy term but is sufficiently simple to allow precise solution, through numerical evaluation, of its partition function. As shown in figure 4.2 the system is composed of three atoms and three potentials in a one dimensional non-periodic box of length L. Atom 1 is kept stationary against the left wall of the box leaving a system with 2 degrees of freedom. The configuration of the atoms is also maintained such that $r_2 \geq 0$, i.e. the third atom is not able to pass through the central atom and maintains its position to the right. The potential energy of this system is given by:

$$U^{model}(r_1, r_2) = Hm(r_1) + LJ^A(r_2) + LJ^B(r_1 + r_2) \tag{4.49}$$

$$\begin{aligned} &= k(b^0 - r_1)^2 + \epsilon^A \left[\left(\frac{\sigma^A}{r_2} \right)^{12} - 2 \left(\frac{\sigma^A}{r_2} \right)^6 \right] \\ &+ \epsilon^B \left[\left(\frac{\sigma^B}{r_1 + r_2} \right)^{12} - 2 \left(\frac{\sigma^B}{r_1 + r_2} \right)^6 \right] \end{aligned} \tag{4.50}$$

where the systems potential energy U^{model} is the sum of a harmonic potential term, Hm, and two Lennard-Jones terms, LJ. The LJ^A and LJ^B potentials correspond to those labelled in figure 4.2. For the harmonic potential the parameters k and b^0, specify the spring constant and equilibrium bond length respectively, whilst for the LJ potentials ϵ and σ give the well-depth and the displacement of the potential's minimum. Hm is considered to comprise the intramolecular energy of the system whilst the Lennard-Jones terms constitute the interaction energy. We can use this model system to mimic the relative free energy of hydration calculation from the previous section. Using equation (4.49) to give the stages of the cycle equivalent to equations (4.19) to (4.22):

$$
\begin{aligned}
U_1^{model}(r_1) =& k_P(b_P^0 - r_1)^2 && (4.51)\\
U_2^{model}(r_1) =& k_Q(b_Q^0 - r_1)^2 && (4.52)\\
U_3^{model}(r_1, r_2) =& k_P(b_P^0 - r_1)^2 + \epsilon_3^A \left[\left(\frac{\sigma_3^A}{r_2} \right)^{12} - 2 \left(\frac{\sigma_3^A}{r_2} \right)^6 \right] \\
&+ \epsilon_3^B \left[\left(\frac{\sigma_3^B}{r_1 + r_2} \right)^{12} - 2 \left(\frac{\sigma_3^B}{r_1 + r_2} \right)^6 \right] && (4.53)\\
U_4^{model}(r_1, r_2) =& k_Q(b_Q^0 - r_1)^2 + \epsilon_4^A \left[\left(\frac{\sigma_4^A}{r_2} \right)^{12} - 2 \left(\frac{\sigma_4^A}{r_2} \right)^6 \right] \\
&+ \epsilon_4^B \left[\left(\frac{\sigma_4^B}{r_1 + r_2} \right)^{12} - 2 \left(\frac{\sigma_4^B}{r_1 + r_2} \right)^6 \right] && (4.54)
\end{aligned}
$$

gas phase conditions are simulated by setting $\epsilon^A = \epsilon^B = 0$ kcal.mol^{-1}. Adjusting the parameters of calculations made with the model system, different scenarios can be designed that assess the various properties of the interaction energy approximation observed in the previous section. First is the question of how well equations (4.23) and (4.24) are approximated by equations (4.34) and (4.38). This can be directly assessed through the control of LJ^B, as this term (being a function of both r_1 and r_2) precludes the separability of the integrals within the partition function of the system. When $\epsilon_3^B = \epsilon_4^B = 0.0$ the integrals are formally separable and equations (4.23) and (4.24) equate closely with equations (4.34) and (4.38), through cancellation of the intramolecular terms. By increasing the strength of LJ^B therefore the quality of the approximation can be controlled. Second, is the contribution of the equation (4.40), giving the difference between the equations (4.34) and (4.38). This is noted to relate to the averaged geometries of the intramolecular term and hence can be examined through altering the value of b^0 between different states of the cycle.

4.3.1. Solution of Model System Equations

We here give details for exact solution of the equations of the model system, including the limits of the integrals and application of the separated expressions to this system.

From equation (4.29):

$$
\Delta\Delta A_{hyd} = -\frac{1}{\beta} \ln \frac{\dfrac{\int_0^L \int_0^{L-r_1} \exp\left[-\beta U_4^{model}(r_1, r_2)\right] dr_2 dr_1}{\int_0^L \int_0^{L-r_1} \exp\left[-\beta U_3^{model}(r_1, r_2)\right] dr_2 dr_1}}{\dfrac{\int_0^L \exp\left[-\beta U_2^{model}(r_1)\right] dr_1}{\int_0^L \exp\left[-\beta U_1^{model}(r_1)\right] dr_1}} \tag{4.55}
$$

From equation (4.30):

$$\Delta\Delta A_{hyd}^{inter} = -\frac{1}{\beta}\ln\frac{\int_0^L\int_0^{L-r_1}\exp\left[-\beta Hm_3(r_1)+LJ_4^A(r_1)+LJ_4^B(r_1+r_2)\right]dr_2dr_1}{\int_0^L\int_0^{L-r_1}\exp\left[-\beta U_4^{model}\right]dr_2dr_1} \tag{4.56}$$

From equation (4.34):

$$\Delta\Delta A_{hyd}^{(sep)} = -\frac{1}{\beta}\ln\frac{\int_0^{L-b_{\mathrm{Q}}^0}\exp\left[-\beta LJ_4^A(r_1)+LJ_4^B(b_{\mathrm{Q}}^0+r_2)\right]dr_2}{\int_0^{L-b_{\mathrm{P}}^0}\exp\left[-\beta LJ_3^A(r_1)+LJ_3^B(b_{\mathrm{P}}^0+r_2)\right]dr_2} \tag{4.57}$$

From equation (4.38):

$$\Delta\Delta A_{hyd}^{inter(sep)} = -\frac{1}{\beta}\ln\frac{\int_0^{L-b_{\mathrm{P}}^0}\exp\left[-\beta LJ_4^A(r_1)+LJ_4^B(b_{\mathrm{P}}^0+r_2)\right]dr_2}{\int_0^{L-b_{\mathrm{P}}^0}\exp\left[-\beta LJ_3^A(r_1)+LJ_3^B(b_{\mathrm{P}}^0+r_2)\right]dr_2} \tag{4.58}$$

From equation (4.40):

$$\Delta\Delta A_{diff}^{(sep)} = -\frac{1}{\beta}\ln\frac{\int_0^{L-b_{\mathrm{Q}}^0}\exp\left[-\beta LJ_4^A(r_1)+LJ_4^B(b_{\mathrm{P}}^0+r_2)\right]dr_2}{\int_0^{L-b_{\mathrm{P}}^0}\exp\left[-\beta LJ_4^A(r_1)+LJ_4^B(b_{\mathrm{P}}^0+r_2)\right]dr_2} \tag{4.59}$$

Integrals are given over the internal coordinate system (r_1,r_2). As mentioned previously, a lower bound of 0 is enforced for r_2 whilst its upper bound is limited by the size of the system, L, and the current value of r_1. All calculations in this section assume the same temperature of 300 K. All integrals were solved using the adaptive Gaussian quadrature routines (see appendix A) of the SciPy library, which provides wrappers to Fortran QUADPACK routines. Integrals were solved to a relative tolerance of 1.5×10^{-8}.

In addition to these analytical expressions, the total and interaction free energy terms $\Delta\Delta A_{hyd}$, $\Delta\Delta A_{hyd}^{inter}$ are evaluated using Monte Carlo sampling and the Zwanzig equation (equation (2.80)). This is used to provide independent support for the derived results and their application within the model system, as well as to act as verification of the numerical integration routines. For all results within this section, equivalent MC results are given in appendix E, showing excellent agreement, generally to within 1 σ.

It may be noted, that the resulting equations fail to exactly correspond to those of the previous section, where the degrees of freedom are divided into three categories. Restricted to two degrees of freedom the model system forgoes inclusion of a separate solvent intramolecular degree of freedom. As the influence of the ligand degrees of freedom on the interaction energy is not qualitatively distinct from that of the solvent intramolecular degrees of freedom we do not consider that this limits this interpretive analysis.

4.3.2. Integral Separation

The first set of scenarios considered are given in table 4.1 and figure 4.3. Here we have varied two parameters of the model pertaining to LJ^B. This term controls the degree of separability of the integrals given by equations (4.55) and (4.56). When $\epsilon^B = 0$ the intramolecular term Hm and the interaction term LJ^A are independent of one another. However, as LJ^B increases r_1 and r_2 become more correlated. Both of the parameters of LJ^B were varied, three values of σ^B were used, along with 15 values of ϵ^B. Equations (4.55) to (4.57) (under this set of parameters equation (4.57) and equation (4.58) are equivalent as $b^0_{\mathrm{P}} = b^0_{\mathrm{Q}}$) were solved for each set of the parameters. For simplicity the two varied parameters are held to the same values for states 3 and 4 of the cycle. Figures 4.3a to 4.3c show a number of interesting properties. First, they provide support for the interpretation of integral separability as the limit in which $\Delta\Delta A_{hyd}$ approaches equivalence with $\Delta\Delta A^{inter}_{hyd}$. Where the values of these terms coincide, they also agree well with the separated limit expression, $\Delta\Delta A^{sep}_{hyd}$, as enforced by small values of ϵ^B.

Secondly, the error of $\Delta\Delta A^{inter}_{hyd}$ with respect to $\Delta\Delta A_{hyd}$ is shown to monotonically increase with the parameter ϵ^B, (figure 4.3d). Thus, increasing the strength of the correlated term LJ^B leads to interaction energies acting as an increasingly poor proxy for total energies. This effect is particularly pronounced where $\sigma^B = 1.7$, much larger errors being observed in this case. As noted in section 4.2 the separated expressions form good approximations where the probability distributions of the intramolecular and interaction degrees of freedom are not strongly influenced by one another. This will allow the cancellation of the isolated intramolecular terms in $\Delta\Delta A^{(sep)}_{hyd}$ to approximate the full and interaction energy expressions. In the case of the model system where $\sigma^B < b^0 + \sigma^A$ the effect of LJ^B is to crowd Hm to occupy shorter bonds lengths than it would typically adopt in isolation. As LJ^B is made stronger Hm is crowded further and the interaction energy result deviates further from the total energy result.

This observation is supported through a second set of simulations, with parameters given in table 4.2. This is based on the collection of parameters from the previous set of simulations that give an absolute error of around 0.05 kcal.mol^{-1}. Here we introduce the parameter k_m that is used a modifier to the parameter k, such that $k_{\mathrm{P}} = 50 + k_m$ and $k_{\mathrm{Q}} = 100 + k_m$. This formulation allows the stiffness of the Hm term to be increased whilst keeping the size of the perturbation between states fairly similar. $\Delta\Delta A_{hyd}$, $\Delta\Delta A^{inter}_{hyd}$ and $\Delta\Delta A^{(sep)}_{hyd}$ for a range of values of k_m are given in figure 4.4. The value of all three terms are shown to converge as the value of k_m increases, the stiffer harmonic bonds being better able to resist distortion by the

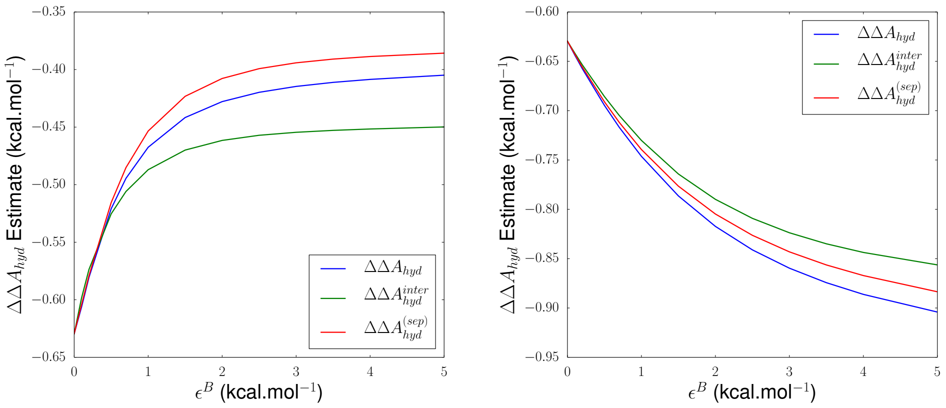

(a) $\sigma^B = 2.3$

(b) $\sigma^B = 2.0$

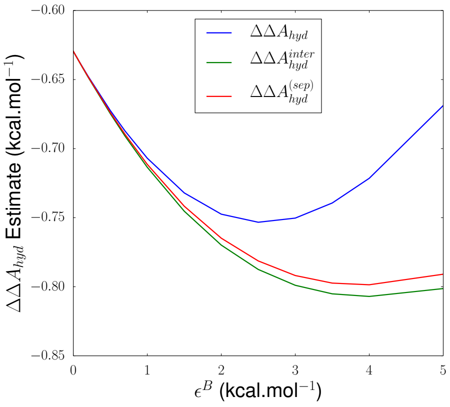

(c) $\sigma^B = 1.7$

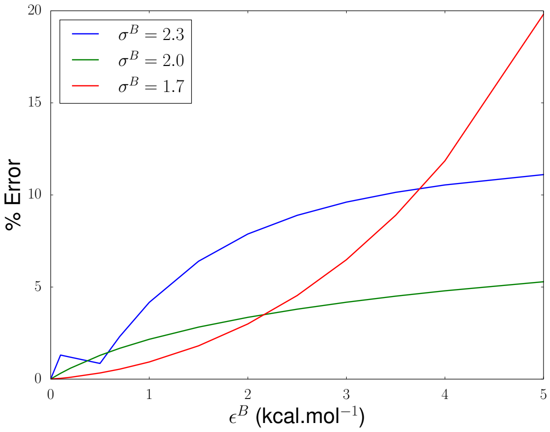

(d) Interaction energy error

Figure 4.3.: Results of model system free energy of hydration calculations using parameters from table 4.1.

Parameter	Value
k_{P}	50 kcal.mol^{-1}.Å^{-2}
k_{Q}	100 kcal.mol^{-1}.Å^{-2}
b^0_{P},b^0_{Q}	1.0 Å
ϵ^A_3	1.0 kcal.mol^{-1}
ϵ^A_4	2.0 kcal.mol^{-1}
ϵ^B_3,ϵ^B_4	var kcal.mol^{-1}
σ^A_3,σ^A_4	1.0 Å
σ^B_3,σ^B_4	var Å
L	3.0

Table 4.1.: Parameter sets used in the calculations with the model system for figure 4.3.

Lennard-Jones terms. $\Delta\Delta A^{(sep)}_{hyd}$ is in fact invariant with respect to the values of k_{P} and k_{Q} as it only takes into account the values of b^0.

Interestingly, the deviation of the average of the bond length r_1, under MC sampling of state 3, from b^0 shows a linear correlation with the difference between $\Delta\Delta A_{hyd}$, $\Delta\Delta A^{inter}_{hyd}$. This offers direct support that alterations in the probability distribution of r_1 contribute to error in the use of interaction energies. We note that the range of spring constants considered cover the typical range of harmonic bond strengths within force field models. However, in this model system the harmonic term is intended to stand in for the various potentials used to describe the intramolecular energy of biomolecules, including softer bond and dihedral potentials.

4.3.3. Difference in Separated Terms

A third set of calculations using the values in table 4.3 were used to explore the implications of equation (4.40), results are shown in figure 4.5. By choosing $\epsilon^B = 0$, we ensure that the integral expressions are equivalent to their separated counterparts. For this set of calculations therefore $\Delta\Delta A_{hyd} = \Delta\Delta A^{(sep)}_{hyd}$ (shown by figure 4.5b) and $\Delta\Delta A^{inter}_{hyd} = \Delta\Delta A^{inter(sep)}_{hyd}$, the difference between $\Delta\Delta A_{hyd}$ and $\Delta\Delta A^{inter}_{hyd}$ being precisely captured by equation (4.40). The interaction energy results (figure 4.5a) are again invariant with respect to changes in the system intramolecular parameters however, we note that changing the bond length b^0_{Q} does not change the free energy contribution of the intramolecular term itself. These contributions cancel precisely in both the interaction and total energy case. The difference given by equation (4.40) instead arises due to the changes in the average value of r_1 between states, and the effect this has on the distribution of r_2. These data demonstrate that even in the limit of complete separability there is an independent source of error arising in the use of interaction energies, caused by differences in intramolecular geometry between states.

The results presented here support the conclusions from the theoretical analysis that

Parameter	Value
k_P	$50 + k_m$ kcal.mol^{-1}.Å^{-2}
k_Q	$100 + k_m$ kcal.mol^{-1}.Å^{-2}
b^0_P,b^0_Q	1.0 Å
ϵ^A_3	1.0 kcal.mol^{-1}
ϵ^A_4	2.0 kcal.mol^{-1}
ϵ^B_3,ϵ^B_4	3.0 kcal.mol^{-1}
σ^A_3,σ^A_4	1.0 Å
σ^B_3,σ^B_4	1.7 Å
L	3.0

Table 4.2.: Parameter sets used in the calculations with the model system for figure 4.4.

Parameter	Value
k_P	50 kcal.mol^{-1}.Å^{-2}
k_Q	100 kcal.mol^{-1}.Å^{-2}
b^0_P	1.0 Å
b^0_Q	var Å
ϵ^A_3	1.0 kcal.mol^{-1}
ϵ^A_4	2.0 kcal.mol^{-1}
ϵ^B_3,ϵ^B_4	0.0 kcal.mol^{-1}
σ^A_3,σ^A_4	1.0 Å
σ^B_3,σ^B_4	2.0 Å
L	3.0

Table 4.3.: Parameter sets used in the calculations with the model system for figure 4.5.

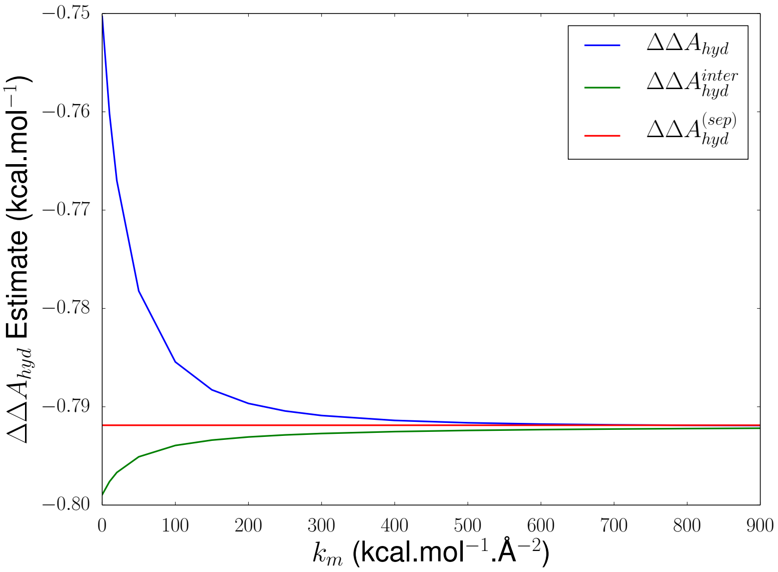

(a) Free energy of hydration with varying strengths of harmonic bond.

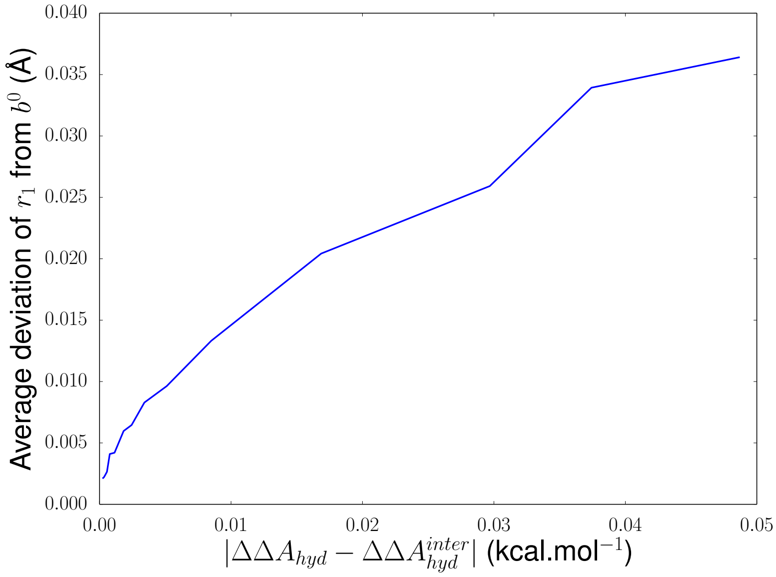

(b) Average values of r_1 derived from MC sampling of state 3.

Figure 4.4.: Results of model system free energy of hydration calculations using parameters from table 4.2.

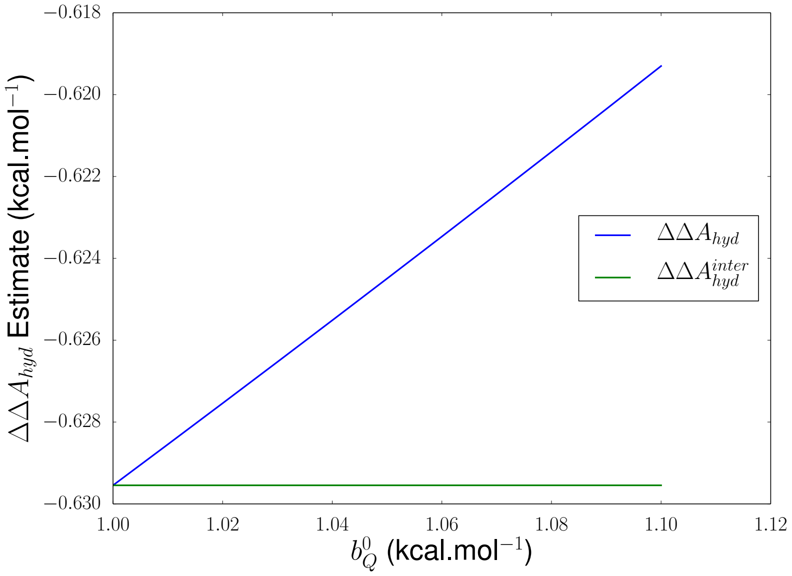

(a) Varying free energy differences with respect to changes in the harmonic bond length of states 2 and 4.

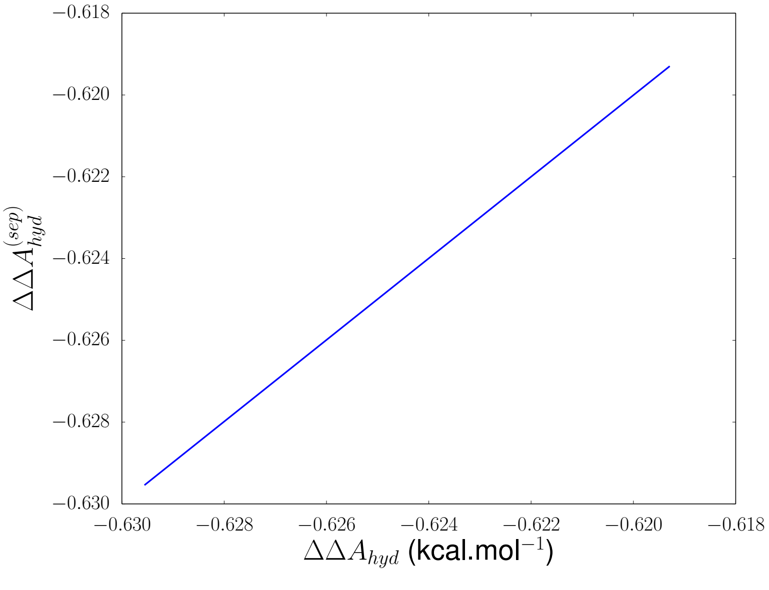

(b) Equivalence of total free energy difference and its separated counterpart. Data shown are from calculations of figure 4.5a.

Figure 4.5.: Results of model system free energy of hydration calculations using parameters from table 4.3.

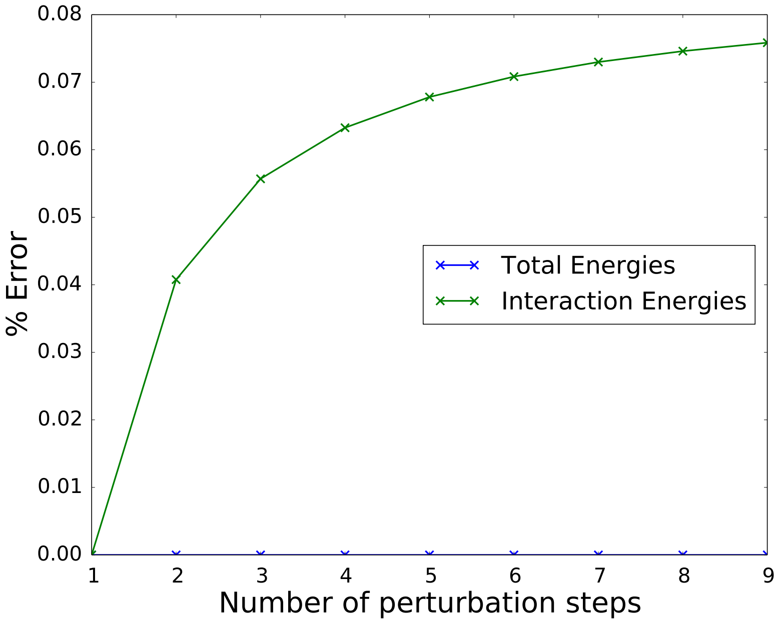

Figure 4.6.: Difference in free energies using multiple λ-windows. All perturbations using more than a single step employ additional λ-windows. Percentage errors are given against the total free energy difference using 1 step.

we have identified the contributions to the difference between the total and interaction free energy results. Questions remain as to what extent the issues highlighted come in to play in the context of more realistic biomolecular free energy calculations.

4.3.4. Effect of λ-windows

Further to the analysis of section 4.2, we here demonstrate the effect of including additional λ-windows in the calculation of interaction free energy differences, shown in figure 4.6. Calculations were carried out using parameters from table 4.1, with $\epsilon_3^B = \epsilon_4^B = 5.0$ kcal.mol^{-1} and $\sigma_3^B = \sigma_4^B = 1.7$ Å. As expected total energy results do not vary with the number of λ-windows, but interaction energy results do. Although the effects demonstrated here are comparitively small compared to the those reported previously, it is sufficient to demonstrate in principle the path-dependence of interaction energies.

4.4. Cyclooxygenase 2

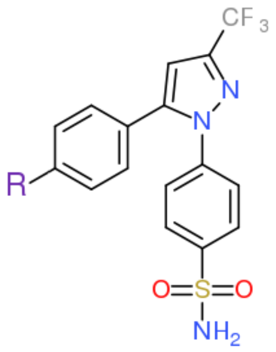

(a) The ligand scaffold for the set of congeneric COX2 inhibitors[99]

R_1	CH_3
R_2	CH_2CH_3
R_4	SCH_3
R_5	OCH_3
R_6	CF_3
R_7	OH

(b) Different R groups for COX2 ligands

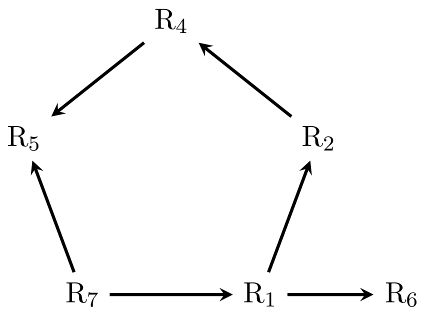

(c) The different R groups used in free energy calculations arranged in a perturbation cycle. Arrows indicate the direction in which free energy differences are given.

Figure 4.7.: Cyclooxygenase 2 Ligands

The results presented in section 4.3 are qualitatively enlightening but their implication for calculations in more realistic systems is not clear. For this reason we have considered the relative free energies of binding and hydration of a set of congeneric inhibitors of cyclooxygenase 2 (COX2). This system was chosen as it has been the subject of previous calculations within this research group[100].

4.4.1. System Setup and Calculations

Calculations were performed using ProtoMS 2.3[76] with simulation inputs prepared for a previous study[100], with a slightly altered simulation protocol. Ligands were prepared using the ANTECHAMBER module of AMBER 8 to calculate AM1-bcc charges for each atom. These were subsequently used to construct template files for use with ProtoMS. Aqueous phase simulations were prepared by adding a sphere of TIP4P waters with a radius of 22 Å around the centre of geometry of the ligand.

The protein structure for the bound-phase simulations was derived from the pdb structure 1CX2. Hydrogen atoms for this structure are already given within the pdb database entry, although the protonation state of all histidine residues was determined by visual inspection (see Michel *et al.* for details[100]). A ‘scoop’ of the protein was prepared such that only residues with a heavy atom with 15 Å of a heavy atom of ligand 2 were retained, leaving a structure of only 155 residues. The protein scoop complexed with ligand 2 was then energy minimised using the SANDER module of AMBER 8. Some minor additional details are given in Michel *et al.*[100].

Free energy differences between different R groups were calculated using FDTI and Hamiltonian replica exchange (see section 2.4). λ-windows between 0.0 and 0.9 were evenly spaced by 0.1, with an additional 2 windows at 0.95 and 1.0 to give a total of 12. The trapezium rule was used to give the final free energy differences between states. Perturbations were conducted using the single topology approach. Dummy atoms were withdrawn within the van der Waals radii of neighbouring atoms, to a bond length of 0.2 Å and were not sampled at the appropriate end point of the perturbation. Boundary conditions for solvent and bound phase simulations were non-periodic. A half-harmonic potential was applied to all solvent molecules more than 22 Å away from the geometric centre of the ligand to prevent diffusion.

Simulation protocols are given as follows. 20 million Metropolis-Hasting MC moves of the solvent alone were performed to remove clashes with protein and/or ligand. For each individual perturbation a further separate equilibration of 20 million MC moves was conducted, using the system with the largest ligand of the perturbation and with moves for all system components attempted. Production simulations were then conducted with a further 10 million moves of equilibration to allow the system to accommodate to the corresponding geometry of the ligand at each λ-window, statistics for TI were then collected for 30 million moves. Ten independent repeats with different random number seeds were carried out for each perturbation. MC moves for different system components were attempted with the following probabilities. In the aqueous phase - solvent 98.4%, ligand 1.6%. In the bound phase - solvent 85.8%, protein 12.8% and ligand 1.4%. Bond length moves for the entire system along with bond and torsion angle moves within rings were not attempted and these degrees of freedom were hence constrained to their initial simulation values. The primary difference between the work presented here and the simulation protocol used previously is that MC moves for all protein sidechain and backbone atoms were attempted. In contrast, Michel *et al.*[100] kept protein backbones rigid and only attempted moves for sidechains with a heavy atom within 10 Å of the ligand.

4.4.2. Results

Results for the relative free energies of hydration and binding calculations are given in Tables 4.4 and 4.5. Using total energies these are in good agreement with those reported by Michel *et al.* using total energies. Differences between results arising through the use of interaction energies and total energies are shown in figure 4.8

Perturbation	$\Delta A_{(g)}$	$\Delta A_{(aq)}$	$\Delta\Delta A_{hyd}$	$\Delta\Delta A_{hyd}^{inter}$
1→2	0.66 ± 0.00	1.18 ± 0.03	0.52 ± 0.03	0.48 ± 0.02
2→4	-1.09 ± 0.00	-1.61 ± 0.01	-0.52 ± 0.01	-0.36 ± 0.01
4→5	-3.43 ± 0.00	-4.93 ± 0.01	-1.50 ± 0.01	-1.66 ± 0.01
7→5	8.96 ± 0.00	12.33 ± 0.04	3.37 ± 0.04	2.93 ± 0.04
7→1	13.54 ± 0.00	18.25 ± 0.01	4.71 ± 0.01	4.69 ± 0.01
1→6	16.35 ± 0.00	16.92 ± 0.01	0.57 ± 0.01	0.45 ± 0.01

Table 4.4.: Results for relative free energy of hydration calculations using COX2 ligands. All values are given in kcal.mol^{-1}.

for both hydration and binding calculations. Generally it can be seen than the error introduced by the use of interaction energies is non-zero. The solvation calculations appear far more robust, the greatest discrepancy being 0.44 kcal.mol^{-1} in the perturbation from group R_7 to R_5, and the mean unsigned average being only 0.16 kcal.mol^{-1}. In the case of the binding free energies however notably larger errors are observed. The two largest being the R_2 to R_4 and R_1 to R_2 perturbations at 0.73 kcal.mol^{-1} and 0.76 kcal.mol^{-1} respectively. The mean unsigned error for binding is also significantly inflated at 0.44 kcal.mol^{-1}.

All perturbations appear well converged, giving standard errors less than 0.1. Cycle closures can be calculated for the loop given by the ligands R_1, R_2, R_4, R_5 and R_7. These are given in table 4.6 for each of calculations performed. Despite the small standard errors given for individual perturbations, the closures for certain cycles are relatively poor, in particular the total energy results in the gas and aqueous phases. Standard errors reported between independent repeats are small and are reported for ten repeats from the same initial configuration. Poor closure of the cycles suggests that the standard errors are overestimating the convergence of these results. As expected the interaction energy results show much better cycle closure, suggesting these values are better converged.

4.4.3. Discussion

Poor convergence of the total energy results makes it difficult to draw any firm conclusions from these data. This said, the difference between results reported for total and interaction energies differ markedly for individual perturbations. Sizeable errors of greater than 0.5 kcal.mol^{-1} arise in three of the five perturbations. A notable improvement is observed for the closure of the thermodynamic cycle in figure 4.7c, for calculations using interaction energies versus those with total energies (table 4.6). This may be rationalised by the observation made in the previous section regarding the decreased requirement for phase space overlap provided by the use of interaction energies.

No particular trends can be observed in those perturbations associated with large

Perturbation	$\Delta A_{(aq)}$	$\Delta A_{(b)}$	$\Delta\Delta A_{bind}$
1→2	1.18 ± 0.03	2.97 ± 0.02	1.78 ± 0.03
2→4	-1.61 ± 0.01	-3.69 ± 0.01	-2.08 ± 0.02
4→5	-4.93 ± 0.01	-5.46 ± 0.03	-0.53 ± 0.03
7→5	12.33 ± 0.04	10.26 ± 0.04	-2.06 ± 0.06
7→1	18.25 ± 0.01	16.18 ± 0.02	-2.06 ± 0.02
1→6	16.92 ± 0.01	21.18 ± 0.02	4.27 ± 0.02

(a) Results for total energy relative free energy of binding calculations using COX2 ligands. $\Delta A_{(b)}$ gives the free energy difference between ligands in the bound state.

Perturbation	$\Delta A^{inter}_{(aq)}$	$\Delta A^{inter}_{(b)}$	$\Delta\Delta A^{inter}_{bind}$
1→2	0.48 ± 0.02	1.49 ± 0.02	1.02 ± 0.03
2→4	-0.36 ± 0.01	-1.71 ± 0.01	-1.35 ± 0.02
4→5	-1.66 ± 0.01	-1.63 ± 0.02	0.03 ± 0.02
7→5	2.93 ± 0.04	0.79 ± 0.03	-2.13 ± 0.05
7→1	4.69 ± 0.01	2.73 ± 0.02	-1.96 ± 0.03
1→6	0.45 ± 0.01	4.50 ± 0.02	4.05 ± 0.02

(b) Results for interaction energy relative free energy of binding calculations using COX2 ligands.

Table 4.5.: Results for relative free energy of binding calculations using COX2 ligands. All values are given in kcal.mol^{-1}.

Cycle	Closure (kcal.mol^{-1})
$\Delta A_{(g)}$	0.72 ± 0.00
$\Delta A_{(aq)}$	0.56 ± 0.05
$\Delta A_{(b)}$	-0.26 ± 0.06
$\Delta A^{inter}_{(aq)}$	0.22 ± 0.05
$\Delta A^{inter}_{(b)}$	0.09 ± 0.05

Table 4.6.: Closure of thermodynamic cycle given in figure 4.7c according to calculation of different free energy differences.

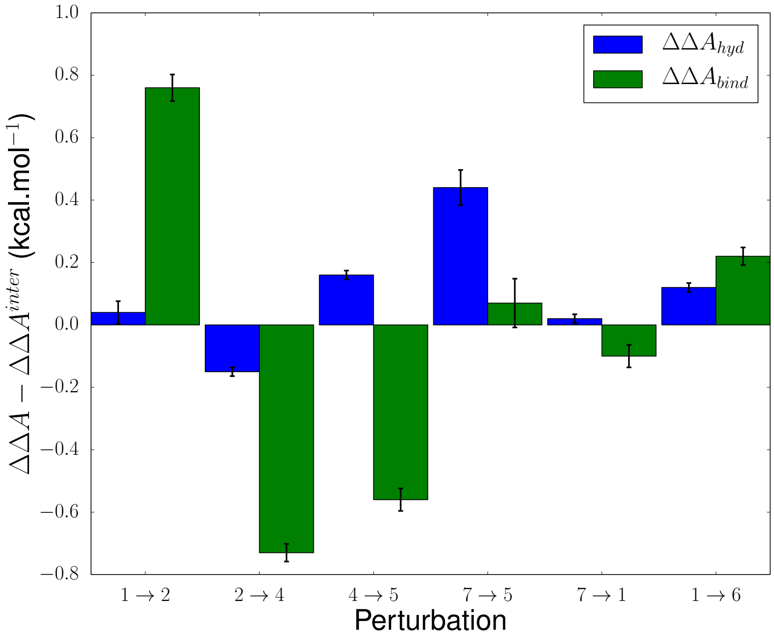

(a) Comparison of relative free energy differences calculated using total energies and interaction energies.

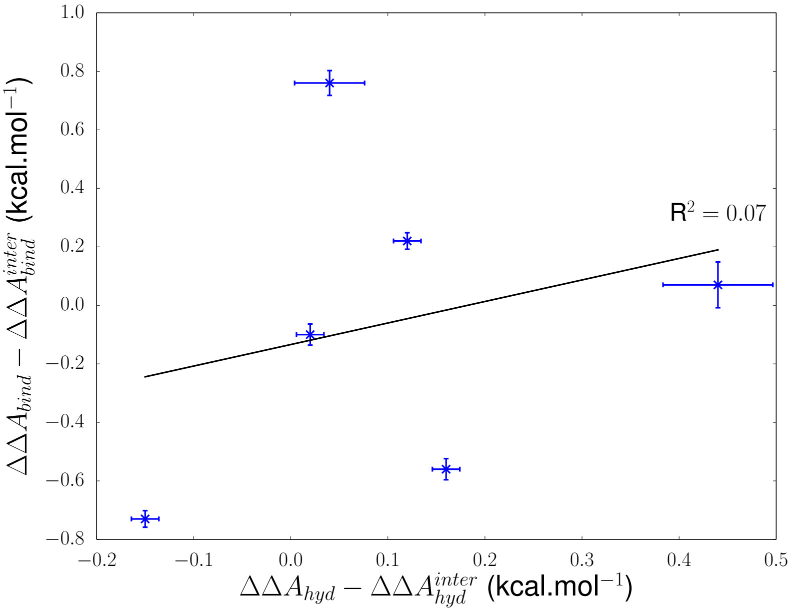

(b) Comparison of interaction errors arising for matching perturbations in the hydration and binding calculations. Line of best fit shown in black.

Figure 4.8.: COX2 interaction energy errors

error terms. Changes in ligand size or number of atoms are not associated with larger errors. Additionally errors are not consistent between binding and hydration calculations, the worst performing hydration being one of the best performing binding. A linear fit between the two for each perturbation gives an R^2 value of only 0.07 (see figure 4.8b).

The physical meaning of equation (4.40) was considered previously, however a further observation in the light of the above results is illuminating. It relates to the effect of the internal geometry intramolecular degrees of freedom on the interaction potentials. It is an indirect effect whereby the difference of internal geometry from state 3 leads to a difference in probability density experienced by the interaction potentials of state 4. This effect might be reasonably expected to be much more sigificant in a bound state than in solvent. Consider the perturbative growth of a ligand in the aqueous phase. As the dummy atoms extend out into the system, the solvent structure is freely able to adjust to accommodate the ligand, allowing the interaction potentials to satisfy their minima. However the environment of a protein binding site is far less homogeneous, as the perturbed group grows out of the ligand it experiences changes in environment, according to the conditions of the binding site. Thus differences in geometry between perturbations end-states may have a more pronounced effect in the context of binding calculations. Although somewhat speculative, it is interesting to note that this aligns with relative errors of hydration and binding in the COX2 system.

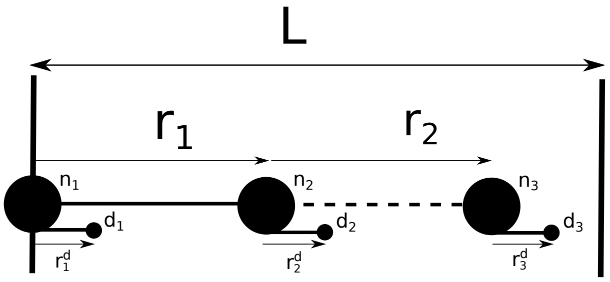

Figure 4.9.: Polarisable model system adapted from figure 4.2. Large circles denote atoms, whilst small circles denote Drude particles.

4.5. Polarisable Model System

A key advantage of QM theory is the inclusion of polarisation effects between molecules. These were not considered in the former model system nor in the calculations on COX2, making it difficult to draw conclusions from these that extend specifically to hybrid MM and QM work. A qualitative examination of polarisation effects is considered here through modification of the model system to include Drude particles[55]. The theoretical background of the use of Drude particles to model polarisation effects in classical energy models is given in section 2.2.2.5. The model system was modified from that of section 4.3 through the addition of a point charge and a mobile Drude particle for each atom. The energetic contribution of the Drude particles is given by:

$$U^{Drude}(\mathbf{r}^d;\mathbf{x}) = \sum_{i=1}^{3} k^d|\mathbf{r}_i^d|^2 + \sum_{i=1}^{3}\sum_{j>i}^{3} \frac{k_e d_i d_j}{r_{ij}} + \sum_{I=1}^{3}\sum_{j\neq I}^{3} \frac{k_e n_I d_j}{r_{Ij}} + \sum_{I=1}^{3}\sum_{J>i}^{3} \frac{k_e n_I n_J}{r_{IJ}} \tag{4.60}$$

This represents a specialised application of the more general expression given by equation (2.51), that has been modified to include the number of atoms within the model system and to restrict the value of k^d to be the same for all atoms. The potential energy function for the polarisable model system is given as:

$$U^{polar}(r_1, r_2) = Hm(r_1) + LJ^A(r_2) + \min\{U^{Drude}(r_1, r_2)\} \tag{4.61}$$

Where the final term represents the energy of the minimised configuration of Drude particle displacements for the given values of r_1 and r_2. For simplicity LJ^B is not considered within this version of the model system, as the introduction of point charges within the system is sufficient to prevent separability of the configuration integrals, through the introduction of an electrostatic interaction between the 1st and 3rd atoms. As noted through its formal definition (equation (2.53)), the isotropic polarisability of the model can be controlled through the combination of the parameters

d_i and k^d.

4.5.1. Solution of Model System Equations

Solution of equation (4.14) and equation (4.18) is complicated by numerical difficulties arising through the introduction of the Drude particles. Firstly for any given configuration of the nuclei, an algorithm is required to find the corresponding minimum of the Drude particle configuration space. This is achieved through the use of a non-linear conjugate gradient algorithm as given by Polak and Ribiere[101] and implemented within Scipy 0.12.0. No detailed explanation of the algorithm is provided here except to note it being chosen for its robustness and numerical stability.

Second, evaluating the energy of system configurations where nuclei are in close proximity to one another can lead to numerical infinities (the so called polarisation catastrophe effect[102]). This is avoided by introducing a buffer region between atoms over which the integrals are not solved, taking the form of a lower bound of greater than zero. This procedure does not alter the free energies calculated so long as the energetic contribution of Hm and LJ^A in these regions are large and positive. This will generally be true due to the short range repulsive property of the Lennard-Jones potential, but restricts the values that are suitable for use as a spring constant for Hm. Too soft a harmonic potential will cause the introduced buffer regions to have non-negligible contributions to the configurational integrals and lead to erroneous results. Care is taken to avoid this effect by assessing consistency of the calculated free energies with small changes in the size of the buffer regions.

4.5.2. Results and Discussion

Calculations on this system were carried out using the parameters in table 4.7a with additional parameters pertaining to the Drude particles given in table 4.7b. Each ligand has been assigned a set of charges q_i and the values of d_i are varied to control the polarisability of the system. Values of n_i are chosen such that the net charge q_i is preserved for each Drude-nucleus pair. All values of d_i were the same across different states within a particular cycle and are referred to as d. 'Gas phase' states involving only atoms 1 and 2 are neutral. States including the third 'solvent' atom obtain a net charge. This is designed such that inclusion of the third atom in the system is able to polarise the neutral 'solute'.

The results of calculations using this system are given in figure 4.10. Where $d = 0$, polarisation effects within the system hav been turned off, leaving $n_i = q_i$. The discrepancy between interaction and total energy results at this point is due to the presence of the nuclear point charges precluding integral separability, due to charge-charge interactions between the first and third nuclei. It is seen that as the value of

Parameter	Value
k_{P}	600 kcal.mol^{-1}.Å^{-2}
k_{Q}	1200 kcal.mol^{-1}.Å^{-2}
b^0_{P},b^0_{Q}	1.0 Å
ϵ^A_3	2.0 kcal.mol^{-1}
ϵ^A_4	4.0 kcal.mol^{-1}
σ^A_3,σ^A_4	1.0 Å
L	3.0

(a) Parameters for non-polarisable components of the model system.

Molecule	Parameters	Value
P	q_1	0.1 e
	q_2	-0.1 e
	d_1,d_2	d (var)
	n_1	$q_1 - d$
	n_2	$q_2 - d$
Q	q_1	0.15 e
	q_2	-0.15 e
	d_1,d_2	var
	n_1	$q_1 - d$
	n_2	$q_2 - d$
N/A	q_3	0.2 e
	d_3	var
	n_3	$q_3 - d$
	k^d	10^5 kcal.mol^{-1}.Å^{-2}

(b) Here e denotes units of elementary charge. Whilst subscripts refer to atom numbers rather than thermodynamic states within the cycle. Different states of the cycle are conveyed through the parameters associated with the two molecules P and Q. Parameters for the third 'solvent' atom are not varied between states 3 and 4 of the thermodynamic cycle.

Table 4.7.: Polarisable model system parameters

d becomes more negative, and hence polarisability becomes more pronounced, the error associated with the use of interaction energies increases. This effect is comparable to the use of the potential LJ^B. As the polarisability of the system is increased, the presence of the charged 'solvent' atom has a more pronounced effect on the internal geometry of the 'solute' atoms, leading to the observed increase in error. The magnitudes of the errors observed are not large compared to the magnitudes of the associated free energy differences. However, we note that the introduction of polarisation more than doubles the error associated with the unpolarisable system ($d = 0$) in the most extreme case. The range of polarisabilities explored in these calculations are similar to those employed in more realistic contexts[55].

One of the major rationales for the use of the hybrid free energy calculations is the inclusion of QM theory to accurately model polarisation effects. However, the use of interaction energies is a practical necessity as demonstrated in chapter 3 in the context of hybrid work. This represents an inherent conflict as interaction energies lose accuracy in precisely the high polarisability conditions that QM theory is invoked to model. Although enough to demonstrate the effect in principle, the results presented here using a model system cannot predict the extent of this problem in realistic biomolecular contexts.

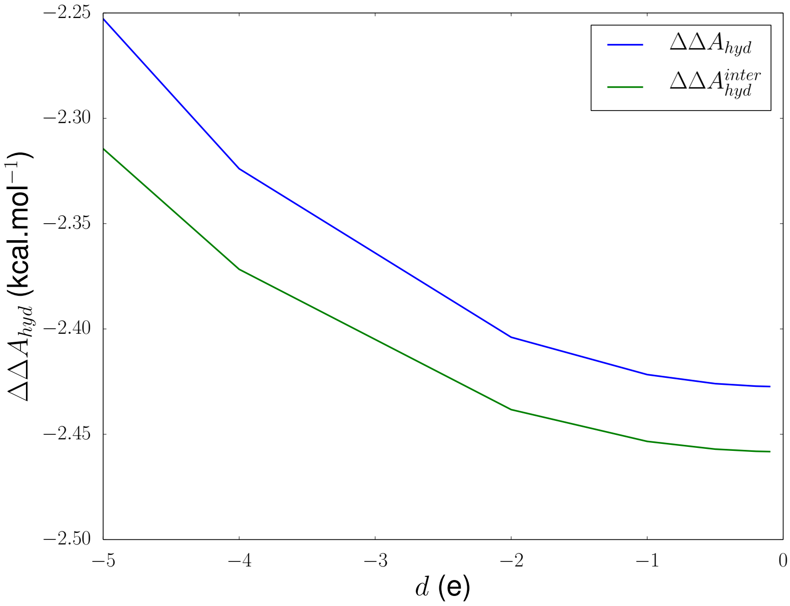

(a) Relative free energy of hydration results with increasing polarisability.

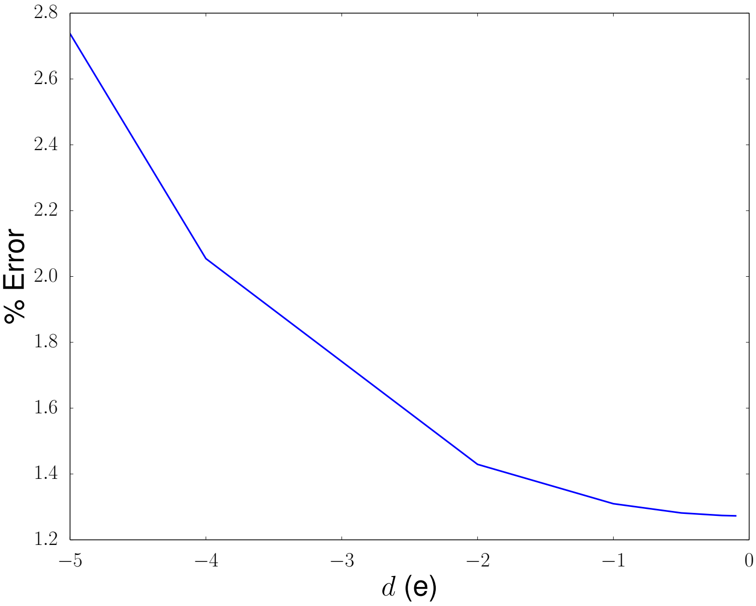

(b) Changes in the interaction energy error with increasing polarisability of the system.

Figure 4.10.: Results from calculations with polarisable model system.

4.6. Predicting the Error of Interaction Energies

Having established that the quality of the interaction energy approximation can vary depending on the specific properties of the system to which it is applied, it is desirable that we should be able to estimate the extent of this error. Comparison of interaction energy results directly with total energy results is not always possible as total energy calculations cannot always be carried out, for instance in the case of the MM to QM SSFEP. For this reason it would be advantageous to formulate a diagnostic approach capable of highlighting problematic calculations with minimial computational effort.

We will first consider the theoretical results of section 4.2 to motivate the form of such a diagnostic. Here we identify two independent contributions to the error arising through the use of interaction energies. It is required that the interaction and intramolecular degrees of freedom of the system be well approximated if they are assumed to be separable, i.e. $\Delta\Delta A_{hyd} \approx \Delta\Delta A_{hyd}^{(sep)}$ and $\Delta\Delta A_{hyd}^{inter} \approx \Delta\Delta A_{hyd}^{inter(sep)}$. Even in the limiting case of this being true it is shown that interaction and total energies results differ and it is hence required that equation (4.40) be small in value. There are multiple integrals within the proposed expressions that require separation. We can highlight those terms that will contribute directly to the error through the following:

$$
\begin{aligned}
&\Delta\Delta A_{hyd} - \Delta\Delta A_{hyd}^{inter} \\
&= -\frac{1}{\beta} \ln \frac{\dfrac{\int \exp\left[-\beta(U_{\mathrm{Q}}^{l} + U^{s} + U_{4}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_{\mathrm{P}}^{l} + U^{s} + U_{3}^{inter})\right] d\mathbf{r}d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U_{\mathrm{Q}}^{l}\right] d\mathbf{r}^{l}}{\int \exp\left[-\beta U_{\mathrm{P}}^{l}\right] d\mathbf{r}^{l}}} + \frac{1}{\beta} \ln \frac{\int \exp\left[-\beta(U_{\mathrm{P}}^{l} + U^{s} + U_{4}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_{\mathrm{P}}^{l} + U^{s} + U_{3}^{inter})\right] d\mathbf{r}d\mathbf{q}}
\end{aligned}
\tag{4.62}
$$

$$
= -\frac{1}{\beta} \ln \frac{\dfrac{\int \exp\left[-\beta(U_{\mathrm{Q}}^{l} + U^{s} + U_{4}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_{\mathrm{P}}^{l} + U^{s} + U_{4}^{inter})\right] d\mathbf{r}d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U_{\mathrm{Q}}^{l}\right] d\mathbf{r}^{l}}{\int \exp\left[-\beta U_{\mathrm{P}}^{l}\right] d\mathbf{r}^{l}}} \tag{4.63}
$$

Clearly interaction energies and total energies are equivalent where equation (4.63) is zero. It is interesting to note that this term does not contain the partition function of state 3. Instead the partition function of the notional state combining the intramolecular terms of state 3 with the interaction energies of state 4 is included, along with the partition function of state 4. As the interaction energy terms control the degree of correlation between the intra and intermolecular degrees of freedom this suggests that the properties of state 4 most significantly influence the error between interaction and total energies (this is supported through a series of model system calculations in appendix F). From equation (4.63):

$$-\frac{1}{\beta}\ln\frac{\dfrac{\int\exp\left[-\beta(U^l_{\mathrm{Q}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_4(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}{\int\exp\left[-\beta(U^l_{\mathrm{P}}(\mathbf{r}^l)+U^s(\mathbf{r}^s)+U^{inter}_4(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q}))\right]d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}}{\dfrac{\int\exp\left[-\beta U^l_{\mathrm{Q}}(\mathbf{r}^l)\right]d\mathbf{r}^l}{\int\exp\left[-\beta U^l_{\mathrm{P}}(\mathbf{r}^l)\right]d\mathbf{r}^l}} \tag{4.64}$$

Separate integrals by introducing the averages $\langle\mathbf{r}\rangle$ and introduce the scaling factors c^{corr}_1 and c^{corr}_2

$$=-\frac{1}{\beta}\ln\frac{\dfrac{c^{corr}_1\int\exp\left[-\beta U^l_{\mathrm{Q}}(\mathbf{r}^l)\right]d\mathbf{r}^l\int\exp\left[-\beta U^s(\mathbf{r}^s)\right]d\mathbf{r}^s\int\exp\left[-\beta U^{inter}_4(\langle\mathbf{r}\rangle_{\mathrm{P}},\mathbf{q}))\right]d\mathbf{q}}{c^{corr}_2\int\exp\left[-\beta U^l_{\mathrm{P}}(\mathbf{r}^l)\right]d\mathbf{r}^l\int\exp\left[-\beta U^s(\mathbf{r}^s)\right]d\mathbf{r}^s\int\exp\left[-\beta U^{inter}_4(\langle\mathbf{r}\rangle_{\mathrm{P}},\mathbf{q}))\right]d\mathbf{q}}}{\dfrac{\int\exp\left[-\beta U^l_{\mathrm{Q}}(\mathbf{r}^l)\right]d\mathbf{r}^l}{\int\exp\left[-\beta U^l_{\mathrm{P}}(\mathbf{r}^l)\right]d\mathbf{r}^l}} \tag{4.65}$$

Cancel terms

$$=-\frac{1}{\beta}\ln\frac{c^{corr}_1\int\exp\left[-\beta U^{inter}_4(\langle\mathbf{r}\rangle_{\mathrm{Q}},\mathbf{q})\right]d\mathbf{q}}{c^{corr}_2\int\exp\left[-\beta U^{inter}_4(\langle\mathbf{r}\rangle_{\mathrm{P}},\mathbf{q})\right]d\mathbf{q}} \tag{4.66}$$

Use equation (4.40)

$$=\Delta\Delta A^{(sep)}_{diff}-\frac{1}{\beta}\ln\frac{c^{corr}_1}{c^{corr}_2} \tag{4.67}$$

Rename ratio of scaling factors

$$=\Delta\Delta A^{(sep)}_{diff}+\Delta\Delta A^{corr} \tag{4.68}$$

Where we have introduced the terms c^{corr}_1 and c^{corr}_2. These are scale factors that recover the value of the correlated integrals from their uncorrelated approximations. We are also left with the contribution of the difference in separated terms, $\Delta\Delta A^{(sep)}_{diff}$, as given previously by equation (4.40). In practice the exact values of c^{corr}_1 and c^{corr}_2 are inaccesible. This is simply used as a convenient way to express the individual error contributions of each integral as a single term. For the value of $\Delta\Delta A^{corr}$ to approach zero we require the condition $c^{corr}_1 \approx c^{corr}_2$. We note this does not in fact require the configurational integrals to be perfectly separable, although this represents one situation where the condition is met.

Predicting the error associated with interaction energies can therefore be approached through attempting to approximate the terms $\Delta\Delta A^{(sep)}_{diff}$ and $\Delta\Delta A^{corr}$. We note that this analysis is limited to the case of a hydration free energy cycle. In the case of a binding calculation, $\Delta\Delta A^{(sep)}_{diff}$ will have a different form and two additional values of c^{corr} are introduced through the inclusion of the bound state within the thermodynamic cycle.

It is common in the case of purely classical free energy calculations to make use of the λ windowing approach (see section 2.4). As demonstrated by the results of section 4.2 and section 4.3.4 the value of the interaction free energy difference is dependent on the number of λ-windows used in a perturbation. We can extend the analysis of equations (4.64) to (4.68) to consider the use of additional λ-windows:

$$\Delta\Delta A^{2\lambda}_{hyd} - \Delta\Delta A^{2\lambda,inter}_{hyd} \tag{4.69}$$

$$\begin{aligned} &= -\frac{1}{\beta}\ln\frac{\frac{Z_{3.5}}{Z_3}}{\frac{Z_{1.5}}{Z_1}} - \frac{1}{\beta}\ln\frac{\frac{Z_4}{Z_{3.5}}}{\frac{Z_2}{Z_{1.5}}} \\ &\quad + \frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_3^{intra} + U_{3.5}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta U_3\right] d\mathbf{r}d\mathbf{q}} + \frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_{3.5}^{intra} + U_4^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta U_{3.5}\right] d\mathbf{r}d\mathbf{q}} \end{aligned} \tag{4.70}$$

$$= -\frac{1}{\beta}\ln\frac{\frac{\int \exp\left[-\beta U_{3.5}\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_3^{intra} + U_{3.5}^{inter})\right] d\mathbf{r}d\mathbf{q}}}{\frac{Z_{1.5}}{Z_1}} - \frac{1}{\beta}\ln\frac{\frac{\int \exp\left[-\beta U_4\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_{3.5}^{intra} + U_4^{inter})\right] d\mathbf{r}d\mathbf{q}}}{\frac{Z_2}{Z_{1.5}}} \tag{4.71}$$

$$= -\frac{1}{\beta}\ln\frac{c_1^{corr}\int \exp\left[-\beta U_{3.5}^{inter}(\langle\mathbf{r}\rangle_{\mathrm{P/Q}}, \mathbf{q})\right] d\mathbf{q}}{c_2^{corr}\int \exp\left[-\beta U_{3.5}^{inter}(\langle\mathbf{r}\rangle_{\mathrm{P}}, \mathbf{q})\right] d\mathbf{q}} - \frac{1}{\beta}\ln\frac{c_3^{corr}\int \exp\left[-\beta U_4^{inter}(\langle\mathbf{r}\rangle_{\mathrm{Q}}, \mathbf{q})\right] d\mathbf{q}}{c_4^{corr}\int \exp\left[-\beta U_4^{inter}(\langle\mathbf{r}\rangle_{\mathrm{P/Q}}, \mathbf{q})\right] d\mathbf{q}} \tag{4.72}$$

Where we introduce the scale factors $c_3^{corr}$and c_4^{corr} for the perturbation from the intermediate λ state 3.5 to state 4. The notation P/Q is used to denote the intermediate ligand states between P and Q. From this we conclude that errors arising from intermediate λ-windows also contribute to the interaction energy error of the perturbation. This does not necessarily mean that introducing λ-windows into a perturbation increases the overall error however. Although additional values of c^{corr} are introduced with additional λ-windows, it might be reasonably expected that the ratios $\frac{c_1^{corr}}{c_2^{corr}}$ and $\frac{c_3^{corr}}{c_4^{corr}}$ will more closely approach unity where their corresponding thermodynamic states are more similar to one another. Although each λ-window contributes individually to the error, as the perturbation performed is smaller the contribution of the error is also reduced. This suggests that in attempting to formulate a diagnostic for the use of interaction energies the individual error contribution of each λ-window must be considered.

In section 4.3.4 we demonstrated an increase in interaction energy error with additional λ windows. However, the effect is very small. The practical effect of the number of λ windows on the use of interaction energies remains an open question therefore.

In addition to the above considerations, we can specify the properties that an ideal diagnostic test would have. First and perhaps foremost it should be simple and feasible to calculate without considerable additional computational work to that required for the calculation itself. Preferably it would be applicable using data typically collected in the course of a free energy simulation run, allowing its application to historical and previously published work. In the case of being able to predict not only the magnitude but the sign of the interaction energy approximation, values from the diagnostic could be used to provide a correction term that restores the total free energy difference.

In the above analysis we note that the partition function of state 3 does not strictly enter into the error term given by equation (4.68). In the context of a hybrid calculation using SSFEP state 3 is simulated at the MM level whilst state 4, the QM, is not sampled at all. This means that we have access to ensemble average properties from the MM state but not the QM. Configurations from the MM state are post-processed to the QM level, allowing access to a certain amount of information about QM properties, but true QM ensemble averages are not available as configurations are weighted by an MM Boltzmann factor. The dependence of the error term equation (4.68) primarily on state 4 (the QM level), is problematic as this is not sampled directly. In such a case a successful diagnostic that retains the efficiency of the SSFEP approach must be accessible using only the QM processed configurations of state 3, or must be inferable from the ensemble average properties of state 3.

Also in the context of MM to QM perturbations, we note in section 4.5 that polarisation effects are correlated with an increase in interaction energy error. Methods for the calculation of polarisabilities at the QM level may therefore be used profitably to infer diagnostic information from configurations generated at the MM level.

In section 4.2 it is noted that the interaction energy approximation holds where the presence of intermolecular potentials does not alter the sampling of the ligand away from that which it adopts in the gaseous phase. Comparison of the sampling of the ligand with and without its surrounding environment may therefore yield information that can be used as a diagnostic. This approach has the drawback however, of requiring additional work in the form of a gas phase calculation. As these are not required for relative free energy of binding calculations this requires additional computational investment. Additionally, if using interaction energies to carry out a free energy of hydration calculation a great benefit is avoiding the need to simulate the ligand in isolation. A complication of this idea is that comparing ligand sampling requires the comparison of two N-dimensioned distributions in configuration space, the interpretation of which may be complex and unclear. An alternative may be to consider changes in the sampled energy distributions that are conveniently single dimensioned.

Building from this we consider the meaning of the scale factors c^{corr}. Hypothesising that these are directly proportional to the separability of their integrals, we note that the physical manifestation of this effect may be related to the degree of correlation between the intra and intermolecular degrees of freedom. Where the integrals under consideration are perfectly separable, the different degrees of freedom of the system will sample independently of one another. The correlation of the intra and intermolecular degrees of freedom can be measured directly through the calculation of correlation coefficients between the corresponding energy distributions. This can be easily calculated from data gathered from an MD or MC run.

We have applied several of the above approaches to the model systems with some success. However, this work remains at a tentative stage and hence is not presented here. In particular further work is required to refine approaches suitable for application to the COX2 system.

4.7. Conclusions

The results presented here provide a strong argument for caution when making use of interaction energies in free energy calculations however they are frequently used within free energy techniques as a practical necessity. In the case of the MM to QM FEP the use of interaction energies not only offers numerical tractability but also provides much greater overlap between the end points of the calculation. We demonstrate strongly in chapter 3 that poor overlaps with total energies requires the use of interaction energies making their use unavoidable.

The results presented here provide insight to the consequences of using interaction energies within free energy calculations. A theoretical elucidation of the underlying approximate nature of interaction energy calculations is given. The use interaction energies with the Zwanzig equation is shown to equate to a perturbation from the sampled state to a notional hybrid state, composed of the intramolecular terms of the source state and the interaction energy of the target state. We demonstrate that this perturbation approaches the total relative free energy of hydration under the limit of separability of the intramolecular and interaction degrees of freedom of the system and where the averaged intramolecular distributions between source and target states are similar.

We demonstrate the validity of these results through application to a series of systems of increasing complexity. The model system employed allows exact quantification of the interaction energy error, and is used to directly support the derived theoretical results. We then move on to consider comparable calculations within a realistic biological system. Poor convergence of the total energy results precludes any firm conclusions but the size of the interaction energy error observed for individual perturbations is compelling. In particular for the binding energy calculations discrepancies of up to 0.8 kcal.mol^{-1} between interaction and total energies were observed for particular perturbations.

One of the most compelling cases for hybrid free energy are systems displaying strong polarisation effects that cannot be appropriately described at the MM level. For this reason we extend the model system to account directly for polarisation effects through the addition of Drude particles. This allows a qualitative examination of the unique effects at play for the case of perturbations involving a QM Hamiltonian. We demonstrate that increasing the strength of polarisation effects is associated with increases in interaction energy error. This suggests an unfortunate tension between the need to use QM energy models to describe polarisable systems and the introduction of error through interaction energies.

Finally we move on to consider the need for a diagnostic approach capable of highlighting calculations that are problematic with the use of interaction energies. We extend the theoretical analysis to this end and provide a number of valuable observations and criteria that such a test should meet. Unfortunately despite discussion of several possibilities, we have been unable to demonstrate a working example within the more realistic COX2 system. As such, this remains a focus for future work.

t

5. Ensemble Generation/Reweighting Approaches

5.1. Introduction

So far this work has focussed on the practical and theoretical aspects of SSFEP as the most commonly used hybrid free energy technique. Recently however, several more sophisticated approaches have been proposed and proof of principle calculations have been made[7,8,69,75]. These were discussed in some detail and their theoretical background developed in section 2.5.

In this chapter we will focus on the application of molecular mechanics based importance function (MMBIF) sampling with multi-'timestepping' Monte Carlo and the use of the non-Boltzmann Zwanzig equation. The primary difference between these methodologies is that MMBIF is able to actively guide sampling of the MM state to match the QM, whilst non-Boltzmann approaches are carried out purely through post-processing of an existing MM ensemble. Calculations are carried out with the model thermodynamic system of chapter 4 and the DNA base pair MD dataset presented in chapter 3. Particular attention is paid to the combination of these approaches with interaction energies and some theoretical results to this end are given in addition to the computational work.

We introduce a novel post-processing method based upon resampling of an equilibrium MM distribution and use of a MC acceptance criteria. We refer to this resampling approach and MMBIF collectively as ensemble building techniques. Resampling provides an option for ensemble building as a purely post-processing technique, thus avoiding the interdependency of QM calculations necessitated by MMBIF. This is discussed more fully in section 5.3.

The techniques from this section were developed in section 2.5 in the context of calculations between MM and QM states. Many of the calculations performed here however, are not explicitly between MM and QM states, but also between MM Hamiltonians and states of the model system. Wherever convenient we have maintained the terminology of MM and QM in order to make apparent the application of these techniques in the context of actual hybrid free energy work.

$$\begin{tikzcd}
\text{QM} & & \text{MM} \\
1{:}\mathrm{P}_{(g)} \arrow[rr, "\Delta A^{(g)}_{P\to Q}"] \arrow[d] & & 2{:}\mathrm{Q}_{(g)} \arrow[ll, bend right] \arrow[d] \\
3{:}\mathrm{P}_{(aq)} \arrow[rr, "\Delta A^{(aq)}_{P\to Q}"'] & & 4{:}\mathrm{Q}_{(aq)} \arrow[ll, bend left]
\end{tikzcd}$$

Figure 5.1.: Thermodynamic cycle used in model system calculations. Cycle states in red are sampled by the auxiliary Markov chain, the blue arrows denoting the construction of an ensemble with MMBIF. The blue labels MM and QM are given for analogy with a hybrid free energy calculation.

5.2. Molecular Mechanics Based Importance Function

As developed fully in section 2.5.3, MMBIF techniques provide the ability to construct a rigorously correct QM ensemble from sampling of only a reference MM ensemble. A specialised acceptance test is used to remove the biasing influence caused by sampling under the MM Hamiltonian and produce an ensemble with the correct QM Boltzmann factor for each configuration. We consider here application of this methodology within the model thermodynamic system of section 4.3 and consider its extension to the use of interaction energies.

5.2.1. Model System

An MMBIF implementation based on multi-“timestepping” Monte Carlo is here employed in conjunction with the model thermodynamic system. This presents an ideal test bed, as it is possible to directly calculate overlaps in configuration space using the overlap integral described in section 3.2. Additionally, free energy differences computed with the MMBIF approach may be compared with direct solution of the partition function for the model system as previously (section 4.3).

The results given here represent application of the standard MMBIF approach de-

Parameter	Value
k_{P}	50 kcal.mol^{-1}.Å^{-2}
k_{Q}	100 kcal.mol^{-1}.Å^{-2}
b^0_{P},b^0_{Q}	1.0 Å
ϵ^A_3	1.0 kcal.mol^{-1}
ϵ^A_4	1.5 kcal.mol^{-1}
ϵ^B_3,ϵ^B_4	var kcal.mol^{-1}
σ^A_3,σ^A_4	1.0 Å
σ^B_3,σ^B_4	1.7 Å
L	3.0

Table 5.1.: Parameter sets used in calculations with the model system from section 4.3 to produce figure 5.3.

scribed in section 2.5.3. The standard Metropolis-Hastings algorithm is employed for the auxiliary Markov chain used to generate structures. After every ten moves of the auxiliary chain an acceptance test for the QM chain is attempted as given by:

$$P_{MMBIF}(\mathbf{x}_1 \to \mathbf{x}_2) = \min\left(1, \exp\left[-\beta \Delta\Delta U\right]\right) \tag{5.1}$$

$$\Delta\Delta U = \left(U_{QM}(\mathbf{x}_2) - U_{MM}(\mathbf{x}_2)\right) - \left(U_{QM}(\mathbf{x}_1) - U_{MM}(\mathbf{x}_1)\right) \tag{5.2}$$

overlap integrals of the model system degrees of freedom r_1 and r_2 between different states are calculated using equation (3.5) as described in detail in section 3.4.6.

The descriptions of MM and QM in this context are somewhat meaningless, as we have no true QM Hamiltonian for the model system. Instead we shall use an alternative set of parameters to define distinct states with the system that shall be considered analogous to the QM. Free energy cycles will therefore be computed as shown in figure 5.1. To complete the analogy, P states stand in for the QM system, whilst Q states represent the classical. Only the states 2 and 4, are sampled and used to build the states 1 and 3, respectively. The free energy differences $\Delta A^{(g)}_{P\to Q}$ and $\Delta A^{(aq)}_{P\to Q}$ are then computed using the ensembles constructed using MMBIF. Hence we are using a reference state to construct a target ensemble in order to calculate the perturbation back to the reference ensemble. This represents a somewhat trivial use of MMBIF, but it serves to illustrate its application without the need to include additional thermodynamic states within the cycle. Free energy differences are calculated using the Zwanzig equation to complete the perturbation from the constructed ensemble back to the reference ensemble.

5.2.1.1. Total Energies

Results for total energy MMBIF calculations are shown in figure 5.3 and were carried out using the parameters set given in table 5.1. A simple visual demonstration of the effectiveness of the ensemble building approach is given in figure 5.2. As can

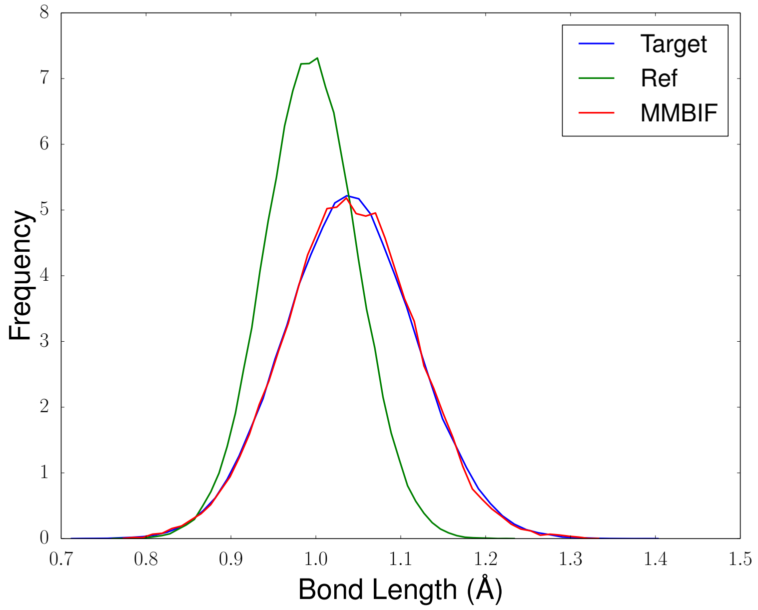

(a) r_1

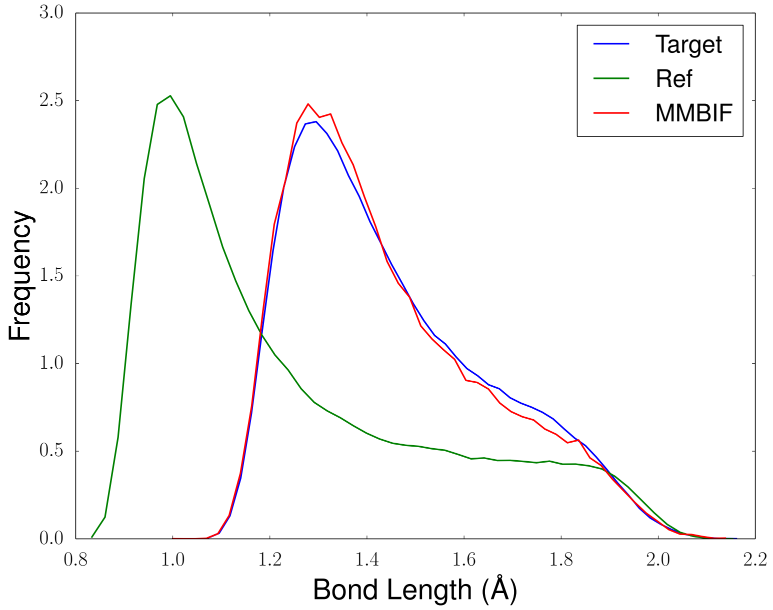

(b) r_2

Figure 5.2.: Probability distributions of r_1 and r_2. The distribution Ref is the system state on which the auxiliary MMBIF Markov chain is constructed. Target is the distribution that MMBIF attempted to reproduce, and MMBIF denotes the actual distribution produced by ensemble building.

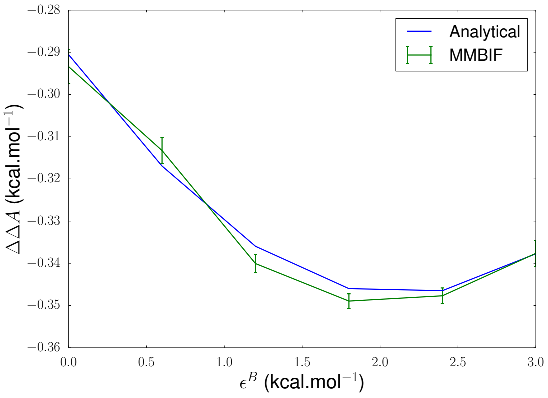

(a) Free energies of hydration

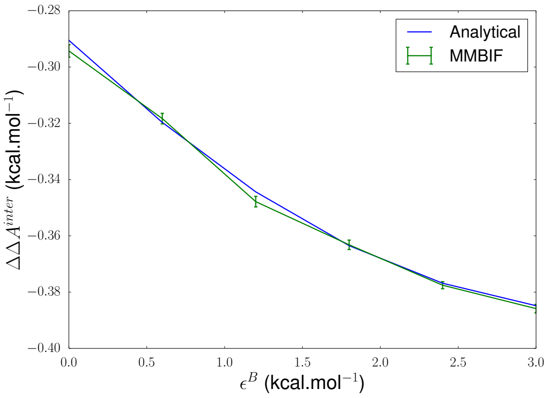

(b) Interaction free energies of hydration

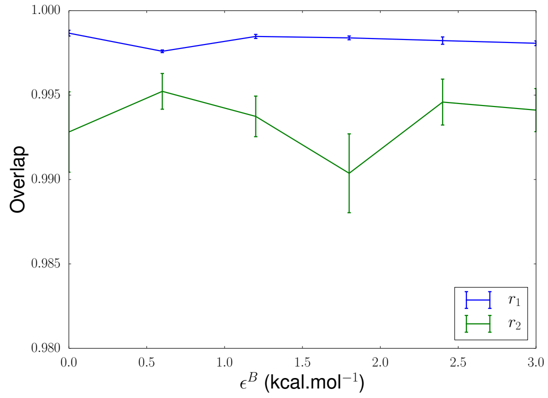

(c) Overlaps between target and MMBIF derived distribution for the model system degrees of freedom r_1 and r_2, calculated using equation (3.5)

Figure 5.3.: Model system calculations using MMBIF. Results derived from direct solution of the partition function are compared with results calculated with MMBIF using parameters as given in table 5.1

be seen, the distribution of bond lengths produced by MMBIF, closely resembles the target distribution, although composed only of configurations generated by the reference distribution Markov chain. Excellent overlap between the MMBIF and target distributions is shown over the full parameter range considered (see figure 5.3c). Excellent agreement is also seen in calculated free energies of hydration using both total and interaction energies between MMBIF and direct numerical solution. The results presented here are intended only to serve as a general purpose demonstration of an application of the MMBIF methodology and to validate our implementation of the algorithm as correct.

5.2.2. Interaction Energies

Having demonstrated a working example MMBIF calculations within the model system we move on to consider the use of interaction energies with MMBIF. This is carried out by maintaining the use of the standard (total energy) Metropolis-Hastings test for the auxiliary 'MM' Markov chain, but using a modified acceptance test at the 'QM' level, given by:

$$\begin{aligned} P^{inter}_{MMBIF}(\mathbf{x}_1 \to \mathbf{x}_2) &= \min\left(1, \exp\left[-\beta\Delta\Delta U^{inter}\right]\right) && (5.3) \\ \Delta\Delta U^{inter} &= (U^{inter}_{QM}(\mathbf{x}_2) - U^{inter}_{MM}(\mathbf{x}_2)) - (U^{inter}_{QM}(\mathbf{x}_1) - U^{inter}_{MM}(\mathbf{x}_1)) && (5.4) \end{aligned}$$

we here make no rigorous attempt to theoretically justify an acceptance test of this form, but its application to the model system is illustrative in the context of the results presented in section 4.2. Two conditions were imposed on the reliability of interaction energies in calculations using FEP, the independence of the intra and intermolecular degrees of freedom, discussed in section 4.3.2, and consistency of the intramolecular equilibrium geometry between states, discussed in section 4.3.3. Violation of either of these conditions causes interaction energies to give divergent results from total energies. The relation of these two conditions to the behaviour of MMBIF with the use of interaction energies are considered individually with two sets of simulations, designed to be analogous to the scenarios considered in sections 4.3.2 and 4.3.3.

Starting with the criteria of separability and reusing the parameters given in table 5.1, gives the results of figure 5.5. A compelling demonstration of the behaviour of MMBIF is given in figure 5.4. Here $\epsilon^B = 1.5$ kcal.mol^{-1}, so the system intra and intermolecular degrees of freedom are not perfectly separable. Nonetheless the target state distribution of the interaction degree of freedom, r_2, is well reproduced by interaction energy MMBIF. Figure 5.5c shows that overlap of the r_2 distribution between interaction energy MMBIF and direct MC sampling of the target distribution remains excellent over the entire range of parameters considered and is not contingent on the separability of the degrees of freedom of the system. The intramolecular distribution of the system meanwhile remains completely unaltered from the reference state regardless of the value of ϵ^B. This result is somewhat intuitive; removing any information relating to the intramolecular degrees of freedom from the acceptance test leaves the technique unable to reweight the relevant intramolecular degrees of freedom. The observation regarding overlaps is borne out through the calculated free energy differences compared against direct numerical solution. Whilst matching total free energies is dependent on the established separability condition, interaction free energy differences appear to be accurately recovered through the use of equation (5.4) as the acceptance test.

We now move to consider the second requirement for the viability of interaction energies. This can be related to the consistency of internal geometries between states, but is more formally expressed as the requirement that equation (4.40) is zero. As in section 4.3 this effect is explored within the model system by changing values of b^0 at different states of the cycle, i.e. $b^0_\mathrm{P} \neq b^0_\mathrm{Q}$. Results presented here employ the parameters given by table 5.2.

In contrast to the results relating to separability, changes in internal equilibrium geometry between target and references states, disrupt the behaviour of MMBIF with the use of interaction energies. Overlaps between MMBIF and target intermolecular distributions are shown in figure 5.7c. As b^0_Q, and therefore the discrepancy between internal equilibrium geometries, is increased the observed overlaps drop off. The reason for this can be intuitively demonstrated by the results shown in figure 5.6. As before, the intramolecular distribution remains unaltered from the reference state. However, within the constructed ensemble the distribution of r_1 is shifted with regards to the target state. This shift combined with the boundary conditions of the system, prevent the generation of configurations of r_2 of greater than 2.0 Å, as these configurations are not generated, the MMBIF procedure is not able to include them within the constructed ensemble. This effect is counteracted when using MMBIF with total energies as configurations for the intramolecular distribution are also correctly reweighted, allowing generation and acceptance of appropriate configurations of the intermolecular distribution. Despite largely matching the target distribution, the final result of interaction energy MMBIF is in fact something of a hybrid between the target and reference states. This failure to correctly reconstitute the interaction degree of freedom distribution leads to an accumulation of error in the calculated free energy differences, not just using total energies, but also interaction energies.

A note of caution perhaps should be injected into interpretation of these results however. Whilst it is interesting to note the effect observed in figures 5.5 and 5.7 this in part arises due to the boundary conditions applied within the model system. The use of periodic boundary conditions would preclude this effect for instance, by removing the dependency of the maximum length of r_2 on r_1.

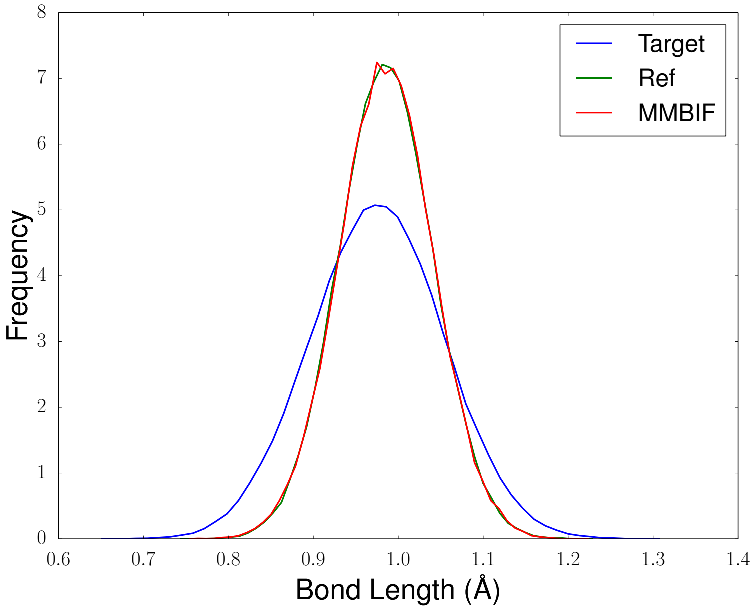

(a) r_1

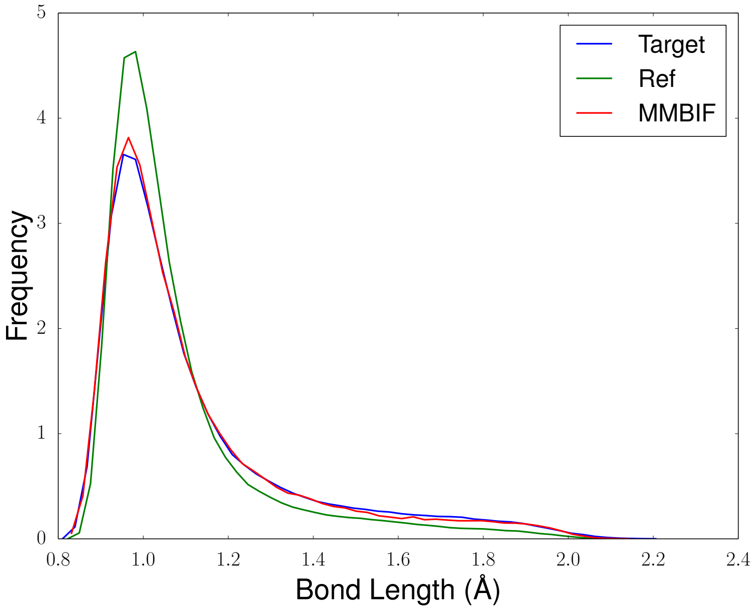

(b) r_2

Figure 5.4.: Probability distributions of r_1 and r_2 constructed with interaction energy MMBIF. The distribution Ref is the system state on which the auxiliary Markov chain is constructed. Target is the distribution that MMBIF attempted to reproduce, and MMBIF is the actual distribution produced by ensemble building.

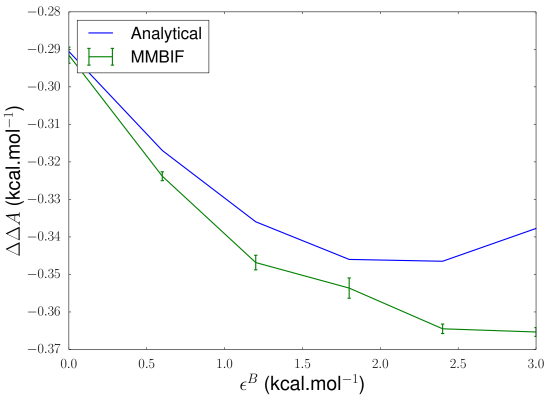

(a) Free energies of hydration

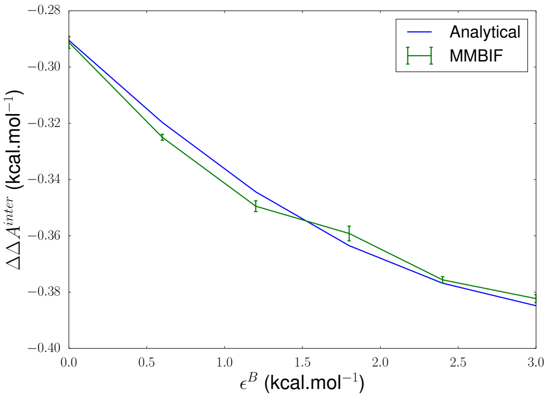

(b) Interaction free energies of hydration

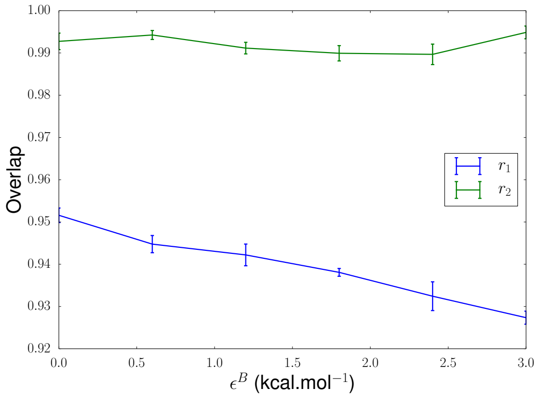

(c) Overlaps between target and MMBIF derived distribution for the model system degrees of freedom r_1 and r_2.

Figure 5.5.: Model system calculations using interaction energy MMBIF. Results derived from direct solution of the partition function are compared with results calculated with MMBIF using parameters as given in table 5.2

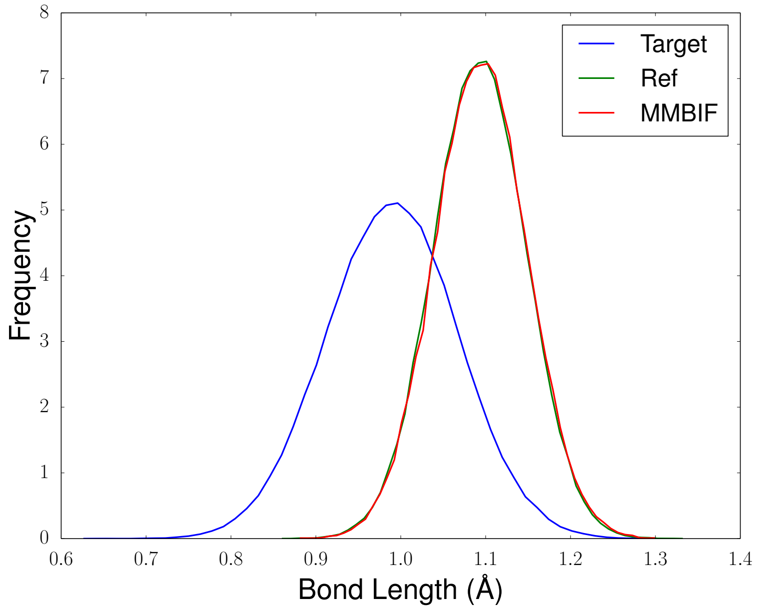

(a) r_1

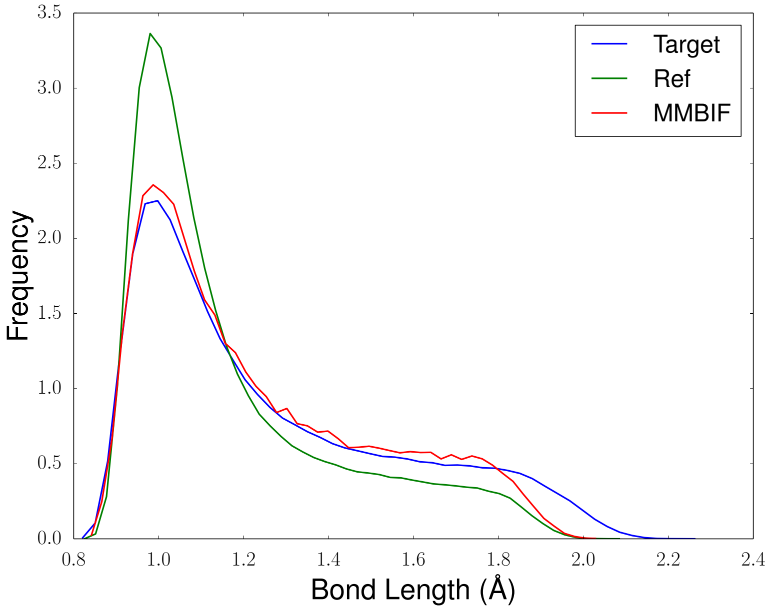

(b) r_2

Figure 5.6.: Probability distributions of r_1 and r_2. The distribution Ref is the system state on which the auxiliary MMBIF Markov chain is constructed. Target is the distribution that MMBIF attempted to reproduce, and MMBIF is the actual distribution produced by ensemble building.

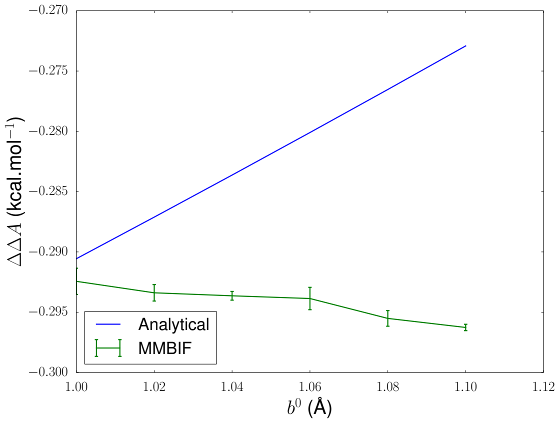

(a) Free energies of hydration

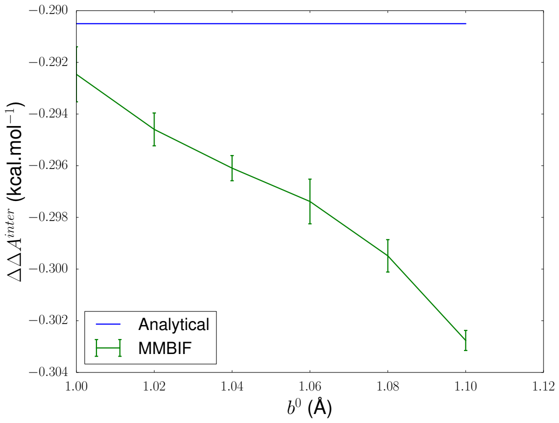

(b) Interaction free energies of hydration

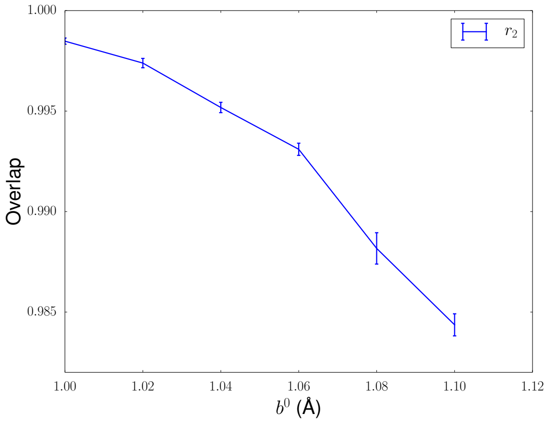

(c) Overlaps between target and MMBIF derived distribution for the intermolecular degrees of freedom r_2. Here we omit the intramolecular degree of freedom r_1, as it displays unvarying poor overlap.

Figure 5.7.: Model system calculations using interaction energy MMBIF. Results derived from direct solution of the partition function are compared with results calculated with MMBIF using the parameters as given in table 5.2

Parameter	Value
k_P	50 kcal.mol^{-1}
k_Q	100 kcal.mol^{-1}
b^0_P	1.0 Å
b^0_Q	var Å
ϵ^A_3	1.0 kcal.mol^{-1}
ϵ^A_4	1.5 kcal.mol^{-1}
ϵ^B_3,ϵ^B_4	0.0 kcal.mol^{-1}
σ^A_3,σ^A_4	1.0 Å
σ^B_3,σ^B_4	2.0 Å
L	3.0

Table 5.2.: Parameter sets used in the calculations with the model system for figure 5.7.

5.3. Monte Carlo Resampling

Despite the similarity of the MMBIF and non-Boltzmann approaches, a significant advantage claimed for MMBIF is the biasing of sampling toward regions of the phase space suitable to the QM ensemble (this is discussed in section 2.5.5). The extent to which this is beneficial however is unclear. We here propose an alternative MC based resampling method, similar to MMBIF but without the sampling bias. This is similar therefore to the non-Boltzmann approaches that reweight MM ensembles to recover QM properties, except a full QM ensemble is reconstituted. The MC resampling procedure is given as follows:

1. Generate an ensemble of MM structures with MD or MC, $\mathbf{x}_{MM,i}$.
2. Randomly select a configuration from $\mathbf{x}_{MM,i}$ to be $\mathbf{x}_1$.
3. Randomly select a configuration from $\mathbf{x}_{MM,i}$ to be $\mathbf{x}_2$.
4. Apply the acceptance test

$$P^{acpt}_{rsmp}(\mathbf{x}_1 \to \mathbf{x}_2) = \min(1, \exp\left[-\beta\Delta\Delta U\right]) \tag{5.5}$$

$$\Delta\Delta U = (U_{QM}(\mathbf{x}_2) - U_{MM}(\mathbf{x}_2)) - (U_{QM}(\mathbf{x}_1) - U_{MM}(\mathbf{x}_1)) \tag{5.6}$$

5. If test passed, $\mathbf{x}_2$ is added to the Markov chain and becomes $\mathbf{x}_1$
6. If test failed, discard $\mathbf{x}_2$ and add another copy of $\mathbf{x}_1$ to the Markov chain
7. Repeat from 3.

The acceptance test proposed has the same form as equation (5.1), the acceptance test for MMBIF. The general form of the Metropolis test is given by equation (2.62). As previously (see sections 2.3.2 and 2.5.3), the full form of equation (5.5) is dependent on the corresponding value of:

$$\frac{P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)} \tag{5.7}$$

where $P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)$ is the probability of generating the trial configuration $\mathbf{x}_2$ from $\mathbf{x}_1$. In this case the term has the form:

$$P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \frac{1}{N_{MM}} P_{MM}(\mathbf{x}_2) \tag{5.8}$$

where, N_{MM} is the number of configurations within the previously generated classical ensemble, $\mathbf{x}_{MM,i}$. In addition $P_{MM}(\mathbf{x}_2)$ is the absolute probability of $\mathbf{x}_2$ arising within the classical Boltzmann distribution. Expanding this to encompass both directions:

$$\begin{align} \frac{P_g(\mathbf{x}_2 \rightarrow \mathbf{x}_1)}{P_g(\mathbf{x}_1 \rightarrow \mathbf{x}_2)} &= \frac{\frac{1}{N_{MM}} P_{MM}(\mathbf{x}_1)}{\frac{1}{N_{MM}} P_{MM}(\mathbf{x}_2)} \tag{5.9} \\ &= \frac{P_{MM}(\mathbf{x}_1)}{P_{MM}(\mathbf{x}_2)} \tag{5.10} \\ &= \exp\left[-\beta(U_{MM}(\mathbf{x}_1) - U_{MM}(\mathbf{x}_2))\right] \tag{5.11} \end{align}$$

inserting this into equation (2.62) gives the full test as in equation (5.5). As with MMBIF, once the ensemble has been constructed the free energy difference is then calculated through the application of the Zwanzig equation. This gives the free energy difference back to the reference ensemble from the constructed target ensemble.

5.3.1. Model System

Validation of the resampling approach is provided through application to the model system of section 4.3, as discussed fully in section 5.2.1. Calculations presented in the following are resampled from 'MM' ensembles constructed using 10^6 moves of the standard Metropolis-Hastings algorithm. Every one in ten configurations, giving a total of 10^5, was then used with the resampling procedure to build an ensemble of 10^6 configurations. As we are carrying out sampling with replacement from the 10^5 structures, the resulting ensemble will contain multiple instances of certain structures. This is designed to simulate the post-processing of only a subset of the MM configurations generated by MD or MCs to the QM level.

5.3.1.1. Total Energies

The results shown here make use of the same parameters employed previously for calculations with MMBIF and are given in table 5.1. The results shown here demon-

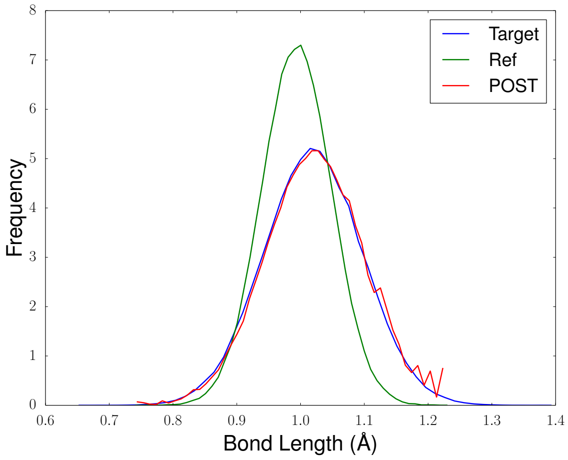

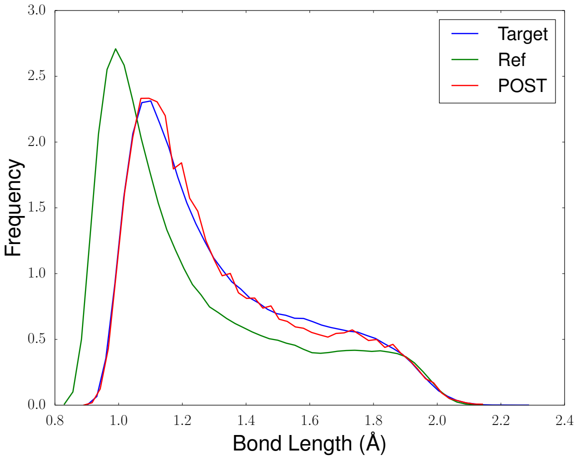

Figure 5.8.: Probability distributions of r_1 and r_2 constructed with interaction energy resampling. The distribution Ref is the resampled system state. Target is the distribution that resampling MC attempts to reproduce, and POST is the actual distribution produced by ensemble building.

strate good agreement with results derived from direct solution of the partition function and the total energy MMBIF results. Similarly these results are only intented to provide a basic proof of principle for this approach and validation of the algorithm's implementation.

5.3.1.2. Interaction Energies

By analogy with the use of interaction energies with MMBIF considered previously, the use of the interaction energies with MC resampling is demonstrated here. As before the interaction acceptance test is given by:

$$P_{rsmp}^{acpt,inter}(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = \min(1, \exp\left[-\beta\Delta\Delta U^{inter}\right]) \tag{5.12}$$

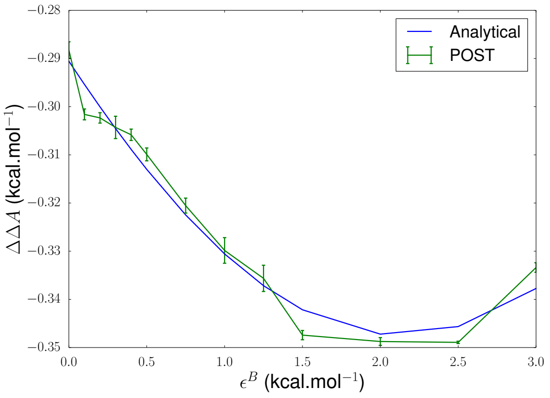

(a) Free energies of hydration

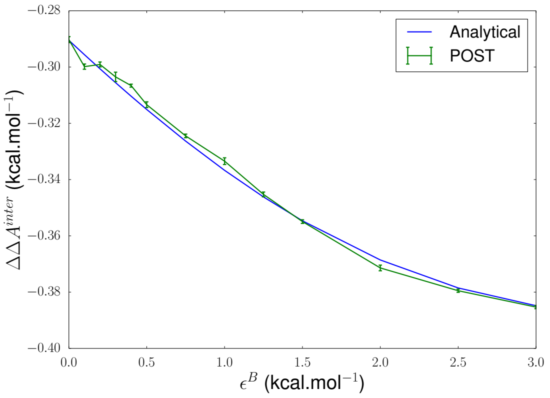

(b) Interaction free energies of hydration

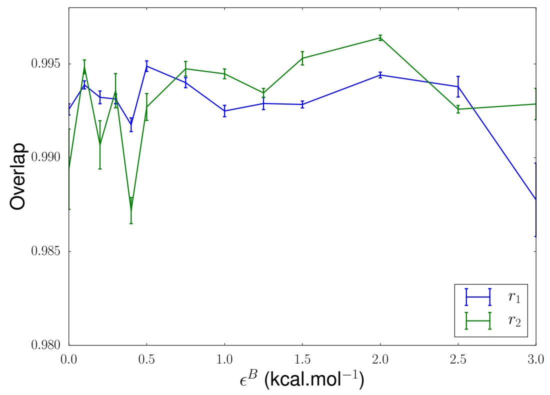

(c) Overlaps between target and resampling derived distribution for the intra and intermolecular degrees of freedom r_1 and r_2 respectively.

Figure 5.9.: Model system calculations using MC resampling. Results derived from direct solution of the partition function are compared with results calculated with MC resampling using parameters as given in table 5.1

$$\Delta\Delta U^{inter} = (U_{QM}^{inter}(\mathbf{x}_2) - U_{MM}^{inter}(\mathbf{x}_2)) - (U_{QM}^{inter}(\mathbf{x}_1) - U_{MM}^{inter}(\mathbf{x}_1)) \quad (5.13)$$

calculations using MC resampling with interaction energies were carried out with the same test sets as employed previously with MMBIF (section 5.2). Compellingly, precisely the same properties as observed for MMBIF with interaction energies are observed in the case of the resampling procedure. Interaction free energy differences are accurately reproduced despite the correlation between the intramolecular and interaction energy terms (figure 5.10), whilst changes of internal geometry disrupt the reconstitution of both interaction and total free energy differences (figure 5.11).

Owing to the similarity of the procedures employed it is not surprising that MC resampling and MMBIF give identical properties with the use of interaction energies. The main advantage offered by the proposed resampling approach, is that it can be applied purely as a post-processing technique.

5.3.2. DNA Base Pair

The data presented in chapter 3 using the adenosine-thymidine DNA base pair in vacuum is amenable to post-processing using the MC resampling procedure presented here. The nature of the data set with four interconnected perturbations allows extensive post-processing with resampling. Three different calculations were performed, resampling with equation (5.5) was used to calculate total and interaction free energy differences (table 5.3 and figure 5.12) and resampling with equation (5.12) was used to calculate interaction free energy differences. For each of these sets of calculations multiple different perturbations involving three thermodynamic states were carried out. This allows many different calculations to be constructed from subsets of the four states of the cycle. In essence, any state can be resampled to produce any other state, and then the free energy of perturbation to any third state can be calculated. Each application of resampling used 450 configurations from the MD trajectories to build an ensemble of 200000 configurations. All standard errors shown are between the five independent repeat MD trajectories that were generated for each state of the cycle.

Results for the total energy calculations (table 5.3) are of limited value due to the calculations being unconverged in the case of the single-step perturbation, it is thus difficult to assess the quality of the resampled free energy differences as no reliable estimate of the true value is available. However, considering the variability of different estimates for each perturbation, the source ensemble chosen to resample impacts considerably the resulting free energy difference. We would conclude therefore that the use of resampling in these approaches does not offer any improvement over the single-step perturbation approach. Considered in the light of the extremely poor total phase space overlap between states of the base pair system observed in sec-

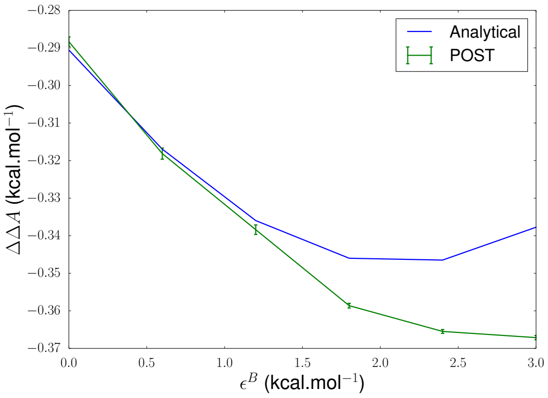

(a) Free energies of hydration

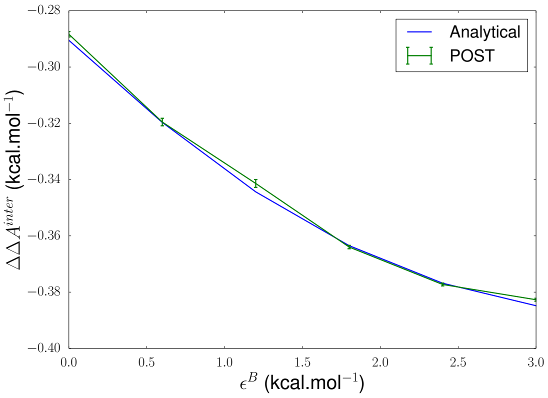

(b) Interaction free energies of hydration

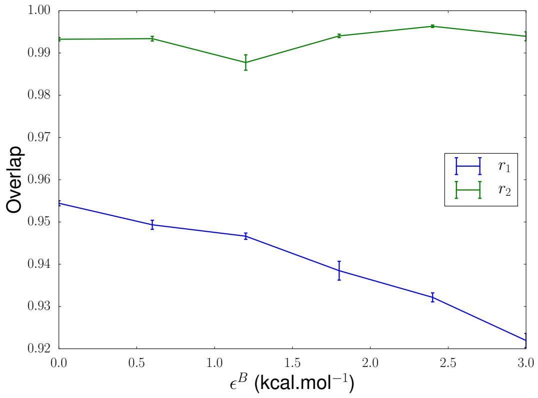

(c) Overlaps between target and resampling derived distribution for the intra and intermolecular degrees of freedom r_1 and r_2 respectively.

Figure 5.10.: Model system calculations using MC resampling with interaction energies. Results derived from direct solution of the partition function are compared with results calculated with MMBIF using parameters as given in table 5.1

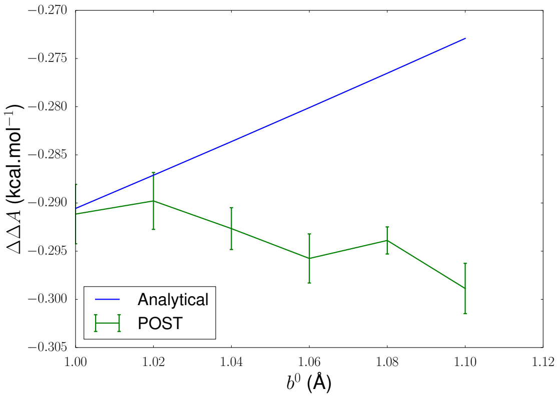

(a) Free energies of hydration

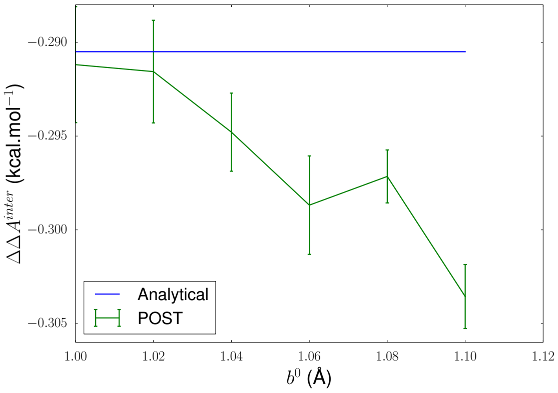

(b) Interaction free energies of hydration

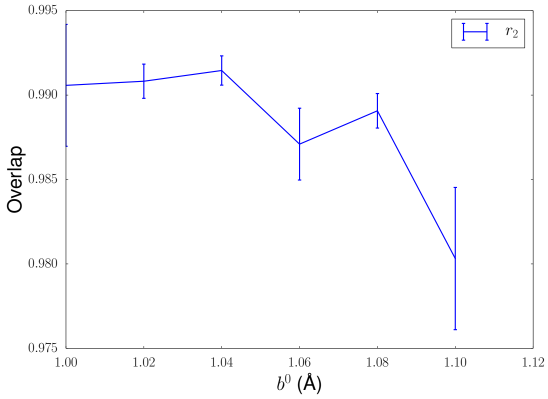

(c) Overlaps between target and MC resampling derived distribution for the intra and intermolecular degrees of freedom r_1 and r_2 respectively.

Figure 5.11.: Model system calculations using MC resampling with interaction energies. Results derived from direct solution of the partition function are compared with results calculated with resampling using parameters as given in table 5.2

tion 3.4.6 however, this result is perhaps unsurprising. Encouragingly however all free energy differences fall within the correct magnitude and are thus entirely consistent with the SSFEP results presented.

Interaction energy differences calculated through resampling with total energies are given in figure 5.12. Agreement with SSFEP results is generally reasonable except in the case of the ff99SB→ LDA perturbation. The mean unsigned error for reweighted results against SSFEP is 1.56 kcal.mol^{-1}. The interaction free energy differences derived from resampling with equation (5.12) however, are markedly improved (see figure 5.13), the majority of the perturbations showing good agreement with SSFEP. These give a mean unsigned error of 1.04 kcal.mol^{-1}.

	Ensemble built from			
Perturbation	PBE	LDA	ff99SB	gaff
PBE→LDA	1780.62 ± 1.14	1766.35 ± 0.47	1779.77 ± 1.71	1780.16 ± 2.60
PBE→ff99SB	207075.46 ± 1.69	207081.83 ± 1.62	207058.00 ± 1.12	207058.64 ± 1.37
PBE→gaff	206843.77 ± 0.98	206849.89 ± 1.85	206828.39 ± 1.48	206797.60 ± 3.47
LDA→PBE	-1780.84 ± 0.93	-1767.20 ± 1.32	-1779.09 ± 1.12	-1780.87 ± 2.25
LDA→ff99SB	205299.82 ± 1.22	205314.64 ± 4.74	205274.40 ± 2.17	205281.43±1.10
LDA→gaff	205070.72 ± 0.75	205080.84 ± 3.43	205049.51 ± 2.32	205021.34±1.28
ff99SB→PBE	-207079.39 ± 0.92	-207090.69 ± 2.74	-207059.68 ± 1.81	-207046.05±3.16
ff99SB→LDA	-205299.82 ± 2.46	-205319.90 ± 2.69	-205280.22 ± 4.42	-205276.90±3.29
ff99SB→gaff	-229.86 ± 1.28	-232.10 ± 1.55	-232.21 ± 2.80	-251.62±1.70
gaff→PBE	-206852.23 ± 1.28	-206851.67 ± 2.01	-206824.09 ± 1.27	-206812.00 ± 1.17
gaff→LDA	-205068.36 ± 2.99	-205096.17 ± 2.80	-205045.65 ± 2.03	-205029.69 ± 4.52
gaff→ff99SB	230.98 ± 1.06	231.92 ± 1.19	229.79 ± 0.33	250.87 ± 1.30

Table 5.3.: Application of the resampling methodology to total free energy differences for the DNA base pair system. Entries shown in red are the results from the use of SSFEP presented in chapter 3.

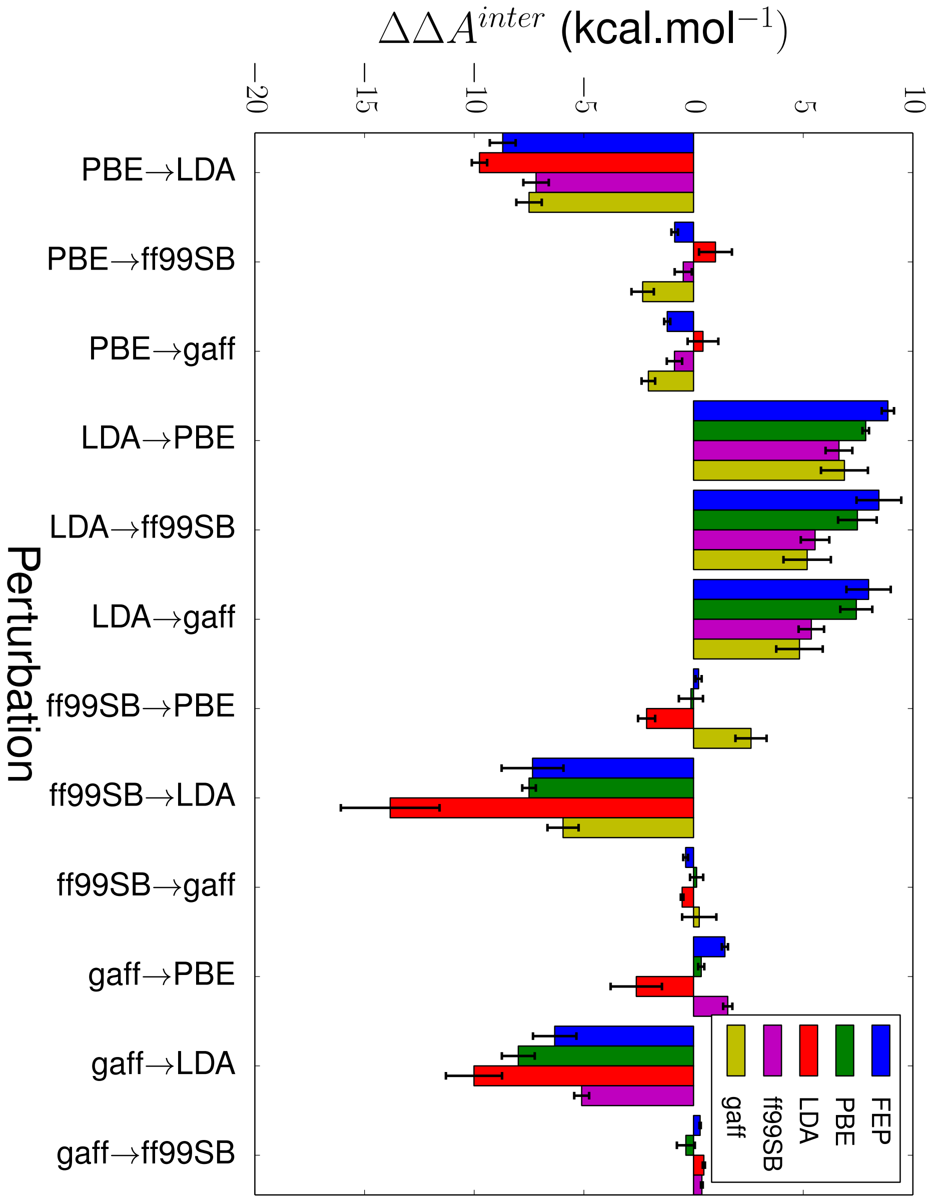

Figure 5.12.: Interaction free energy differences from total energy resampling for the DNA base pair system. Bars are coloured according to the ensemble that was built **from**. The blue FEP data is taken from the results presented in chapter 3. Error bars show a single standard error between repeats.

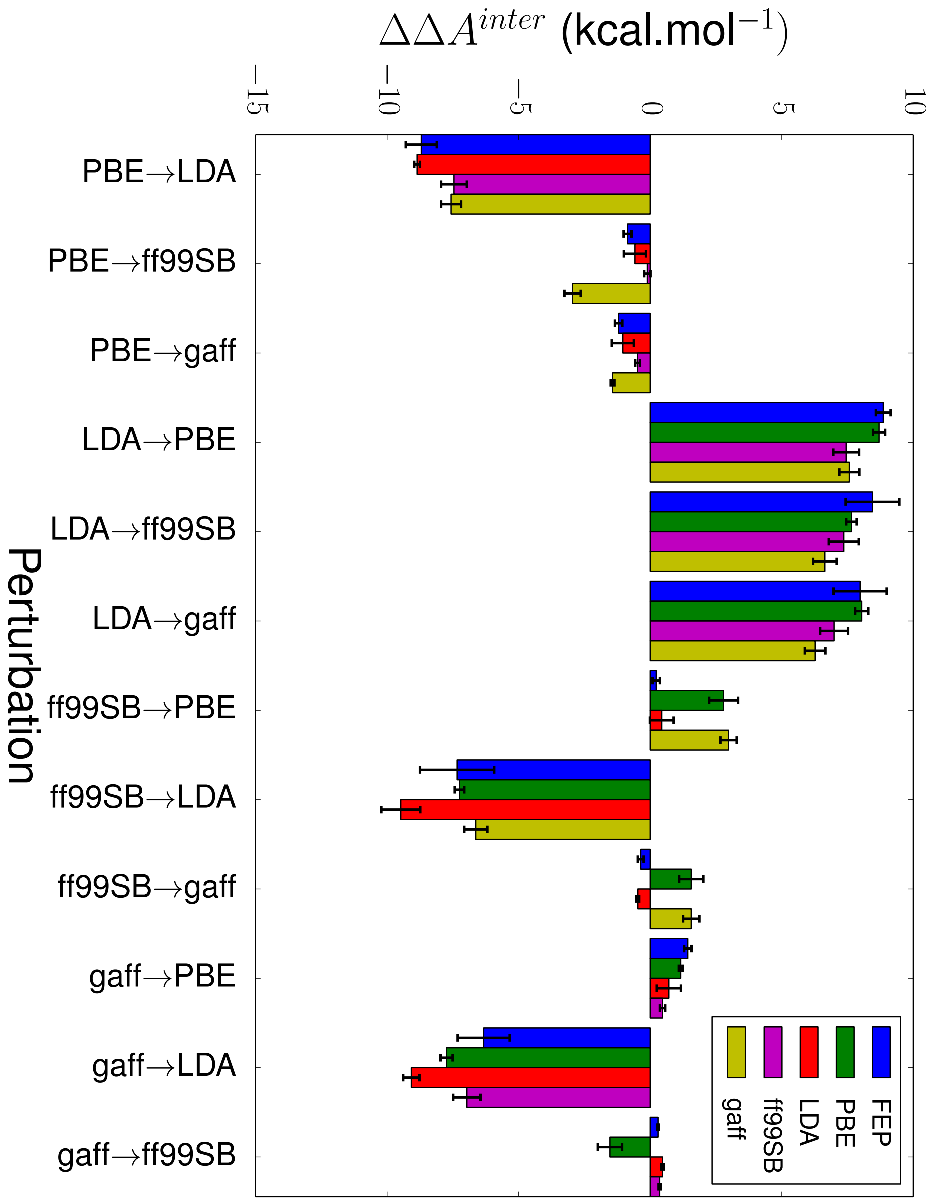

Figure 5.13.: Interaction free energy differences from interaction energy resampling with equation (5.12) for the DNA base pair system. Bars are coloured according to the ensemble that was built **from**. The blue FEP data is taken from the results presented in chapter 3. Error bars show a single standard error between repeats.

5.4. Non-Boltzmann Free Energy Perturbation

Detailed derivation of the non-Boltzmann approaches are given in section 2.5, but we shall again briefly outline the theory here for clarity. Non-Boltzmann approaches are based around the core result:

$$\langle X \rangle_{QM} = \frac{\langle X \exp[\beta V] \rangle_{MM}}{\langle \exp[\beta V] \rangle_{MM}} \tag{5.14}$$

this provides a means to recover a true QM ensemble average of the property X from sampling under an MM Hamiltonian. Where V is known as the bias potential and is given by:

$$V = U_{MM} - U_{QM} \tag{5.15}$$

Choosing $X = \exp\left[-\beta(U_{QM,1} - U_{QM,0})\right]$ gives the non-Boltzmann Zwanzig equation:

$$\Delta A_{QM,0\to 1} = -\frac{1}{\beta} \ln \frac{\langle \exp\left[-\beta(U_{QM,1} - U_{QM,0})\right] \exp[\beta V_0] \rangle_{MM,0}}{\langle \exp[\beta V_0] \rangle_{MM,0}} \tag{5.16}$$

here $\Delta A_{QM,0\to 1}$ is the free energy difference between the chemically distinct states 0 and 1 at the QM level of theory. This result is only dependent however on an ensemble average of state 0 at the MM level. This result can be applied therefore by generating an ensemble of state 0 at the MM level and postprocessing configurations to obtain $U_{QM,0}$ and $U_{QM,1}$. This result was reported and applied, along with its BAR counterpart, by König *et al.*[8,79].

5.4.1. Interaction Energy Theory

We shall here consider the novel extension of the theory to the use of interaction energies in the context of the Zwanzig equation. We propose two novel results, based around the combinatino of interaction energies with equation (5.16) in two ways. Firstly, and perhaps slightly trivially, we choose the value of X to be the interaction energy difference, i.e. $X = \exp\left[-\beta(U^{inter}_{QM,1} - U^{inter}_{QM,0})\right]$. Substituting this into equation (5.16) we can derive:

$$\begin{aligned} &\exp[\beta \Delta A^{inter}_{QM,0\to 1}] \\ =&\frac{\left\langle \exp\left[-\beta(U^{inter}_{QM,1} - U^{inter}_{QM,0})\right] \exp[\beta V_0] \right\rangle_{MM,0}}{\langle \exp[\beta V_0] \rangle_{MM,0}} \end{aligned} \tag{5.17}$$

$$=\frac{\dfrac{\int \exp\left[-\beta U_{MM,0}\right]\exp\left[-\beta(U^{inter}_{QM,1}-U^{inter}_{QM,0})\right]\exp\left[-\beta(U_{QM,0}-U_{MM,0})\right]d\mathbf{x}}{\int \exp\left[-\beta U_{MM,0}\right]d\mathbf{x}}}{\dfrac{\int \exp\left[-\beta U_{MM,0}\right]\exp\left[-\beta(U_{QM,0}-U_{MM,0})\right]d\mathbf{x}}{\int \exp\left[-\beta U_{MM,0}\right]d\mathbf{x}}} \tag{5.18}$$

$$=\frac{\int \exp\left[-\beta(U_{MM,0}-U_{MM,0})\right]\exp\left[-\beta U_{QM,0}\right]\exp\left[-\beta(U^{inter}_{QM,1}-U^{inter}_{QM,0})\right]d\mathbf{x}}{\int \exp\left[-\beta(U_{MM,0}-U_{MM,0})\right]\exp\left[-\beta U_{QM,0}\right]d\mathbf{x}} \tag{5.19}$$

$$=\frac{\int \exp\left[-\beta U_{QM,0}\right]\exp\left[-\beta(U^{inter}_{QM,1}-U^{inter}_{QM,0})\right]d\mathbf{x}}{\int \exp\left[-\beta U_{QM,0}\right]d\mathbf{x}} \tag{5.20}$$

$$=\left\langle \exp\left[-\beta(U^{inter}_{QM,1}-U^{inter}_{QM,0})\right]\right\rangle_{QM,0} \tag{5.21}$$

at this point we have simply reproduced the interaction Zwanzig equation at the QM level. Using a series of steps analogous to those of section 4.2 gives:

$$\Delta A^{inter}_{QM,0\to1}=-\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U^{intra}_{QM,0}+U^{inter}_{QM,1})\right]d\mathbf{x}}{\int \exp\left[-\beta(U^{intra}_{QM,0}+U^{inter}_{QM,0})\right]d\mathbf{x}} \tag{5.22}$$

this is completely analogous to equation (4.24), differences in notation not withstanding, and we therefore expect use of interaction energies in this context to accurately reproduce interaction free energy differences. Use of equation (5.17) will be referred to as the **total bias interaction energy** approach.

The second novel expression based on equation (5.16) arises through considering the extension of interaction energies to include the bias potential which we now give as:

$$V^{inter}=U^{inter}_{MM}-U^{inter}_{QM} \tag{5.23}$$

Replacing V_0 in equation (5.23) with equation (5.17) gives:

$$\Delta A^{interb}_{QM,0\to1}=-\frac{1}{\beta}\ln\frac{\left\langle \exp\left[-\beta(U^{inter}_{QM,1}-U^{inter}_{QM,0})\right]\exp[\beta V^{inter}_0]\right\rangle_{MM,0}}{\left\langle \exp[\beta V^{inter}_0]\right\rangle_{MM,0}} \tag{5.24}$$

this will be referred to as the **interaction bias** approach, the associated free energy difference denoted as $\Delta A^{interb}_{QM,0\to1}$. Replacing the total energy bias with the interaction energy bias, is possible but how does $\Delta A^{interb}_{QM,0\to1}$ relate to the total free energy difference? Continuing from equation (5.24):

$$\exp\left[-\beta\Delta A^{interb}_{QM,0\to1}\right]$$

$$=\frac{\left\langle \exp\left[-\beta(U^{inter}_{QM,1}-U^{inter}_{QM,0})\right]\exp[\beta V^{inter}_0]\right\rangle_{MM,0}}{\left\langle \exp[\beta V^{inter}_0]\right\rangle_{MM,0}} \tag{5.25}$$

$$=\frac{\dfrac{\int \exp\left[-\beta U_{MM,0}\right] \exp\left[-\beta(U^{inter}_{QM,1} - U^{inter}_{QM,0})\right] \exp\left[-\beta(U^{inter}_{QM,0} - U^{inter}_{MM,0})\right] d\mathbf{x}}{\int \exp\left[-\beta U_{MM,0}\right] d\mathbf{x}}}{\dfrac{\int \exp\left[-\beta U_{MM,0}\right] \exp\left[-\beta(U^{inter}_{QM,0} - U^{inter}_{MM,0})\right] d\mathbf{x}}{\int \exp\left[-\beta U_{MM,0}\right] d\mathbf{x}}} \tag{5.26}$$

$$=\frac{\int \exp\left[-\beta(U_{MM,0} - U^{inter}_{MM,0})\right] \exp\left[-\beta U^{inter}_{QM,1}\right] d\mathbf{x}}{\int \exp\left[-\beta(U_{MM,0} - U^{inter}_{MM,0})\right] \exp\left[-\beta U^{inter}_{QM,0}\right] d\mathbf{x}} \tag{5.27}$$

$$=\frac{\int \exp\left[-\beta(U^{intra}_{MM,0} + U^{inter}_{QM,1})\right] d\mathbf{x}}{\int \exp\left[-\beta(U^{intra}_{MM,0} + U^{inter}_{QM,0})\right] d\mathbf{x}} \tag{5.28}$$

Interestingly this result is analogous to equation (5.22) however with the MM intramolecular terms in place of the QM intramolecular terms. Equation (5.28) therefore contains energy terms that do not pertain to either end-state of the perturbation, i.e. the QM states. Introduction of the separability approximation, as previously (section 4.2), gives, from equation (5.22):

$$\frac{\int \exp\left[-\beta(U^{intra}_{QM,0}(\mathbf{r}) + U^{inter}_{QM,1}(\mathbf{r},\mathbf{q}))\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U^{intra}_{QM,0}(\mathbf{r}) + U^{inter}_{QM,0}(\mathbf{r},\mathbf{q}))\right] d\mathbf{r}d\mathbf{q}} \tag{5.29}$$

$$\approx\frac{\int \exp\left[-\beta U^{intra}_{QM,0}(\mathbf{r})\right] d\mathbf{r} \int \exp\left[-\beta U^{inter}_{QM,1}(\langle\mathbf{r}\rangle_{QM,0}, \mathbf{q})\right] d\mathbf{q}}{\int \exp\left[-\beta U^{intra}_{QM,0}(\mathbf{r})\right] d\mathbf{r} \int \exp\left[-\beta U^{inter}_{QM,0}(\langle\mathbf{r}\rangle_{QM,0}, \mathbf{q})\right] d\mathbf{q}} \tag{5.30}$$

$$=\frac{\int \exp\left[-\beta U^{inter}_{QM,1}(\langle\mathbf{r}\rangle_{QM,0}, \mathbf{q})\right] d\mathbf{q}}{\int \exp\left[-\beta U^{inter}_{QM,0}(\langle\mathbf{r}\rangle_{QM,0}, \mathbf{q})\right] d\mathbf{q}} \tag{5.31}$$

and from equation (5.28):

$$\frac{\int \exp\left[-\beta(U^{intra}_{MM,0}(\mathbf{r}) + U^{inter}_{QM,1}(\mathbf{r},\mathbf{q}))\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U^{intra}_{MM,0}(\mathbf{r}) + U^{inter}_{QM,0}(\mathbf{r},\mathbf{q}))\right] d\mathbf{r}d\mathbf{q}} \tag{5.32}$$

$$\approx\frac{\int \exp\left[-\beta U^{intra}_{MM,0}(\mathbf{r})\right] d\mathbf{r} \int \exp\left[-\beta U^{inter}_{QM,1}(\langle\mathbf{r}\rangle_{MM,0}, \mathbf{q})\right] d\mathbf{q}}{\int \exp\left[-\beta U^{intra}_{MM,0}(\mathbf{r})\right] d\mathbf{r} \int \exp\left[-\beta U^{inter}_{QM,0}(\langle\mathbf{r}\rangle_{MM,0}, \mathbf{q})\right] d\mathbf{q}} \tag{5.33}$$

$$=\frac{\int \exp\left[-\beta U^{inter}_{QM,1}(\langle\mathbf{r}\rangle_{MM,0}, \mathbf{q})\right] d\mathbf{q}}{\int \exp\left[-\beta U^{inter}_{QM,0}(\langle\mathbf{r}\rangle_{MM,0}, \mathbf{q})\right] d\mathbf{q}} \tag{5.34}$$

In the limit of separability both equations (5.22) and (5.28) approach the desired free energy difference. The difference lies in the implicit effect of the intramolecular terms. Equation (5.31) corresponds to the interaction energy result equation (4.24), however the viability of replacing the averaged QM intramolecular geometries with

MM geometries is not readily apparent. In principle, so long as the effect on the interaction energy terms are equivalent, equation (5.34) will give comparable results to equation (5.31). This should be expected to vary on a case by case basis and so we move on to consider the practical application of this approach.

5.4.2. Base Pair

Calculations were carried out through post-processing of the dataset presented in chapter 3. Three sets of calculations were made using equations (5.16), (5.17) and (5.24). Results using the total bias total energy approach are comparable with those of section 5.3.2. As discussed there, the poor convergence of the SSFEP results prevents any firm conclusions from being drawn. However, general consistency in the size of the perturbations gives general validation of the approach. Results for the total bias interaction energy approach are given in figure 5.14. The quality of these results compared to SSFEP is markedly variable. Certain perturbations show reasonable agreement with SSFEP, for instance PBE→LDA, LDA→PBE, gaff→ff99SB. However the majority show fairly poor agreement, and in some cases considerable errors are seen, eg. gaff→LDA. In particular the ff99SB→LDA result reweighted from the gaff ensemble is off by greater than 10 kcal.mol^{-1}. The mean unsigned error of the non-Boltzmann results against SSFEP is 1.93 kcal.mol^{-1}.

Improvement on the total energy and total bias energy results results is seen in the use of the interaction bias approach (see figure 5.15), with a mean unsigned error of 1.05 kcal.mol^{-1}. The disastrous ff99SB→LDA reweighted from the gaff ensemble is improved considerably in particular. We would speculate that the reason for this improvement may attributed to the presence of the classical intramolecular energies within equation (5.28). In the case of total bias interaction energies the total energies within the bias still require the reweighting of intramolecular energy distributions to produce the correct ensemble averages. The complete removal of the intramolecular energies in the interaction bias method therefore simplifies the requirement for phase space overlap. Only the average effect of the MM intramolecular energies are required in equation (5.34) and these are determined directly from sampling of the MM state, rather than through reweighting.

5.4.3. Comparison of MC Resampling and non-Boltzmann approaches

As the two post-processing techniques considered here it is interesting to compare the results derived from the resampling and non-Boltzmann approaches. We consider here the accuracy with which each methodology recreates the well converged interaction energy SSFEP results from the base pair system (figure 3.9b). As can be seen from table 5.5, use of interaction energies with either resampling or non-Boltzmann

	Ensemble built from			
Perturbation	PBE	LDA	ff99SB	gaff
PBE→LDA	1780.62 ± 1.14	1767.20 ± 0.53	1778.01 ± 1.67	1781.92 ± 2.34
PBE→ff99SB	207075.46 ± 1.69	207081.74 ± 1.45	207058.23 ± 0.78	207049.28 ± 1.33
PBE→gaff	206843.77 ± 0.98	206848.03 ± 1.24	206826.02 ± 1.24	206811.94 ± 0.46
LDA→PBE	-1780.63 ± 0.46	-1767.20 ± 1.32	-1778.01 ± 1.67	-1780.16 ± 2.62
LDA→ff99SB	205297.17 ± 1.24	205314.64 ± 4.74	205280.22 ± 1.77	205280.89 ± 1.52
LDA→gaff	205065.98 ± 1.17	205080.84 ± 3.43	205048.01 ± 1.45	205029.69 ± 1.96
ff99SB→PBE	-207075.12 ± 0.44	-207081.24 ± 1.73	-207059.68 ± 1.81	-207048.51 ± 1.96
ff99SB→LDA	-205294.50 ± 0.86	-205313.94 ± 2.27	-205280.22 ± 4.42	-205266.10 ± 3.13
ff99SB→gaff	-231.30 ± 0.71	-233.71 ± 1.36	-232.21 ± 2.80	-237.34 ± 1.35
gaff→PBE	-206843.82 ± 0.40	-206848.40 ± 1.50	-206826.02 ± 1.24	-206812.00 ± 1.17
gaff→LDA	-205063.19 ± 0.55	-205081.10 ± 1.69	-205048.02 ± 1.45	-205029.69 ± 4.52
gaff→ff99SB	231.30 ± 0.71	233.71 ± 1.36	232.21 ± 1.12	250.87 ± 1.30

Table 5.4.: Application of the non-Boltzmann Zwanzig equation to total free energy differences for the DNA base pair system. Entries shown in red are the results from the use of SSFEP presented in chapter 3.

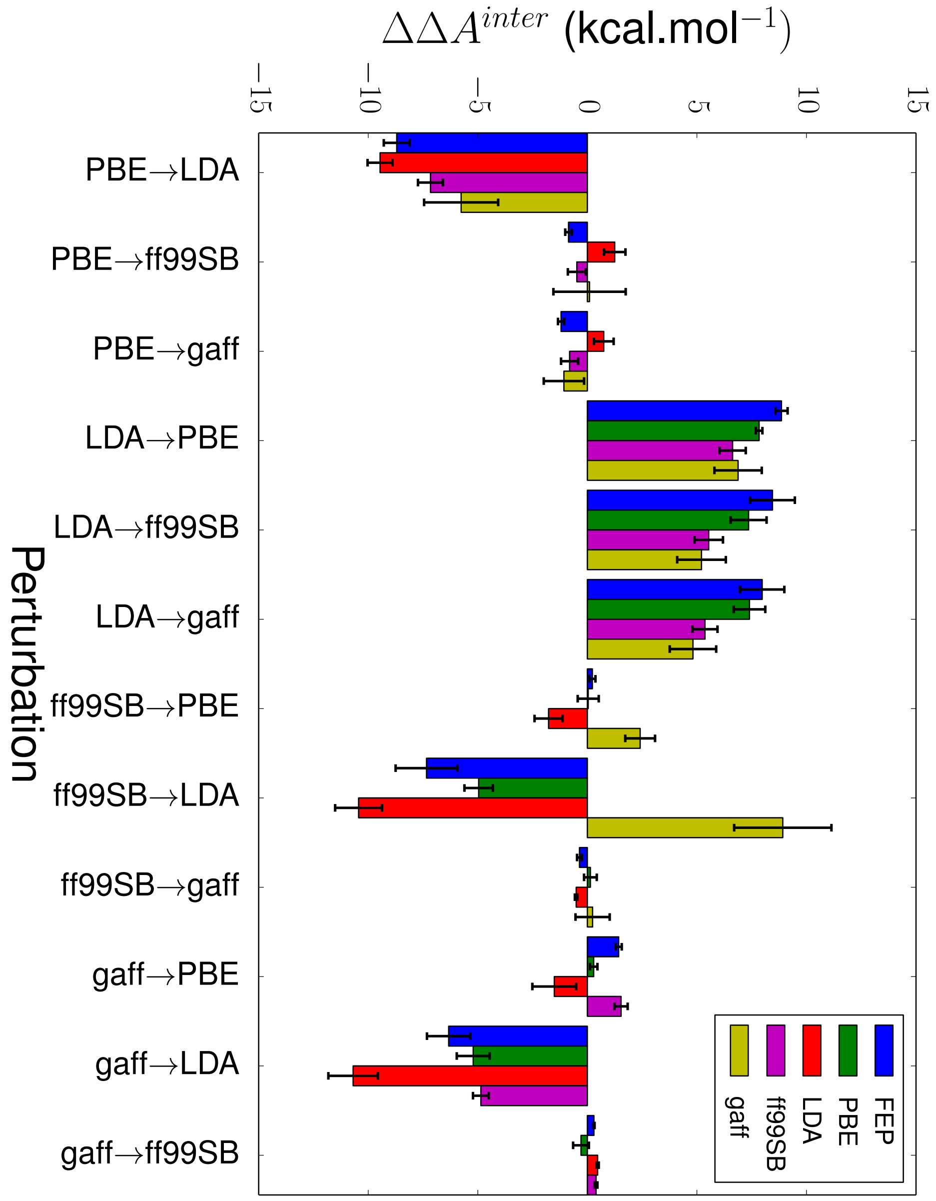

Figure 5.14.: Interaction free energy differences from total bias interaction energy approach. Bars are coloured according to the ensemble that was built **from**. The blue FEP data is taken from the results presented in chapter 3.

Method	MUE (kcal.mol^{-1})
Total energy bias	1.93
Interaction energy bias	1.05
Total energy resampling	1.56
Interaction energy resampling	1.04

Table 5.5.: Mean unsigned error for different hybrid methodologies compared with SSFEP.

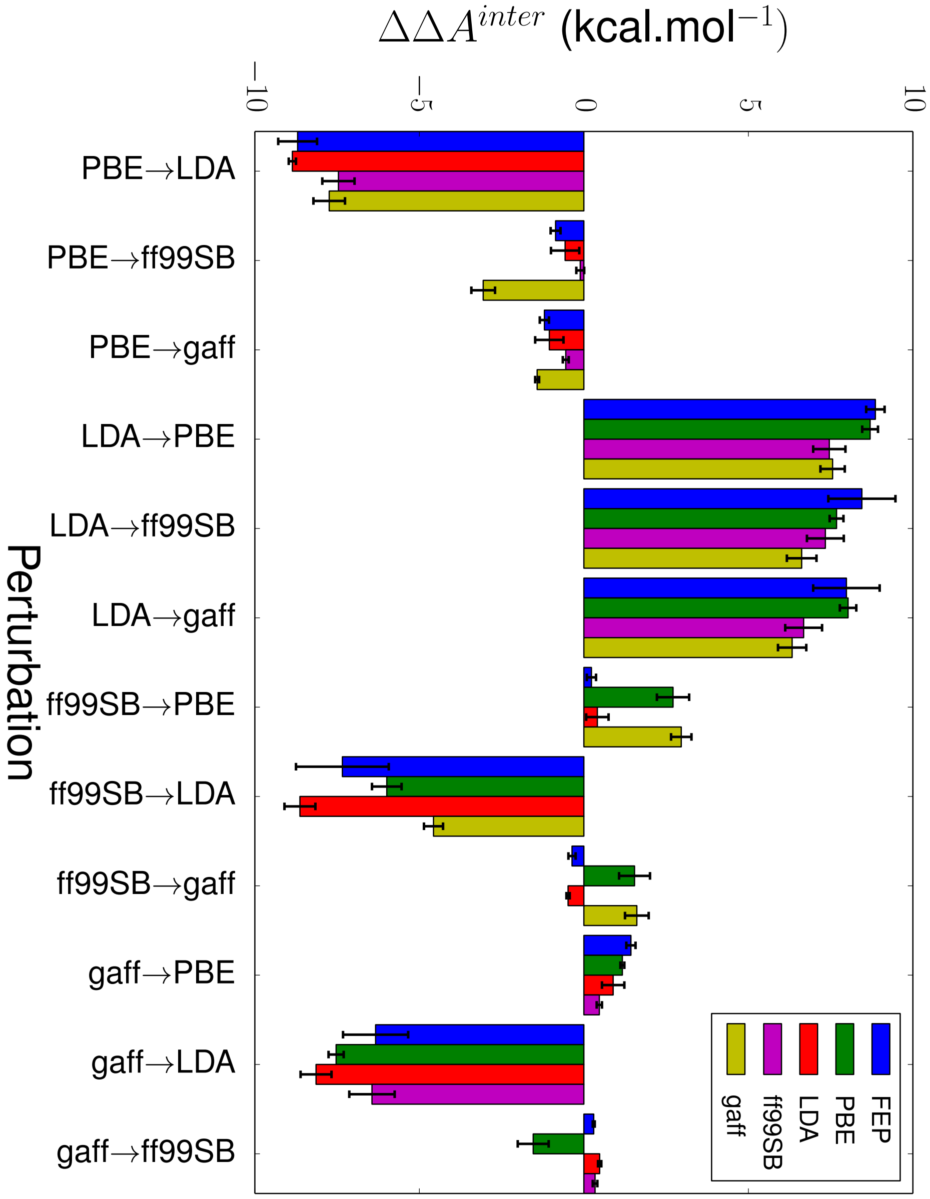

Figure 5.15.: Interaction free energy differences from the interaction bias approach. Bars are coloured according to the ensemble that was built **from**. The blue FEP data is taken from the results presented in chapter 3.

approaches provides very similar accuracy of around 1 kcal.mol^{-1} mean unsigned error. Moderately better performance is shown for total energy resampling than total energy bias results.

5.5. Discussion and Conclusions

The results presented here provide a practical demonstration and comparison of several state of the art hybrid free energy techniques. Results using these techniques are compared with the well converged SSFEP results presented in chapter 3. Although no improvements in the calculated free energy differences were shown, highly consistent results were found between all approaches.

We introduce a novel ensemble-building technique based on resampling equilibrium MM distributions and use of an MC acceptance criteria. We develop the theory and demonstrate its viability through model system calculations, and application to the base pair data. It shows performance comparable, and in some cases slightly superior, to the non-Boltzmann techniques also considered here. A key advantage provided by resampling is that it can be used purely as a post-processing technique, compared to MMBIF which actively uses QM computations to guide sampling. Considering the practical limitations implied by the guided sampling of MMBIF however, we conclude that, for sufficiently well behaved perturbations, there might be more drawbacks associated with its use than advantages. In this case using a purely post-processing derived techniques such as resampling or non-Boltzmann may be advantageous.

We extend the application of the presented techniques beyond the established literature by considering further usage of interaction energies. This is shown rigorously through theoretical analysis in the case non-Boltzmann and through demonstrative calculations with the model thermodynamic system for the ensemble building approaches. Post-processing the DNA base-pair MD data was considerably improved by the use of interaction energies in all cases, with considerably improved mean unsigned error when compared directly with SSFEP results. Whilst interaction energy application was successful for the data presented here, we note that the same criteria are required for accuracy with post-processing techniques as are considered in chapter 4.

The behaviour of the ensemble building techniques considered here is especially interesting to note in the context of a diagnostic test for the size of the interaction energy error as discussed in section 4.6. In particular, the use of these techniques with interaction energies is noted to produce an ensemble with hybrid properties i.e. composed of the intramolecular distributions of the references state but the interaction distributions of the target state. Recalling the results of section 4.6 we can write the size of the interaction energy error over a single λ-window as:

$$\Delta\Delta A_{hyd} - \Delta\Delta A_{hyd}^{inter} = -\frac{1}{\beta}\ln \frac{\dfrac{\int \exp\left[-\beta(U_{QM}^{intra} + U_{QM}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_{MM}^{intra} + U_{QM}^{inter})\right] d\mathbf{r}d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U_{QM}^{l}\right] d\mathbf{r}^{l}}{\int \exp\left[-\beta U_{MM}^{l}\right] d\mathbf{r}^{l}}} \tag{5.35}$$

Where we have simplified and adapted the notation for the context of the MM to QM perturbation. Notably, the partition function of the solvated MM state does not appear in this expression. Instead only the partition function of the target QM state and a hybrid state of MM and QM energy terms are present. The hybrid state partition function given here and the distributions generated through ensemble building with interaction energies are intriguingly similar. We would suggest therefore that use of ensemble building to generate information regarding this hybrid state is a promising avenue to pursue in the development of a viable diagnostic approach. Ensemble building with total energies is also able in principle to generate the target QM ensemble, and provide information for the diagnostic. The utility of MMBIF in the context of a diagnostic is limited by its need to influence sampling at the MM level, preventing its application as a post-processing approach. This condition is alleviated by the proposed and validated resampling approach.

6. CD1c

6.1. Introduction

One of the key challenges presented to the body's immune system, is the question of self or non-self[103]. How to identify those molecules or cells within the body that do not form a part of normal physiological function but mark the presence of a harmful pathogen or carcinogenic mutation. Immune cells patrol the blood stream and extracellular compartments and have mechanisms to inspect the surface of cells that they encounter. In order to locate cells that are playing host to pathogens or with harmful mutations an active mechanism is required, involving participation of the cell under inspection. This is fulfilled by the role of antigen-presentation pathways. These represent a specialised collection of proteins that sample the intramolecular composition of a cell and present the contents at the cell surface for inspection by the immune system. An antigen is any small molecule that elicits a response from the adaptive immune system[104].

The most commonly discussed and best understood antigen presentation pathways are associated with the major histocompatibility complex (MHC) proteins. These specialise in the presentation of peptide antigens from various intracellular compartments for inspection by cytotoxic T-cells and helper T-cells[105]. Far less well studied and the focus of this simulation work is the CD1 (Cluster of Differentiation 1) family of proteins, composed of the isoforms CD1a through to CD1e. In contrast to MHC pathways, CD1 isoforms present primarily lipid antigens (although presentation of some lipopeptides has been identified[106]), and activate different classes of T-cells[107]. Sitting at this position of the immune system therefore CD1 proteins make tempting therapeutic targets for modulation of immune responses to a wide variety of conditions. This may include amplifying the immune responses to cancers or pathogens, or suppressing activity for the treatment of autoimmune conditions.

The general structure of the CD1 protein family is shown in figure 6.1. The CD1 heavy chain is made up of the three domains $\alpha1$, $\alpha2$ and $\alpha3$. Similar to MHC proteins, the heavy chain forms a complex with β_2-microglobulin. The $\alpha1$ and $\alpha2$ domains cooperate to form the antigen binding domain; this consists of a flat base of six beta sheet strands and a pair of helices, one from each domain, that make up the sides and roof of the binding pockets[108]. The $\alpha2$ helix is broken by a short loop of residues. The CD1 main chain faces outwards from the cell surface and is able to interact with T-cell receptor (TCR) tetramers on the surface of different classes of T-cells[104].

The CD1 binding site, formed by the $\alpha 2$ and $\alpha 3$ domains, is separated into two volumes known as the F' and A' binding pockets with a connecting bridge. The structure of these regions show considerable variability between the known CD1 isoforms. This enhances the diversity of lipid presentation by allowing each isoform to specialise in presenting subclasses of lipid structures[108].

The aim of this work is to aid in the interpretation of X-ray crystallography data to produce a viable model, for the structure of the isoform CD1c. The data produced by X-ray crystallography comes in the form of a 3 dimensional maps that show the electron density within the system to the achieved resolution. The challenge is then to fit atoms within this electron density that correctly correspond to the structure of the crystal. The fitted atoms are therefore an interpretation of the data and can only be considered as a model for the true structure of the crystal. Fitting protein atoms within the structure is relatively easy for several reasons. The majority of proteins usually composed solely of the basic twenty amino acids, providing only a limited number of residues that can be fitted to the observed density. The availability of sequence data for proteins also provides detailed information that aids in fitting. Non-protein components of the system can prove more challenging however, as protein crystals are often prepared under conditions containing many different chemical species and in which impurities and degradation products may arise preventing exact knowledge of the composition of the crystallisation media[109].

Data from MD simulations can be used profitably to validate models derived from crystal structures. MD provides access to the dynamic properties of the system and can assess the stability of structure models and check for consistency with previously reported MD results[110,111]. In the following we present results from MD work with a series of CD1c structures produced by our collaborators Salah Mansour and Stephan Gadola of the Faculty of Medicine and Ivo Tews of the Centre for Biological Sciences.

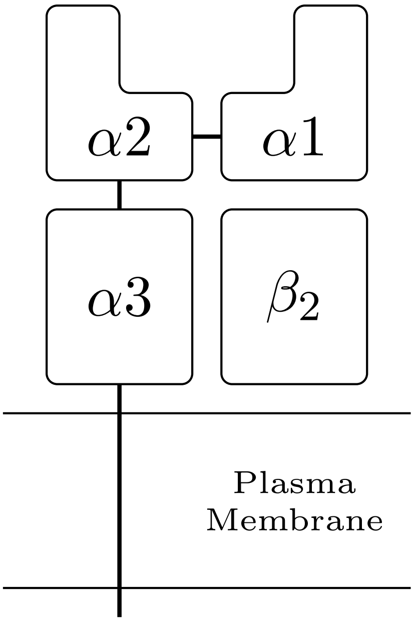

Figure 6.1.: Schematic figure of CD1 isoforms.

6.2. CD1c with PEG

The initial model provided contained three polyethylene glycol (PEG) molecules, two in the F' pocket and one in the A' pocket, along with a stearic acid molecule in the A' (figure 6.2). The structure of PEG is given by H-(O-CH_2-CH_2)$_n$-OH, where for the crystal structure n is 2 or 3. The model provided was based on an electron density map resolved to ~2.7 Å. To achieve crystallisation it proved necessary to produce a hybrid structure with the $\alpha3$ domain of CD1b grafted onto the $\alpha1$ and $\alpha2$ domains of CD1c. The PEG molecules are not naturally occurring biological compounds, but are a significant component of the buffer solution used in the crystallisation process.

A previous structure of CD1c has been published by Scharf *et al.*[108], showing the protein in complex with Mycobacterial lipid. This is available as the pdb structure 3OV6 from the protein databank and will be referred to as the CD1c MPM structure due to the presence of a complexed mannosyl-b1-phosphomycoketide primarily in the A' pocket. Comparison of the CD1c PEG and CD1c MPM two structures is shown in figure 6.3. Figure 6.3a shows an alignment of the first 165 residues (encompassing the entire binding site) of the CD1c chain, giving an RMSD of only 1.18 Å. This demonstrates a large degree of structural similarity between CD1c PEG and MPM.

An interesting feature highlighted by Scharf *et al.* in CD1c MPM is the presence of a F' binding groove with an open roof structure, in the place of a well defined F' binding pocket as seen in other CD1 isoforms. They hypothesise that this open structure enhances the diversity of antigens that CD1c is able to present, similar to MHC class proteins. Figure 6.3b shows the ligands occupying the binding sites of both structures. Occupation of the A' pocket is remarkably similar. The stearic

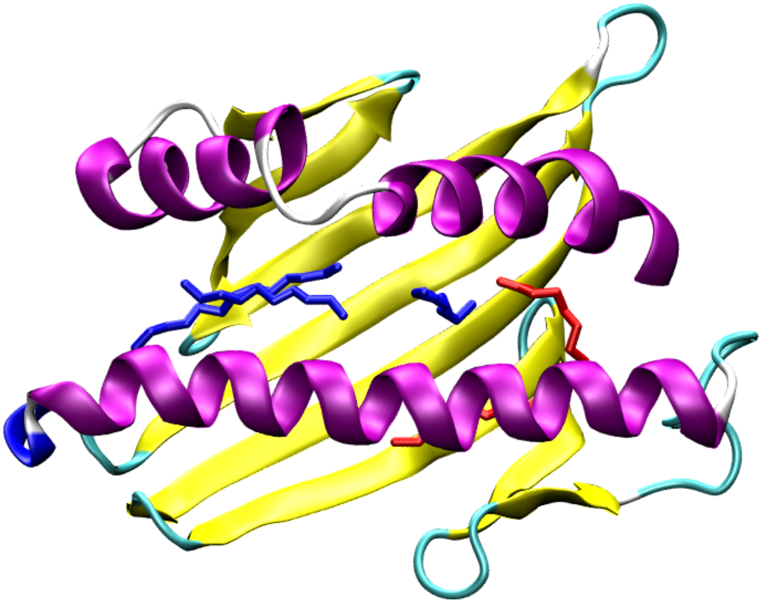

Figure 6.2.: Binding site of CD1c PEG model. PEG molecules are shown in blue, stearic acid in red.

acid and a PEG molecule largely recreate the tail of the MPM lipid. The contents of the F' pocket differ markedly however, with considerably more volume occupied by the PEG molecules than the C12 hydrocarbon chain of CD1c MPM. CD1c PEG therefore demonstrates more fully the binding capacity of the F' site.

The larger volume of ligands in the F' site suggests that CD1c PEG may be exploiting the open F' groove structure. Comparison of the roof residues of the F' site between the two structures suggests otherwise however (figure 6.3c). Rather than the binding groove observed in CD1c MPM and despite the fuller occupation of the F' binding site, CD1c PEG shows a closed roof structure similar to other CD1 isoforms. This F' pocket roof is comprised of the residues TYR152, LEU147, GLU80 and HIS84 that connect the $\alpha1$ and $\alpha2$ helices in a tethered configuration.

Given the differing occupation of the F' pocket it is surprising that the position of the helices agree so well between the structures. How is the highly hydrophobic F' pocket being held open in the CD1c MPM structure? Although Scharf *et al.* describe the fitted C12 chain as 'conservative' there is no suggestion of unaccounted for electron density that would appropriately fill the pocket. Instead it has been suggested that the open conformations of CD1c MPM and crystal structures of other CD1 isoforms are maintained by crystal packing effects[110].

Fitting of PEG to the binding site was carried out on the basis of their excellently matching the observable electron density, and their presence during the crystallisation process. Whilst partially hydrophobic, PEG is also soluble in water and has an overall amphipathic character[112]. This raises questions about its ability to form stable complexes with the highly hydrophobic binding pocket of CD1c. We here undertake MD simulation work of the proposed CD1c PEG structure to assess the stability of the complex and the suitability of PEG as the fitted ligand. In addition we carry out simulations with the CD1c binding site emptied of all ligands and, to examine the effect of ligands with a more hydrophobic character, we use a system employing aliphatic chains in the place of PEG.

6.2.1. System Setup

The first requirement for simulation with the protein structure is the addition of hydrogen atoms that are not have not been resolved within the electron density. This was considered in some detail with comparison of three different protonation software packages (for this analysis see appendix G). For all simulation work presented here, the structure produced by the PROTANATE3D module of the MOE software package[113] was used. Crystallographic waters were retained within the structure and additional solvent was added to the system, with TIP3P water molecules placed in a box such that the minimum distance from protein to box edge was 8 Å. Sodium ions were added to neutralise the overall charge of the system, giving a total system size of ∼50,000 atoms. Three different simulation setups were created. The all PEG setup: the provided crystal structure with all three PEG molecules and stearic acid present. The empty setup: all PEG molecules and stearic acid removed from the binding site leaving it empty. The aliphatic chain setup: all oxygen atoms in PEG molecules were replaced with CH_2 groups to produce purely aliphatic carbon chains.

Simulations were carried out using the AMBER 12 package[80], with the FF99SB force field[43]. Ligands were parameterised with the GAFF force field[48] using the ANTECHAMBER module with partial atomic charges assigned using the AM1-bcc method[86,87]. GAFF atom types and charges were checked manually for consistency with no changes made.

6.2.2. Simulation Protocol and Analysis

The equilibration protocol included a series of successive minimisations, gradually releasing restraints on the heavy atoms of the system. Heavy atoms of the protein were then restrained and the system was gradually heated to 300 K over 200 ps, and equilibrated to the correct pressure for 200 ps under the isothermal-isobaric ensemble. The system was then cooled and the procedure repeated with protein restraints removed.

Temperature control was achieved using the Langevin thermostat[58], with a collision constant of 3.0 ps^{-1}. Pressure regulation employed the Berendsen barostat[114] with a

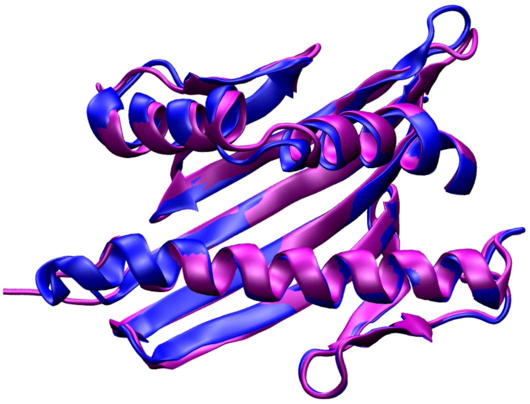

(a) Binding site protein structure

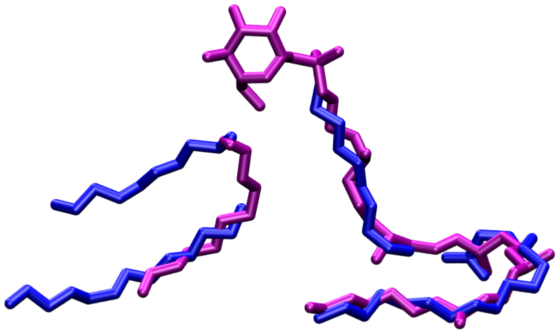

(b) CD1c ligands. Those to the right occupy the A' pocket, to the left the F' pocket.

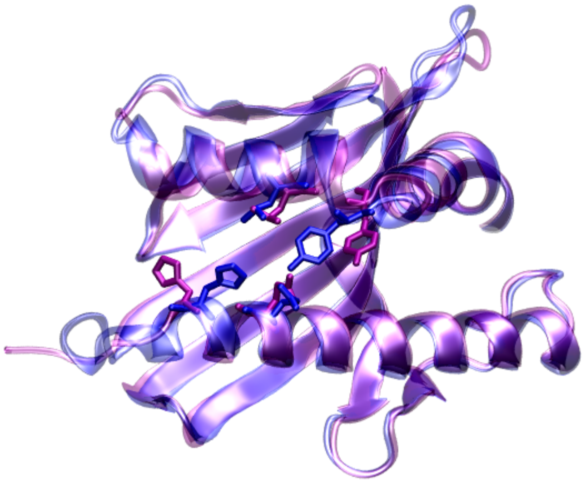

(c) F' pocket roof.

Figure 6.3.: Comparison of CD1c PEG and MPM structures. CD1c MPM is shown in purple, whilst CD1c PEG is shown in blue. Structures were aligned to the backbone of the residue range 5-170 of the CD1c main chain.

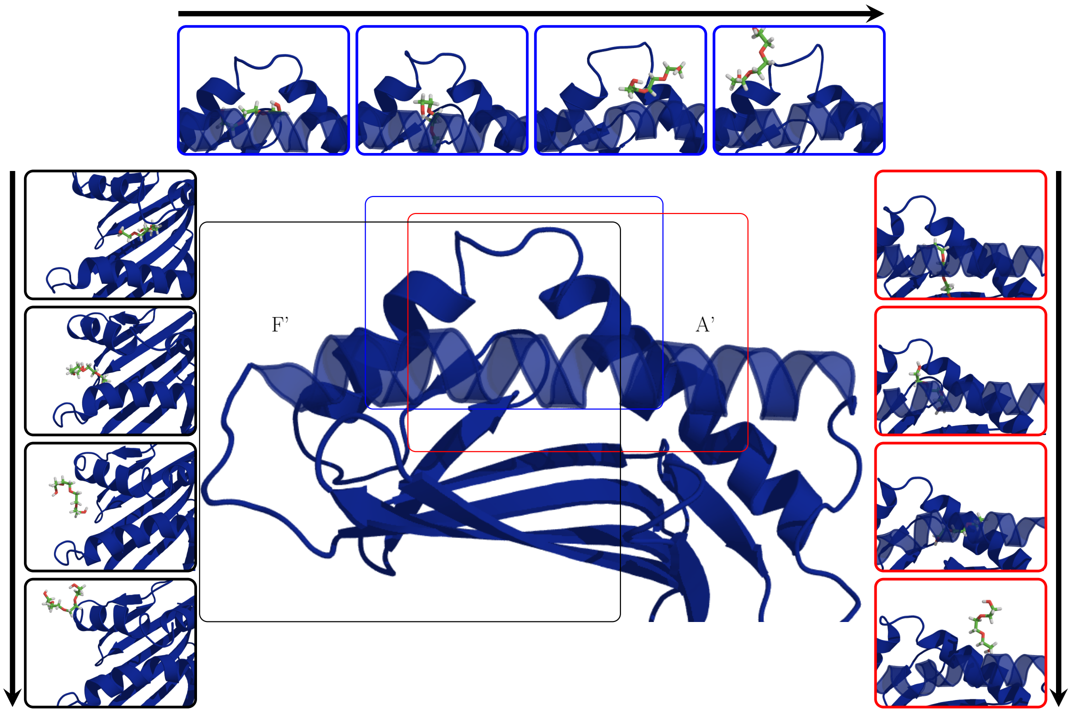

Figure 6.4.: Exit of PEG molecules from CD1c binding sites. The series of stills show exit of PEG molecules through each pathway, with the relevant region highlighted on the central diagram. Black - F' pocket side exit, blue - F' pocket roof exit, red - A' roof exit.

relaxation time of 2.0 ps. Dynamics were carried out using a 2 fs timestep and a cut off of 8 Å for non-bonded interactions. Covalent bonds were constrained through use of the SHAKE algorithm[68]. Production runs were carried out on the Emerald GPU cluster, using the CUDA implementation of PMEMD[115,116]. Production runs were carried out under constant volume and temperature dynamics. Three repeats of each simulation setup were carried out to give a total of nine MD trajectories, each 200 ns in length.

Simulation data was post-processed using the PTRAJ executable of the AMBER software package for calculation of RMSD values. Pocket volume data was produced using the PocketAnalyzer package[117], with minimum number of neighbours set to 15 and degree of buriedness set as 11. This represents the default settings. Additionally, VAL96 was specified as the only residue of interest for the pocket finding algorithm.

6.2.3. Results

The most striking observation that can be made from the simulation data is the rapid exit of the PEG molecules from both the F' and A' binding sites. Three different locations were identified where PEG molecules were able to leave the binding site, two in the F' pocket and one in the A' pocket. These are shown in figure 6.4 along with trajectories taken by PEGs as they leave the binding site. Of the three repeat

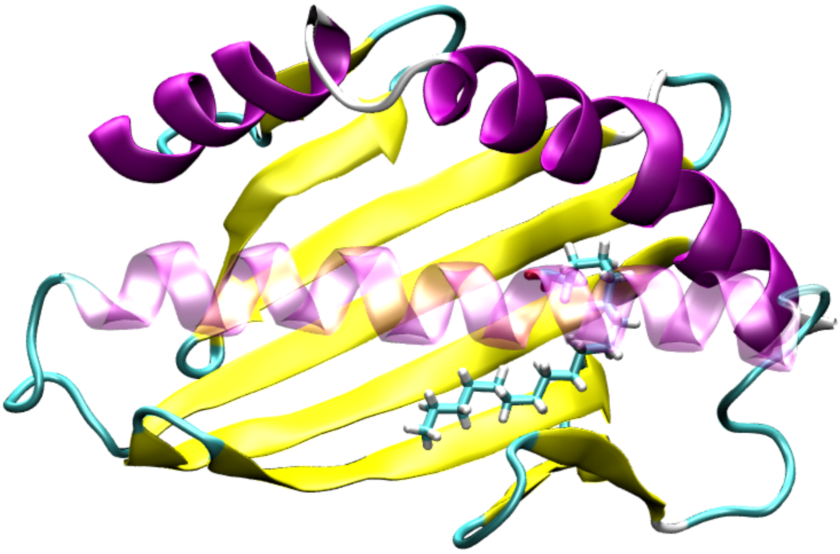

(a) 0 ns

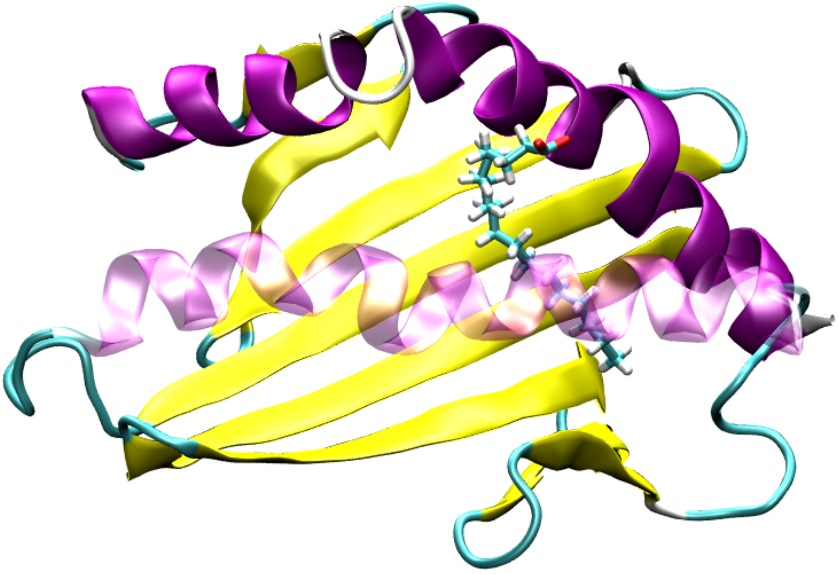

(b) 200 ns

Figure 6.5.: Starting and final configurations of stearic acid within the A' binding pocket.

simulations, two lose all PEGs within the 200 ns time scales simulated. For the third, one PEG molecule remains in the F' pocket by the end.

Within all simulations the stearic acid molecule remains within the binding site, although it is displaced significantly from its initial position. The crystal structure places the aliphatic chain of the stearic acid into a groove beneath the $\alpha 1$ helix, whilst the head group is buried, pointing 'downwards' into the A' pocket. Rather than remaining buried, the acid head group seeks solvent through the roof of the A' pocket; this is perhaps unsurprising given the lack of a polar interactions within the binding site to stabilise the position. The starting and final configurations of the stearic acid are shown in figure 6.5.

In a previous MD study of CD1 isoforms, by Garzón *et al.* working with the CD1c MPM structure, it was noted that hydrophobic collapse of the binding pocket occurs in the absence of ligand. Our MD results with the CD1c PEG structure confirm this observation. In all empty runs the inward collapse of helices $\alpha 1$ and $\alpha 2$ was observed as shown in figure 6.8. This leads directly to the loss of binding pocket volume (figure 6.6). Despite the restraints placed on protein atoms during the majority of the equilibration process, by the start of the production phase, the empty simulation runs already demonstrate much smaller pocket volumes than the PEG simulations. The pockets rapidly collapses further to around zero, with PocketAnalyzer[117] often unable to locate a binding pocket. Interestingly, the all-PEG simulations show much a more progressive collapse as PEG molecules evacuate the binding site. Loss in binding pocket volume is strongly correlated with the exit of PEG molecules from the binding site. One of the PEG simulations maintains the binding volume of the pocket; this is the same trajectory that retains a PEG molecule. It is interesting to note therefore that the presence of even a single PEG in the binding site maintains the structure of the pocket.

As observed above the most prominent difference between the CD1c MPM structure and CD1c PEG is the arrangement of the four roof residues of the F' pocket. As these are hypothesised to relate to the binding state of the protein, special attention was focussed on these residues. Figure 6.7 shows RMSD values for the sidechain heavy atoms of these residues against the starting configuration of each trajectory. For all simulation setups the roof residues remain in the tethered configuration at the start of the production runs. Large increases in RMSD correspond to loss of the tethered configuration to give structures that resemble more closely those of CD1c MPM (figure 6.9). For the empty simulations (figure 6.7b) loss of the tethered configuration is immediate, likely due to the hydrophobic collapse of the pocket disrupting the required binding geometries. For the PEG simulations (figure 6.7a) in accordance with the pocket volume data, loss of tethering is more gradual. The second simulation (green line), where pocket collapse does not occur, after an initial loss, recovers the tethering interactions. This may provide a rationale for the persistence of the

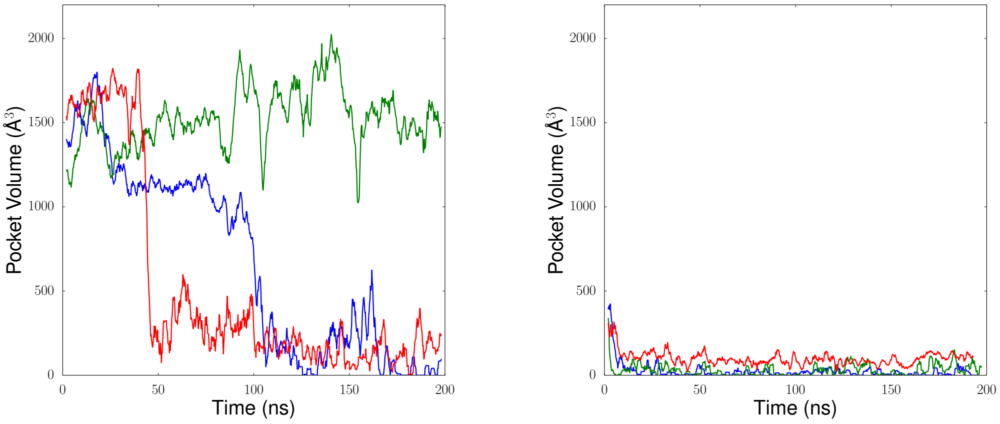

(a) All PEGs

(b) Empty

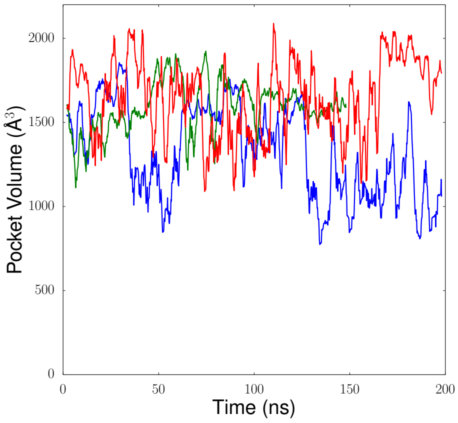

(c) Aliphatic Chains

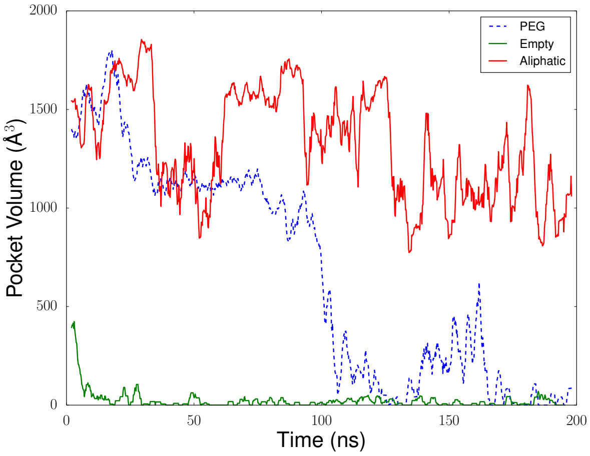

(d) Comparison a single trajectory from each setup

Figure 6.6.: Pocket volumes throughout MD trajectories for PEG simulation setups. Results shown are moving averages over a window of 2 ps.

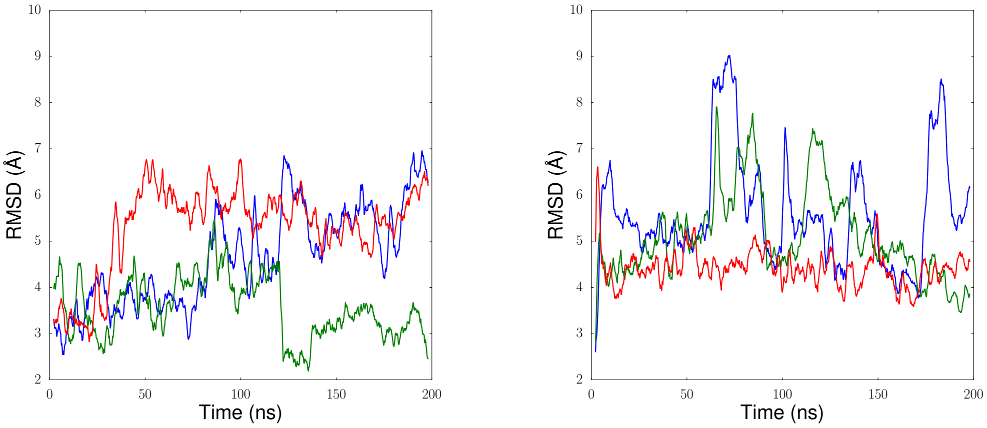

(a) All PEGs

(b) Empty

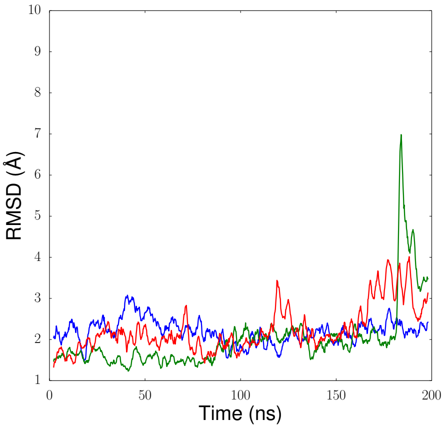

(c) Aliphatic Chains

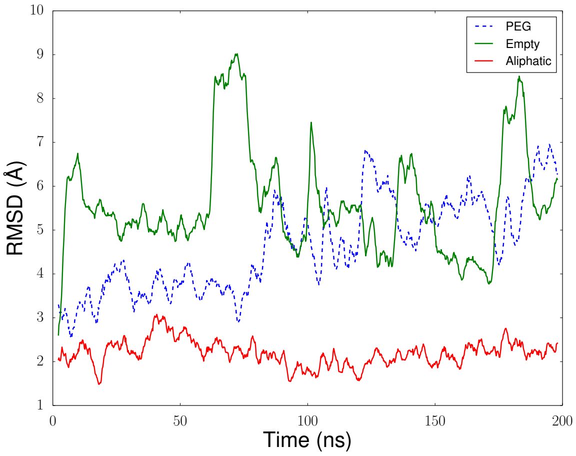

(d) Comparison of a single trajectory from each setup

Figure 6.7.: RMSD of roof residues versus starting frame of simulation. Structures were aligned to the beta sheets that make up the base of the binding site of starting structure of each trajectory and RMSDs were calculated with no additional fitting.

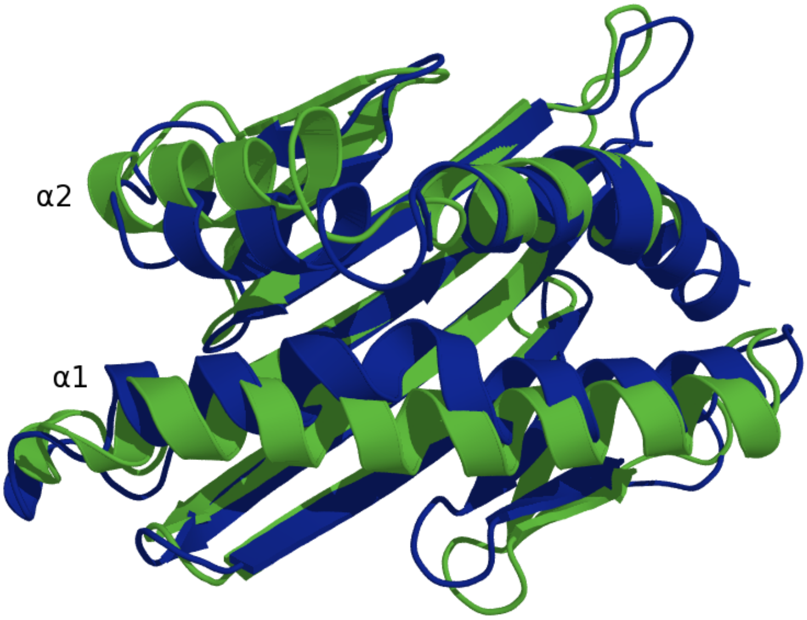

Figure 6.8.: Comparison of CD1c PEG (green) versus a configuration from MD trajectory with empty pockets (blue). Significant displacement of the $\alpha 2$ and $\alpha 1$ helix leads to collapse of the F pocket.

PEGs in this simulation with closure of the roof preventing a route of exit. Observing the eventual collapse of the other repeats however, it is likely that this is only a temporary effect.

The observed instability of PEG as a ligand of CD1c, suggests strongly that these are unlikely to be the ligands present in the molecular crystal. Owing to the presence of the oxygen groups, PEG is a moderately polar species that does not seem to reside for long periods of time in the purely hydrophobic binding site. For this reason it was decided to modify the CD1c structure to contain aliphatic carbon chains in the place of PEG through the introduction of CH_2 groups in the place of oxygen. This was not designed necessarily to represent plausible ligands that might be present in the crystal structure but instead to produce ligands of a comparable volume to PEG but with a more hydrophobic character.

Results of the simulations with aliphatic chains are given in figures 6.6c and 6.7c. The complexes produced by these simulations appear much more stable than those with the PEG molecules. Pocket volumes for all simulations remain large and consistent the crystal structure. Importantly the tethering interactions of the roof are also maintained strikingly well. This shows that stability of the roof structure is dependent on occupation of the binding site by ligands of the correct size and physicochemical character. This is interesting to note in the context of the previous observation that the CD1c MPM structure does not demonstrate the full binding capacity of the CD1c F' pocket.

In the PEG simulations the stearic acid group undergoes significant rearrangement to allow access of the charged head group to solvent. In the aliphatic chain simulations however, the stable presence of the aliphatic chain within the A' pocket prevents access through the A' roof in two of three simulation repeats. In these cases the head groups emerge through a portal beneath the $\alpha 2$ helix. This portal is composed of the residues PHE126, TYR160, LEU163 and ARG164; a positive charge is provided by ARG164 that interacts with the lipid head group. To our knowledge this portal has not been previously observed in any simulation work or crystal structures relating to CD1 isoforms and is not present in the crystal structure, suggesting that it may arise transiently.

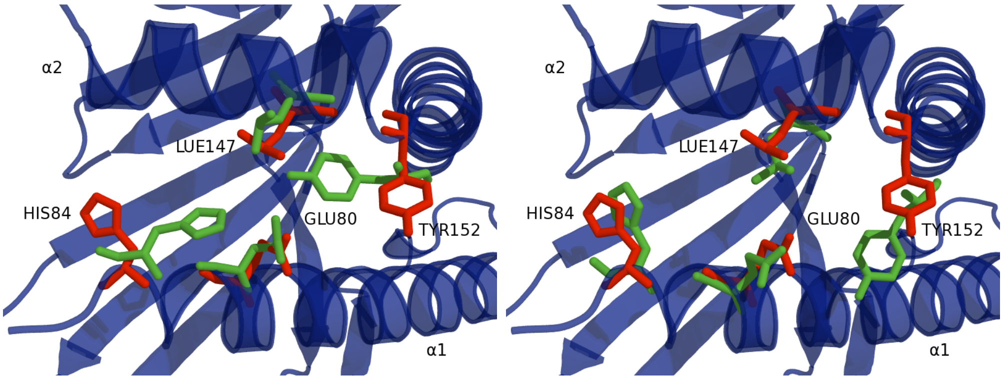

(a) 0 ns

(b) 40 ns

Figure 6.9.: Comparison of F' roof residues from trajectory of CD1c PEG against the CD1c MPM crystal structure. Roof residues of CD1c MPM in red. Configuration first CD1c PEG trajectory shown in green. At 0 ns MD configuration (a) shows the tethered arrangement of the roof residues as seen in CD1c PEG crystal structure. As the simulation progresses, at 40 ns (b) the structure of the F' roof has been lost, giving configurations that more closely resemble the CD1c MPM crystal structure. The considerable flexibility of residues His84 and Tyr152 gives a wide range of configurations for the roof residues, with the above image being selected for its similarity to CD1c MPM.

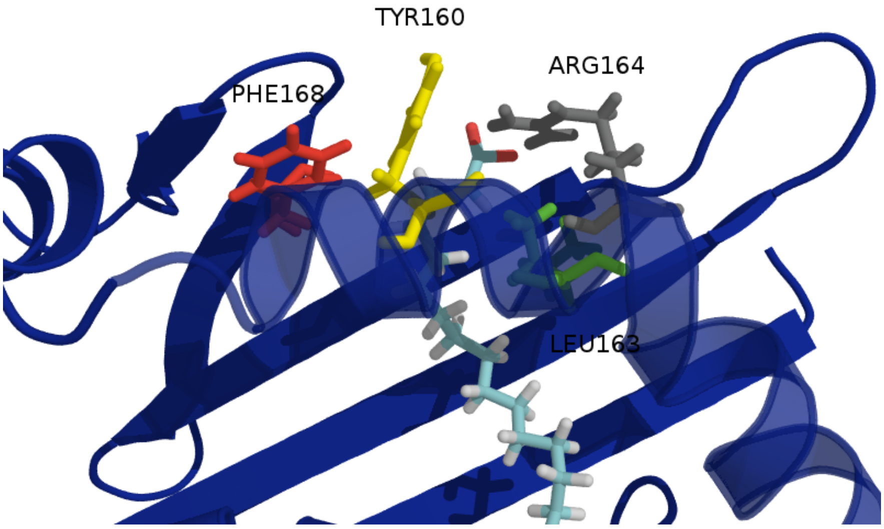

Figure 6.10.: stearic acid emerging through a side portal of the A' pocket comprised of residues PHE126, TYR160, LEU163 and ARG164.

6.3. Crystal Structure with Lauric Acid

On the basis of the simulation work with CD1c PEG, and given the impure nature of crystallisation media it was considered that PEG was a poor choice for the fitted ligand. For this reason the diffraction data were subjected to another round of refinements by our collaborators. This produced a second structure here termed CD1c LAU. This was based on a different crystal than CD1c PEG, although produced under identical conditions, with an electron density map of ~2.8 Å resolution and a unit cell comprising two copies of CD1c with minor differences in sequence. The binding site and fitted ligands for both structures were close to identical. In place of the three PEG molecules, two lauric acid molecules were fitted in the F' binding pocket, with only the stearic acid in the A' pocket. The new configuration of the stearic acid places the head group facing upwards towards the roof of the A' pocket, as opposed to the CD1c PEG where it is buried facing down into the pocket.

The system set up procedure used was the same as given in section 6.2.1. Of the two CD1c structures within the unit cell one was chosen on an arbitrary basis for simulation work due to their similarity.

6.3.1. Results

The pocket volume and roof residue RMSD results for all three lauric acid simulations are given in figure 6.12. Encouragingly these are very similar to the results

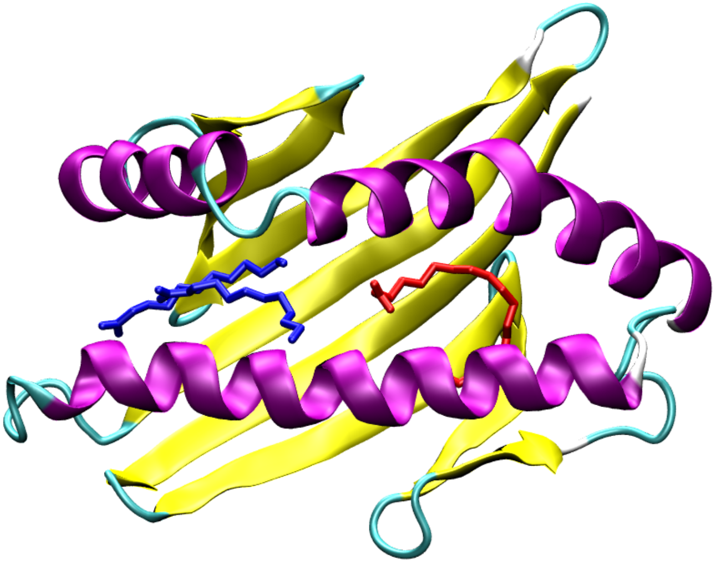

Figure 6.11.: CD1c LAU crystal structure. Lauric acid molecules are shown in blue, stearic acid in red.

provided by the aliphatic chain simulations. Comparison with the other simulation setups is given by figures 6.13 and 6.14. Stable complexes are observed throughout the time scales simulated. Pocket volumes fluctuate less than the aliphatic chain simulations suggesting that the pocket is more stable in the presence of the lauric acid.

Starting and ending configurations from a production run are shown in figure 6.15. Some movement of the ligands away from their position within the crystal structure is seen although generally agreement is good. Notably in the 200 ns configuration the stearic extends out of the A' pocket into the F'; this likely provides a route for the head group to access solvent without disrupting the roof residues of the A' pocket.

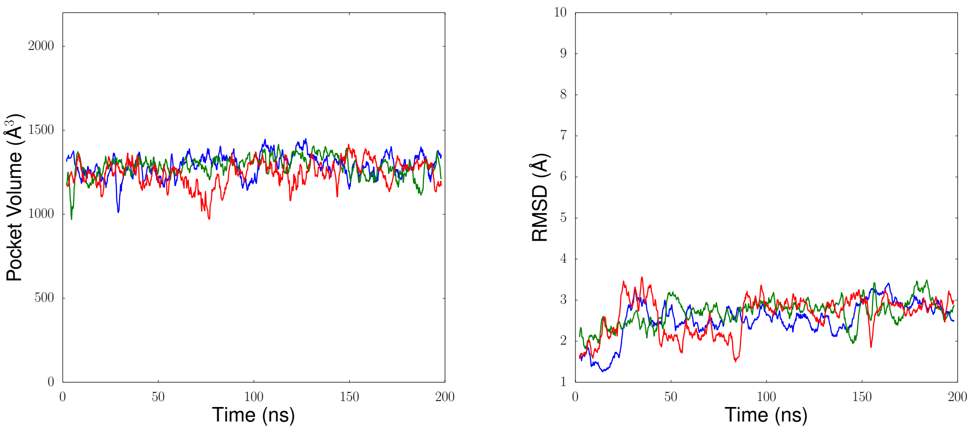

(a) Pocket Volume (b) Roof Residue RMSDs

Figure 6.12.: Lauric acid simulation data. Pocket volumes and RMSDs are calculated as discussed previously in section 6.2 and figure 6.7

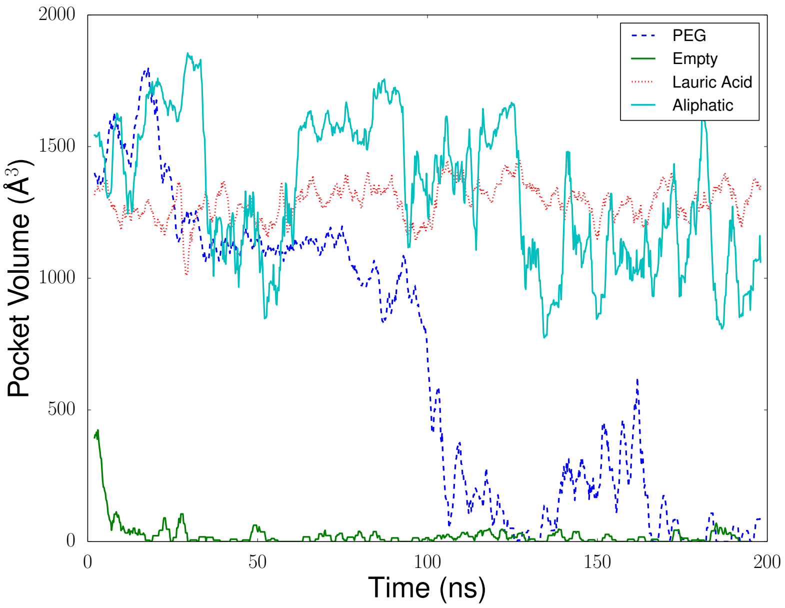

Figure 6.13.: Pocket volumes throughout MD trajectories for all simulation setups. Results shown are moving averages over a window of 2 ps from representative trajectories.

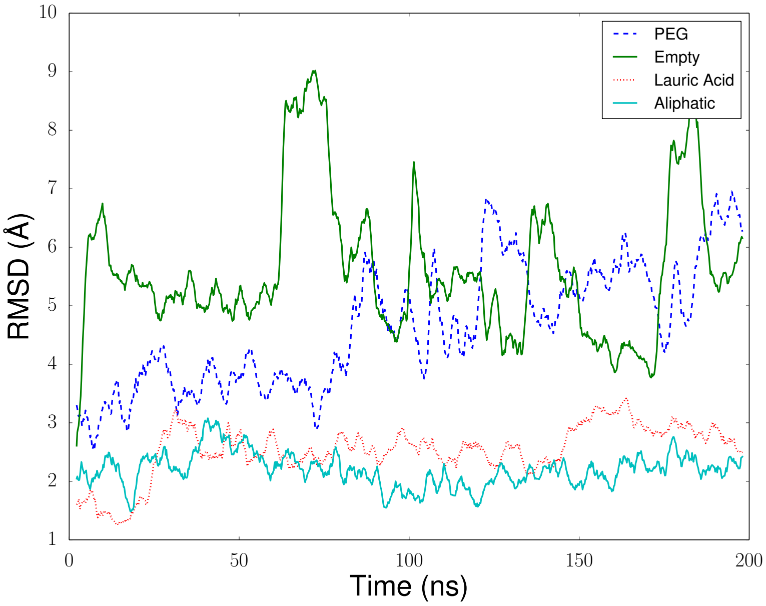

Figure 6.14.: RMSD of F' roof residues versus their starting tethered configuration for all simulation setups. Results shown are moving averages over a window of 2 ps from representative trajectories.

6.4. Discussion and Conclusions

As discussed previously, the CD1 protein family isoforms have an important role in modulating the response of the immune system. CD1c has been shown to present mycobaterial antigens[108,118] as well as lipopeptides[106] and a selection of self-lipids[107]. In particular, tuberculosis is a member of the mycobacterium genus and a notable target for therapeutic intervention. The results presented here have been successfully used to aid the interpretation of X-ray diffraction data to develop an atomic resolution model of CD1c. The generated structure containing lauric and stearic acid species produces a stable protein-ligand complex over the 200 ns time scales simulated.

The initially proposed model with PEG molecules fitted within the binding site was not stable. Rapid loss of PEG molecules occurred in all simulations closely followed by loss of binding pocket volume. Additional simulations were performed with aliphatic chains of equivalent size modelled in the place of PEG, producing stable complexes maintaining the F' roof. It was concluded that PEG does not have a sufficiently hydrophobic character to complex stably within the non-polar pockets of CD1c. This information was used to carry out additional refinements and produce the final CD1c LAU structure. Further simulation work demonstrated the stability of this structure, the maintenance of the F' roof interactions.

These findings contribute to the general understanding of antigen presentation by

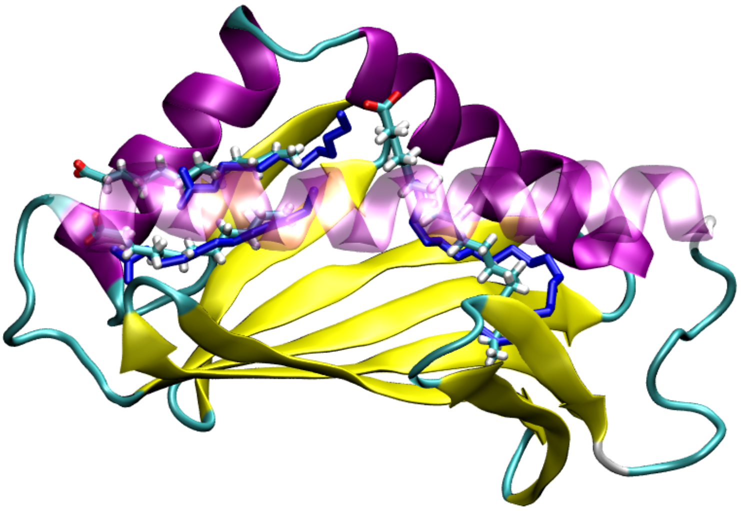

(a) 0 ns

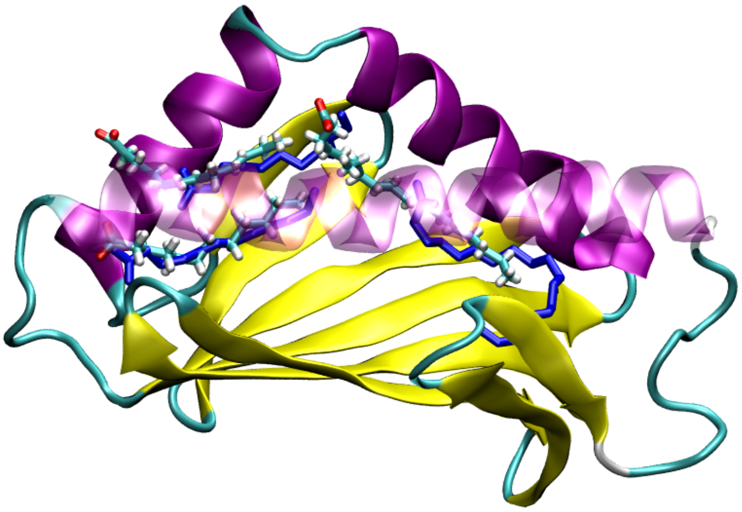

(b) 200 ns

Figure 6.15.: Start and ending configurations for lauric acids from a simulation production run. The ligands shown in blue are the configurations of the crystal structure, along with the visible protein structure.

CD1c and the key role of the F' roof in this regard. As well as acting as a learning exercise in the preparation of protein structures and MD work with large systems. This work is also intended as preparatory work for the application of hybrid free energy techniques to the CD1d isoform as part of an ongoing BBSRC collaboration. This is discussed in more detail in chapter 8.

Comparison of the CD1c MPM and PEG structures gave rise to the suggestion that the MPM structure may represent an intermediate structure that has not adopted its final conformation for TCR tetramer binding. Such a suggestion is interesting in the context of the results presented here. We demonstrate the requirement for the binding site to be filled with ligand of the correct physico-chemical character in order to maintain the interactions of the roof residues. From the data presented here it is unclear if the loss of the F' roof is a necessary prerequisite for collapse of the binding site, or whether the collapse in the absence of ligand causes loss of the roof interactions. Stable complexes of CD1c MPM were seen by Garzón *et al.*[110] over timescales of 100 ns; however no analysis of the F' pocket roof was carried out. Without further work it is not possible to rule out CD1c MPM producing a stable F' roof under simulation, rather than crystallographic, conditions. Strong evidence for the open roof structure as a folding intermediate would require experimental demonstration of the failure of such a structure to elicit a T-cell response in cellular assays.

7. Conclusions

The aim of this work has been to provide important validation work, both practical and theoretical, in the application of hybrid free energy techniques. We present a valuable demonstration of the convergence of the widely used SSFEP approach, with respect to not only the MM to QM, but also the reverse QM to MM perturbation. Although this does not guarantee appropriate convergence of the SSFEP approach in all cases, it represents an important proof of principle for a representative system of biologically relevant complexity.

The convergence of calculations with SSFEP is shown to be dependent on the use of interaction energies in the place of total energies. We demonstrate that is due to poor overlaps between the MM and QM probability distributions, and that overlaps when considering interaction energies are improved by several orders of magnitude. Unfortunately this also leads to a loss of theoretical rigour as free energy estimators are derived using total energies.

Given the seeming practical necessity of the use of interaction energies we proceed with a detailed theoretical analysis, supported with a series of calculations in model systems. We demonstrate that the use of interaction energies is inherently approximate, but note a limit in which interaction and total energies are equivalent. This limit is given by a system in which the intra and intermolecular degrees of the system can be considered independent. This is obviously not generally the case for biologically relevant systems, and we demonstrate through direct comparison of total and interaction free energy differences in a protein-ligand system that considerable differences can arise between the two, although this varies on a case by case basis. The prospect of designing a diagnostic test predict the size of the error associated with the use of interaction energies is discussed and some further theoretical and practical observations are made to that end.

In addition to SSFEP we go on to consider the application of a pair of more recently developed hybrid free energy techniques, in the form of non-Boltzmann reweighting and MMBIF. We consider the validation of these procedures through their application to the model system as used in chapter 4 and the base pair data of chapter 3. Additionally we present a novel technique closely related to MMBIF, based upon resampling of classical distributions to derive QM distributions through use of a QM acceptance test. This approach is demonstrated to give behaviour equivalent to MMBIF but can be used purely as a post-processing technique. We extend all of these techniques to consider their use with interaction energies and demonstrate improved performance in application to the DNA base pair data.

Finally, as preparatory work for future application of hybrid free energy approaches, we carry out setup and molecular dynamics simulations of an isoform of the CD1 protein family. Working in collaboration with other researchers we carry out validation of a proposed crystal structure for CD1c.

8. Future Work

As highlighted throughout this work, the inherent approximation engendered by the use of interaction energies, gives rise to errors that vary on a system by system basis (see section 4.4). Although we present some valuable theoretical observations for the development of a diagnostic approach, capable of estimating the error for each perturbation, considerable additional work will be required to create a viable procedure that meets the criteria detailed in section 4.6. Approaches are suggested by the results presented in chapter 5, exploiting properties of the proposed MC resampling technique.

A striking improvement in convergence of free energy calculations with the use of interaction energies is demonstrated in chapters 3 and 5. It would be highly profitable therefore to expand the use of interaction energies beyond hybrid free energy work, to include their application in alchemical perturbations with purely MM calculations and develop their use as a standard approximation in the field of free energy calculations. This would compromise one of the strengths of free energy estimation as a theoretically rigorous approach. However, combined with an effective diagnostic approach capable at least of highlighting cases that should be extended with further simulation work to converge total free energy differences, considerable computational savings are potentially available. It would be interesting to consider the different convergence characteristics of total and interaction energies with, for instance, the COX2 calculations presented in chapter 4. Savings may be possible, not only in the number of required MC moves, but in the required number of λ windows.

Although the results for MMBIF within the model system are qualitatively enlightening, it would be informative to consider application of the MMBIF methodology to the DNA base pair system for comparison with the directly calculated SSFEP results. This was not carried out here due to the guided sampling of the approach, requiring considerable additional calculations.

A major advantage of the techniques considered in chapter 5 is that they enable the use of more sophisticated free energy difference estimators than the Zwanzig equation i.e. TI or BAR. We limit ourselves to application of the Zwanzig equation however, for comparison to the SSFEP base pair data. Due to the extensive QM MD data generated, it would be interesting to compare application of BAR to these data natively, as well in the context of the reweighting and ensemble building techniques.

It is our intention to consider application of the hybrid methodologies considered here to the development of fluorinated derivatives of the ligand KRN7000 for the C1d isoform, CD1d. This system is the subject of an ongoing BBSRC collaboration at the University of Southampton. We anticipate hybrid free energy techniques to be valuable in this context due to proposed fluorine-carbonyl interactions between the ligand and protein backbone. This interaction is poorly described by classical energy models and we therefore hypothesise that purely MM approaches will be unable to appropriately capture the free energy differences between KRN7000 derivatives. This represents an opportunity to demonstrate the effectiveness of hybrid free energy techniques within a system of interest for clinical developments and not adequately handled by classical theory.

The data presented here cover a number of different approaches to the calculation hybrid free energy differences. Our focus has been the validation and extension of these rather than direct comparison of their merits and efficiency. Further application of the non-Boltzmann and ensemble building techniques, is required to assess their practical viability. In particular these demonstrate superior convergence behaviour with extension to the use interaction energies and we have not here considered the advantages of using estimators like BAR and TI. The ability of MMBIF to bias sampling of the reference potential to the 'QM' phase space is compelling. However, this procedure is shown to fail even in the limited test case considered by Woods *et al.* On the basis of the validation presented in chapter 3 and its being the most extensively applied throughout the literature, we would consider the use of SSFEP to be the most rigorous and robust method for the calculation of free energy differences at the QM level at this time, although more recent approaches show considerable promise.

Bibliography

1. Michel, J. & Essex, J. W. Prediction of protein-ligand binding affinity by free energy simulations: Assumptions, pitfalls and expectations. *J. Comput.-Aided Mol. Design* **24,** 639–658 (2010).
2. Zwanzig, R. W. High-temperature equation of state by a perturbation method. I. Nonpolar gases. *J. Chem. Phys.* **22,** 1420–1426 (1954).
3. Bennett, C. H. Efficent estimation of free energy differences from Monte Carlo data. *J. Comput. Phys.* **22,** 245–268 (1976).
4. Kirkwood, J. G. Statistical mechanics of fluid mixtures. *J. Chem. Phys.* **3,** 300–313 (1935).
5. Warshel, A. & Levitt, M. Theoretical studies of enzymic reactions: Dielectric, electrostatic and steric stabilization of the carbonium ion in the reaction of lysozyme. *J. Mol. Biol.* **103,** 227–249 (1976).
6. Beierlein, F. R., Michel, J. & Essex, J. W. A simple QM/MM approach for capturing polarization effects in protein-ligand binding free energy calculations. *J. Phys. Chem. B* **115,** 4911–4926 (2011).
7. Woods, C. J., Manby, F. R. & Mulholland, A. J. An efficient method for the calculation of quantum mechanics/molecular mechanics free energies. *J. Chem. Phys.* **128,** 014109 (2008).
8. König, G., Hudson, P. S., Boresch, S. & Woodcock, H. L. Multiscale free energy simulations: An efficient method for connecting classical MD simulations to QM or QM/MM free energies using non-Boltzmann Bennett reweighting schemes. *J. Chem. Theory Comput.* **10,** 14061419 (2014).
9. Fox, S. J. *et al.* Free energies of binding from large-scale first-principles quantum mechanical calculations: Application to ligand hydration energies. *J. Phys. Chem. B* **117 (32),** 9478–9485 (2013).
10. Vaidehi, N., Wesolowski, T. A. & Warshel, A. *Ab initio* pseudopotential free-energy perturbation approach. *J. Chem. Phys.* **97,** 4264–4271 (1992).
11. Liu, W., Sakane, S., Wood, R. H. & Doren, D. J. The hydration free energy of aqueous Na+ and Cl- at high temperatures predicted by ab initio/classical free energy perturbation:973 K with 0.535 g/cm^3 and 573 K with 0.725 g/cm^3. *J. Phys. Chem. A* **106,** 1409–1418 (2002).
12. Rod, T. H., Rydberg, P. & Ryde, U. Implicit versus explicit solvent in free energy calculations of enzyme catalysis: Methyl transfer catalyzed by catechol o-methyltransferase. *J. Chem. Phys.* **124,** 174503 (2006).

13. Heimdal, J., Rydberg, P. & Ryde, U. Protonation of the proximal histidine ligand in heme peroxidases. *J. Phys. Chem. B* **112,** 2501–2510 (2008).

14. Shaw, K. E., Woods, C. J. & Mulholland, A. J. Compatibility of quantum chemical methods and empirical (mm) water models in quantum mechanics/molecular mechanics liquid water simulations. *J. Phys. Chem. Lett.* **1,** 219–223 (2010).

15. Wesolowski, T. & Warshel, A. Ab initio free energy perturbation calculations of solvation free energy using the frozen density functional approach. *J. Phys. Chem.* **98,** 5183–5187 (1994).

16. Wood, R. H., Yezdimer, E. M., Sakane, S., Barriocanal, J. A. & Doren, D. J. Free energies of solvation with quantum mechanical interaction energies from classical mechanical simulations. *J. Chem. Phys.* **110,** 1329–1337 (1999).

17. Genheden, S., Cabedo Martinez, A., Criddle, M. & Essex, J. Extensive all-atom Monte Carlo sampling and QM/MM corrections in the sampl4 hydration free energy challenge. *J. Comput.-Aided Mol. Des.* **28,** 187–200 (2014).

18. Tuckerman, M. *Statistical mechanics: Theory and molecular simulation* (Oxford University Press, 2010).

19. Leach, A. R. *Molecular modelling: Principles and applications* (Pearson Education, Edinburgh Gate, Harlow, Essex, CM20 2JE, England, 2001).

20. Stone, A. *The theory of intermolecular forces* 2nd (Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP, UK, 2013).

21. Chipot, C. & Pohorille, A. *Free energy calculations* (Springer, Berlin, 2007).

22. Allen, M. P. & Tildesley, D. J. *Computer simulation of liquids* (Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP, UK, 1987).

23. Christ, C. D., Mark, A. E. & van Gunsteren, W. F. Basic ingredients of free energy calculations: A review. *J. Comput. Chem.* **31,** 1569–1582 (2010).

24. Szabo, A. & Ostlund, N. S. *Modern quantum chemistry* (Dover Publications Inc., 31 East 2nd Street, Mineola, NY 11501, 1996).

25. Rae, A. I. M. *Quantum mechanics* 3rd (Institute of Phys. Publishing, Dirac House, Temple Back, Bristol, BS1 6BE, 1992).

26. Hartree, D. *The wave mechanics of an atom with a non-coulomb central field. Part I. Theory and methods* in *Mathematical proceedings of the Cambridge philosophical society* **24** (Cambridge Univ Press, 1928), 89–110.

27. Yarkony, D. R. *Modern electronic structure theory part II* (World Scientific, P O Box 128, Farrer Road, Singapore 9128, 1995).

28. Argaman, N. & Makov, G. Density functional theory: An introduction. *Am. J. Phys.* **68,** 69 (2000).

29. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140,** 1133–1138 (1965).

30. Springborg, M. *Methods of electronic-structure calculations: From molecules to solids* (John Wiley and Sons Ltd, Baffins Lane, Chichester, West Sussex, PO19 1UD, 2000).

31. Srivastava, G. & Weaire, D. The theory of the cohesive energies of solids. *Adv. Phys.* **36,** 463–517 (1987).
32. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77,** 3865–3868 (1996).
33. Kaschner, R. & Hohl, D. Density functional theory and biomolecules:A study of glycine, alanine, and their oligopeptides. *J. Phys. Chem. A* **102,** 5111–5116 (1998).
34. Fox, S., Wallnoefer, H., Fox, T., Tautermann, C. & Skylaris, C. First principles-based calculations of free energy of binding: Application to ligand binding in a self-assembling superstructure. *J. Chem. Theory Comput.* **7,** 1102–1108 (2011).
35. Robinson, M. & Haynes, P. D. *Ab initio* NMR calculations: Classical force fields fitted to quantum forces. *J. Chem. Phys.* **133,** 084109 (2010).
36. Fonseca Guerra, C., van der Wijst, T., Poater, J., Swart, M. & Bickelhaupt, F. Adenine versus guanine quartets in aqueous solution: Dispersion-corrected DFT study on the differences in π-stacking and hydrogen-bonding behavior. *Theor. Chem. Acc.* **125),** 245–252 (2010).
37. Schwegler, E., Galli, G. & Gygi, F. Conformational dynamics of the dimethyl phosphate anion in solution. *Chem. Phys. Lett.* **342,** 434–440 (2001).
38. Fox, S. J. *et al.* Electrostatic embedding in large-scale first principles quantum mechanical calculations on biomolecules. *J. Chem. Phys.* **135,** 224107 (2011).
39. Payne, M. C., Teter, M. P., Allan, D. C., Arias, T. A. & Joannopoulos, J. D. Iterative minimization techniques for ab initio total-energy calculations: Mol. dynamics and conjugate gradients. *Rev. Mod. Phys.* **64,** 1045–1097 (1992).
40. Cornell, W. D. *et al.* A second generation force field for the simulation of proteins, nucleic acids, and organic molecules. *J. Am. Chem. Soc.* **117,** 5179–5197 (1995).
41. Lennard-Jones, J. E. Cohesion. *Proc. Phys. Soc.* **43,** 461 (1931).
42. Weiner, S. J. *et al.* A new force field for molecular mechanical simulation of nucleic acids and proteins. *J. Am. Chem. Soc.* **106,** 765–784 (1984).
43. Wang, J., Cieplak, P. & Kollman, P. A. How well does a restrained electrostatic potential (RESP) model perform in calculating conformational energies of organic and biological molecules? *J. Comput. Chem.* **21,** 1049–1074 (2000).
44. J.W., P. & D.A., C. Force fields for protein simulations. *Adv. Protein. Chem.* **66,** 27–85 (2003).
45. Weiner, S. J., Kollman, P. A., Nguyen, D. T. & Case, D. A. An all atom force field for simulations of proteins and nucleic acids. *J. Comput. Chem.* **7,** 230–252 (1986).
46. Bayly, C. I., Cieplak, P., Cornell, W. & Kollman, P. A. A well-behaved electrostatic potential based method using charge restraints for deriving atomic charges: The RESP model. *J. Phys. Chem.* **97,** 10269–10280 (1993).

47. Hornak, V. *et al.* Comparison of multiple amber force fields and development of improved protein backbone parameters. *Proteins: Struct., Func., Bioinf.* **65,** 712–725 (2006).
48. Wang, J., Wolf, R. M., Caldwell, J. W., Kollman, P. A. & Case, D. A. Development and testing of a general amber force field. *J. Comput. Chem.* **25,** 1157–1174 (2004).
49. Harmony, M. D. *et al.* Mol. structures of gasphase polyatomic molecules determined by spectroscopic methods. *J. Phys. Chem. Ref. Data* **8,** 619–722 (1979).
50. Allen, F. H. *et al.* Tables of bond lengths determined by X-ray and neutron diffraction. Part 1. Bond lengths in organic compounds. *J. Chem. Soc., Perk-Trans. 2,* S1–S19 (1987).
51. Guillot, B. A reappraisal of what we have learnt during three decades of computer simulations on water. *J. Mol. Liq.* **101.** Mol. Liquids. Water at the New Millenium, 219–260 (2002).
52. Jorgensen, W. L., Chandrasekhar, J., Madura, J. D., Impey, R. W. & Klein, M. L. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.* **79,** 926–935 (1983).
53. Jorgensen, W. L. & Madura, J. D. Temperature and size dependence for Monte Carlo simulations of TIP4P water. *Mol. Phys.* **56,** 1381–1392 (1985).
54. Drude, P. *The theory of optics* (Longmans, Green and Co., New York, 1902).
55. Lamoureux, G. & Roux, B. Modeling induced polarization with classical Drude oscillators: Theory and molecular dynamics simulation algorithm. *J. Chem. Phys.* **119,** 3025 (2003).
56. Ewald, P. P. Die berechnung optischer und elektrostatischer gitterpotentiale. *Ann. Phys.* **369,** 253–287 (1921).
57. Darden, T., York, D. & Pedersen, L. Particle mesh Ewald: An nlog(n) method for Ewald sums in large systems. *J. Chem. Phys.* **98,** 10089–10092 (1993).
58. Grest, G. S. & Kremer, K. Molecular dynamics simulation for polymers in the presence of a heat bath. *Phys. Rev. A* **33,** 3628–3631 (1986).
59. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21,** 1087–1092 (1953).
60. Hastings, W. K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57,** 97–109 (1970).
61. Mezei, M. The finite difference thermodynamic integration, tested on calculating the hydration free energy difference between acetone and dimethylamine in water. *J. Chem. Phys.* **86,** 7084–7088 (1987).
62. Woods, C., King, M. & Essex, J. in *New algorithms for macromolecular simulation* (eds Leimkuhler, B. *et al.*) 251–259 (Springer Berlin Heidelberg, 2006).
63. Woods, C. J., Essex, J. W. & King, M. A. The development of replica-exchange-based free-energy methods. *J. Phys. Chem. B* **107,** 13703–13710 (2003).

64. Van Eerden, J., Briels, W., Harkema, S. & Feil, D. Potential of mean force by thermodynamic integration: Mol.-dynamics simulation of decomplexation. *Chem. Phys. Lett.* **164,** 370–376 (1989).
65. Torrie, G. & Valleau, J. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *J. Comput. Phys.* **23,** 187–199 (1977).
66. Mülders, T., Krüger, P., Swegat, W. & Schlitter, J. Free energy as the potential of mean constraint force. *J. Chem. Phys.* **104,** 4869–4870 (1996).
67. Andersen, H. C. RATTLE: A “velocity” version of the SHAKE algorithm for molecular dynamics calculations. *J. Comput. Phys.* **52,** 24–34 (1983).
68. Ryckaert, J.-P., Ciccotti, G. & Berendsen, H. J. Numerical integration of the Cartesian equations of motion of a system with constraints: Molecular dynamics of n-alkanes. *J. Comput. Phys.* **23,** 327–341 (1977).
69. Iftimie, R., Salahub, D., Wei, D. & Schofield, J. Using a classical potential as an efficient importance function for sampling from an ab initio potential. *J. Chem. Phys.* **113,** 4852–4862 (2000).
70. Lin, H. & Truhlar, D. QM/MM: What have we learned, where are we, and where do we go from here? *Theor. Chem. Acc.* **117,** 185–199 (2007).
71. Senn, H. M. & Thiel, W. QM/MM methods for biomolecular systems. *Angew. Chem., Int. Ed.* **48,** 1198–1229 (2009).
72. Field, M. J., Bash, P. A. & Karplus, M. A combined quantum mechanical and molecular mechanical potential for molecular dynamics simulations. *J. Comput. Chem.* **11,** 700–733 (1990).
73. Bakowies, D. & Thiel, W. Hybrid models for combined quantum mechanical and molecular mechanical approaches. *J. Phys. Chem.* **100,** 10580–10594 (1996).
74. Christ, C. D. & van Gunsteren, W. F. Enveloping distribution sampling: A method to calculate free energy differences from a single simulation. *J. Chem. Phys.* 184110 (2007).
75. Iftimie, R., Salahub, D. & Schofield, J. An efficient Monte Carlo method for calculating ab initio transition state theory reaction rates in solution. *J. Chem. Phys.* **119,** 11285–11297 (2003).
76. Michel, J., Taylor, R. D. & Essex, J. W. Efficient generalized Born models for Monte Carlo simulations. *J. Chem. Theory Comput.* **2,** 732–739 (2006).
77. Hetenyi, B., Bernacki, K. & Berne, B. Multiple “time step” Monte Carlo. *J. Chem. Phys.* **117,** 8203–8207 (2002).
78. Coe, J. D., Sewell, T. D. & Shaw, M. S. Optimal sampling efficiency in Monte Carlo simulation with an approximate potential. *J. Chem. Phys.* **130,** – (2009).
79. König, G., Pickard, I. FrankC., Mei, Y. & Brooks, B. Predicting hydration free energies with a hybrid QM/MM approach: An evaluation of implicit and explicit solvation models in sampl4. *J. Comput.-Aided Mol. Des.* **28,** 245–257 (2014).

80. Case, D. A. *et al. Amber 12* 2012.
81. Clark, S. J. *et al.* First principles methods using castep. *Z Kristallogr.* **220,** 567–570 (2005).
82. Perdew, J. P. & Zunger, A. Self-interaction correction to density-functional approximations for many-electron systems. *Phys. Rev. B* **23,** 5048–5079 (1981).
83. Marx, D. & Hutter, J. in *Modern methods and algorithms of quantum chemistry* (ed Gortendorst, J.) 301–449 (John von Neumann Institute for Computing, 2000).
84. Segall, M. D. *et al.* First-principles simulation: Ideas, illustrations and the castep code. *J. Phys.: Condens. Matter* **14,** 2717–2744 (2002).
85. Case, D. A. *et al.* The amber biomolecular simulation programs. *J. Comput. Chem.* **26 (16),** 1668–1688 (2005).
86. Jakalian, A., Bush, B. L., Jack, D. B. & Bayly, C. I. Fast, efficient generation of high-quality atomic charges. AM1-bcc model: I. Method. *J. Comput. Chem.* **21,** 132–146 (2000).
87. Jakalian, A., Jack, D. B. & Bayly, C. I. Fast, efficient generation of high-quality atomic charges. AM1-bcc model:II. Parameterization and validation. *J. Comput. Chem.* **23,** 1623–1641 (2002).
88. Pfrommer, B. G., Côté, M., Louie, S. G. & Cohen, M. L. Relaxation of crystals with the quasi-Newton method. *J. Comput. Phys.* **131,** 233–240 (1997).
89. Elber, R. Calculation of the potential of mean force using molecular dynamics with linear constraints: An application to a conformational transition in a solvated dipeptide. *J. Chem. Phys.* **93 (6),** 4312–4321 (1990).
90. Lennox, C & Chadwick, M. *Mathematics for engineers and applied scientists* 2nd (Heinemann Educational Books Ltd, 22 Bedford Square, London, WC1B 3HH, 1977).
91. Niklasson, A. M. N., Tymczak, C. J. & Challacombe, M. Time-reversible Born-Oppenheimer molecular dynamics. *Phys. Rev. Lett.* **97,** 123001 (2006).
92. Sadhukhan, S., Munoz, D., Adamo, C. & Scuseria, G. E. Predicting proton transfer barriers with density functional methods. *Chem. Phys. Lett.* **306,** 83–87 (1999).
93. Pohorille, A., Jarzynski, C. & Chipot, C. Good practices in free-energy calculations. *J. Phys. Chem. B* **114,** 10235–10253 (2010).
94. Cave-Ayland, C., Skylaris, C. K. & Essex, J. W. Direct validation of the single step classical to quantum free energy perturbation. *J. Phys. Chem. B* **ASAP.** DOI: 10.1021/jp506459v. eprint: `http://dx.doi.org/10.1021/jp506459v`.
95. Boresch, S. The role of bonded energy terms in free energy simulations - insights from analytical results. *Mol. Simul.* **28,** 13–37 (2002).
96. Evans, L. C. & Gariepy, R. F. *Measure theory and fine properties of functions. Studies in Advanced Mathematics.* (CRC Press, Boca Raton, FL, 1992).
97. Boresch, S. & Karplus, M. The role of bonded terms in free energy simulations: 1. Theoretical analysis. *J. Phys. Chem. A* **103,** 103–118 (1999).

98. Boresch, S. & Karplus, M. The role of bonded terms in free energy simulations. 2. Calculation of their influence on free energy differences of solvation. *J. Phys. Chem. A* **103,** 119–136 (1999).
99. Kurumbail, R. G. *et al.* Structural basis for selective inhibition of cyclooxygenase-2 by anti-inflammatory agents. *Nat.* **384,** 644–648 (1996).
100. Michel, J., Verdonk, M. L. & Essex, J. W. Protein-ligand binding affinity predictions by implicit solvent simulations: A tool for lead optimization? *J. Med. Chem.* **49,** 7427–39 (2006).
101. Nocedal, J. & Wright, S. J. *Numerical optimization* 2nd (Springer, New York, 2006).
102. Thole, B. Molecular polarizabilities calculated with a modified dipole interaction. *Chem. Phys.* **59,** 341–350 (1981).
103. Alberts, B. *et al. Molecular Biology of the Cell* 3rd (Garland Publishing Inc., 717 Fifth Avenue, New York, NY 10022, 1994).
104. Beckman, E. M. *et al.* CD1c restricts responses of mycobacteria-specific T-cells. Evidence for antigen presentation by a second member of the human CD1 family. *J. Immunol.* **157,** 2795–803 (1996).
105. Bach, F., Bach, M. & Sondel, P. Differential function of major histocompatibility complex antigens in T-lymphocyte activation. *Nat.* **259,** 273281 (1976).
106. Rhijn, I. V. *et al.* CD1c bypasses lysosomes to present a lipopeptide antigen with 12 amino acids. *J. Exp. Med.* **206,** 1409–1422 (2009).
107. Haig, N. A., Guan, Z., Li, D., Mcmichael, A. & Raetz, C. R. H. Identification of self-lipids presented by CD1c and CD1d. *J. Biol. Chem.* **286,** 37692–37701 (2011).
108. Scharf, L. *et al.* Structure of CD1c in complex with a mycobacterial lipid reveals an open groove ideally suited for diverse antigen presentation. *Immunity* **33,** 853–862 (2010).
109. Wlodawer, A., Minor, W., Dauter, Z. & Jaskolski, M. Protein crystallography for non-crystallographers, or how to get the best (but not more) from published macromolecular structures. *FEBS J.* **275,** 1–21 (2008).
110. Garzón, D., Anselmi, C., Bond, P. J. & Faraldo-gómez, J. D. Dynamics of the antigen-binding grooves in CD1 proteins. *J. Biol. Chem.* **288,** 19528–19536 (2013).
111. Garzón, D., Bond, P. J. & Faraldo-gómez, J. D. Predicted structural basis for CD1c presentation of mycobacterial branched polyketides and long lipopeptide antigens. *Mol. Immunol.* **47,** 253–260 (2009).
112. Israelachvili, J. The different faces of poly(ethyleneglycol). *Proc. National Academy Sciences* **94,** 8378–8379 (1997).
113. *Chem. Computing Group Inc., 1010 Sherbooke St. West, Suite #910, Montreal, QC, Canada, H3A 2R7. Mol. Operating Environment (MOE), 2013.10* 2013.

114. Berendsen, H. J. C., Postma, J. P. M., van Gunsteren, W. F., DiNola, A & Haak, J. R. Molecular dynamics with coupling to an external bath. *J. Chem. Phys.* **81** (1984).
115. Götz, A. W. *et al.* Routine microsecond molecular dynamics simulations with amber on GPUs. 1. Generalized Born. *J. Chem. Theory Comput.* **8,** 1542–1555 (2012).
116. Salomon-Ferrer, R., Götz, A. W., Poole, D., Le Grand, S. & Walker, R. C. Routine microsecond molecular dynamics simulations with amber on GPUs. 2. Explicit solvent particle mesh Ewald. *J. Chem. Theory Comput.* **9,** 3878–3888 (2013).
117. Craig, I. R., Pfleger, C., Gohlke, H., Essex, J. W. & Spiegel, K. Pocket-space maps to identify novel binding-site conformations in proteins. *J. Chem. Inf. Model.* **51,** 2666–2679 (2011).
118. Moody, D. B., Mu, W., Grant, E., Rosat, J.-p. & Brenner, M. B. CD1c-mediated T-cell recognition of isoprenoid glycolipids in mycobacterium tuberculosis infection. *Nat.* **404,** 884–888 (2000).
119. Swope, W. C., Andersen, H. C., Berens, P. H. & Wilson, K. R. A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters. *J. Chem. Phys.* **76,** 637–649 (1982).
120. Vriend, G. What If: A molecular modeling and drug design program. *J. Mol. Graphics* **8,** 52–56 (1990).
121. Word, J., Lovell, S. C., Richardson, J. S. & Richardson, D. C. Asparagine and glutamine: Using hydrogen atom contacts in the choice of side-chain amide orientation. *J. Mol. Biol.* **285,** 1735–1747 (1999).

Appendices

A. Numerical Quadrature

Quadrature techniques provide approximate solutions to problems of the form:

$$\int_a^b f(x)dx \tag{A.1}$$

Where f is some function integrated over the range a to b. The most basic algorithm is given by the trapezium rule:

$$\int_a^b f(x)dx \approx \frac{(b-a)}{2N}\sum_{i=1}^{N}[f(x_i)+f(x_{i+1})] \tag{A.2}$$

Where the range a to b is divided over N evenly sized windows. The values of x_i are a collection of N+1 points at which f is evaluated and are referred to as nodes. The trapezium rule can be thought of as fitting a linear polynomial to the function at each node and giving the exact area of the fitted function.

Rather than assuming a fixed interval between nodes, a more general collection of quadrature techniques can be derived by generalising the form of the problem.

$$\int_{-1}^{1} f(x)dx \approx \sum_{i}^{N} w_i f(x_i) \tag{A.3}$$

Where w_i is a set of weights paired with the nodes x_i. Gaussian quadrature rules are defined by the number nodes used in their evaluation, known as the order of the rule. Each order of rule has an associated set of nodes and weights. Different forms of Gaussian quadrature are associated with fitting different forms of polynomial, for instance the integration routines from the Scipy library used extensively throughout this work make use of adaptive Chebyshev-Gauss quadrature.

A.1. Adaptive Quadrature

The results given by numerical quadrature rules are inherently approximate. There is a direct trade off between the number of nodes used in a calculation and the accuracy of the answer produced. Increasing the number of nodes within an integral also increases the required number of function evaluations and hence the computational cost of evaluating the integral. Where function evaluations are expensive, it is desirable therefore to be able to solve an integral to a certain degree of accuracy with the minimum required number of nodes. This is carried out through the use of adaptive quadrature. The procedure is given by:

1. Evaluate the integral with some number of nodes, N
2. Increase the value of N
3. Evaluate the integral with new N
4. Compare answers for different values of N
5. If difference between answers is less then the desired accuracy stop
6. Else repeat from 2.

The difference between evaluations of the integral with different values of N is known as the relative error, the desired value of which is specified by the user. Efficient procedures can be defined that make efficient use of function evaluations with previous values of N.

B. Numerical Methods for the Solution of Differential Equations

This section addresses numerical methods for the solution of differential equations of the form:

$$\frac{dx(t)}{dt} = f(x(t)) \tag{B.1}$$

Where a solution takes the form of $x(t)$. The simplest approach to is the application of the Euler method[90]. An intuitive and straightforward process, it is given by:

$$x(t+h) = x(t) + f(x(t))h \tag{B.2}$$

Where h is the freely chosen step size of the algorithm. Here t is discretised and the value of $f(x(t))$ is treated as constant over the interval h. In practice the Euler method requires a very small h to accurately approximate a solution to $x(t)$. For the majority of application far more sophisticated approaches than the Euler method are used (e.g. high order Runge-Kutta methods), however for some problems it remains suitable.

B.1. Integrators for MD

Although in principle its form is the same as equation (B.1), the demanding problem posed by MD, i.e. the solution of equation (2.60), has led to the development of a number of specialised algorithms with desirable numerical properties. A major problem with the Euler method (and its higher order associates) is that error in the solution aggregates globally. This means that the more steps of integration you carry out the further from the desired solution you end up. In the case of MD, this manifests itself in terms of drift in the total energy of the system. This is problematic in ensuring that one's statistics correspond to the correct ensemble.

This is resolved through the use of a class of methods that do not accumulate error globally, known as symplectic integrators. Symplectic methods are specialised in that they depend on a form of symmetry, inherent in the MD problem, however with a sufficiently small step size will conserve the energy of a system within. In practice, because of the discretisation of time inherent to such methods they do not give an exact solution for $x(t)$, however they have proved to be sufficiently stable to allow even protracted MD runs. A drawback of such methods is that they are limited to being 2nd order accurate in the local error.

Here are presented two of the most frequently used algorithms in a variety of popular MD packages. The first of these is the Leapfrog algorithm[50]:

$$\mathbf{x}_i(t+\Delta t) = \mathbf{x}_i(t) + \mathbf{v}_i(t+\frac{1}{2}\Delta t) \tag{B.3}$$

$$\mathbf{v}_i(t+\frac{1}{2}\Delta t) = \mathbf{v}_i(t-\frac{1}{2}\Delta t) + \mathbf{a}_i(t) \tag{B.4}$$

Where Δt is the timestep of the algorithm. This has an unusual form as the velocities are offset from the positions by a half timestep, leading to its unusual name as they hop past each other. The Leapfrog algorithm is symplectic and 2nd order accurate, however the noted offset introduces drawbacks. If the velocities at an integer timestep are required they must be interpolated from the half timesteps. This requires additional computational effort and the storage of two sets of velocities in memory. A popular alternative to the Leapfrog algorithm as it does not suffer from these notable drawbacks is the Velocity Verlet algorithm[119]:

$$\mathbf{x}(t+\Delta t) = \mathbf{x}(t) + \mathbf{v}(t)\Delta t + \frac{1}{2}\mathbf{a}(t)\Delta t^2 \tag{B.5}$$

$$\mathbf{v}(t+\Delta t) = \mathbf{v}(t) + \frac{1}{2}[\mathbf{a}(t) + \mathbf{a}(t+\Delta t)]\Delta t \tag{B.6}$$

In fact this integrator can be derived through straightforward algebraic manipulation of equations (B.3) and (B.4). In this form it retains the desirable properties of Leapfrog integration but can be efficiently implemented to require the storage of only a single set of velocities.

C. Dual Topology and Interaction Energies

We here demonstrate that the use of interaction energies that arises through the use of dual topology approaches (section 2.4) is distinct from their usage in the case of single topology. Interaction energies emerge naturally in dual topology approaches and we here demonstrate rigorously their correctness. In parallel with the development of section 4.2 we consider free energy of hydration calculations of the form given by figure 2.6. We start by defining the two states involved in the calculation:

$$\begin{aligned}
U_3^{dual}(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q},\mathbf{r}^{l2}) &= U_{\mathrm{P}}^l(\mathbf{r}^l) + U^s(\mathbf{r}^s) + U_3^{inter}(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q}) + U_{\mathrm{Q}}^l(\mathbf{r}^{l2}) && \text{(C.1)}\\
&= U_3^{intra}(\mathbf{r}^l,\mathbf{r}^s) + U_3^{inter}(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q}) + U_{\mathrm{Q}}^l(\mathbf{r}^{l2}) && \text{(C.2)}\\
U_4^{dual}(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q},\mathbf{r}^{l2}) &= U_{\mathrm{Q}}^l(\mathbf{r}^l) + U^s(\mathbf{r}^s) + U_4^{inter}(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q}) + U_{\mathrm{P}}^l(\mathbf{r}^{l2}) && \text{(C.3)}\\
&= U_4^{intra}(\mathbf{r}^l,\mathbf{r}^s) + U_4^{inter}(\mathbf{r}^l,\mathbf{r}^s,\mathbf{q}) + U_{\mathrm{P}}^l(\mathbf{r}^{l2}) && \text{(C.4)}
\end{aligned}$$

And consider their application within the Zwanzig equation:

$$\begin{aligned}
\Delta\Delta A_{hyd}^{dual} &= \Delta A_{3\to 4}^{dual} && \text{(C.5)}\\
&= -\frac{1}{\beta}\ln\left\langle \exp\left[-\beta(U_4^{inter} - U_3^{inter})\right]\right\rangle_3 && \text{(C.6)}\\
&= -\frac{1}{\beta}\ln\left\langle \exp\left[-\beta((U_4^{inter} + U^s + U_{\mathrm{Q}}^l + U_{\mathrm{P}}^l) - (U_3^{inter} + U^s + U_{\mathrm{P}}^l + U_{\mathrm{Q}}^l))\right]\right\rangle_3 && \text{(C.7)}\\
&= -\frac{1}{\beta}\ln\left\langle \exp\left[-\beta(U_4^{dual} - U_3^{dual})\right]\right\rangle_3 && \text{(C.8)}\\
&= -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta U_3^{dual}\right]\exp\left[-\beta(U_4^{dual} - U_3^{dual})\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}}{\int \exp\left[-\beta U_3^{dual}\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}} && \text{(C.9)}\\
&= -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta U_3^{dual}\right]\exp\left[-\beta(U_4^{dual} - U_3^{dual})\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}}{\int \exp\left[-\beta U_3^{dual}\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}} && \text{(C.10)}\\
&= -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta U_4^{dual}\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}}{\int \exp\left[-\beta U_3^{dual}\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}} && \text{(C.11)}\\
&= -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_4^{intra} + U_4^{inter})\right]\exp\left[-\beta U_{\mathrm{P}}^l\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}}{\int \exp\left[-\beta(U_3^{intra} + U_3^{inter})\right]\exp\left[-\beta U_{\mathrm{Q}}^l\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} d\mathbf{r}^{l2}} && \text{(C.12)}\\
&= -\frac{1}{\beta}\ln\frac{\int \exp\left[-\beta(U_4^{intra} + U_4^{inter})\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} \int \exp\left[-\beta U_{\mathrm{P}}^l\right] d\mathbf{r}^{l2}}{\int \exp\left[-\beta(U_3^{intra} + U_3^{inter})\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q} \int \exp\left[-\beta U_{\mathrm{Q}}^l\right] d\mathbf{r}^{l2}} && \text{(C.13)}
\end{aligned}$$

$$= -\frac{1}{\beta} \ln \frac{\dfrac{\int \exp\left[-\beta(U_4^{intra} + U_4^{inter})\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}{\int \exp\left[-\beta(U_3^{intra} + U_3^{inter})\right] d\mathbf{r}^l d\mathbf{r}^s d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U_{\mathrm{Q}}^l\right] d\mathbf{r}^{l2}}{\int \exp\left[-\beta U_{\mathrm{P}}^l\right] d\mathbf{r}^{l2}}} \tag{C.14}$$

$$= \Delta\Delta A_{hyd} \tag{C.15}$$

For convenience we have dropped explicit notation of many of the degrees of freedom. However these results surfice to demonstrate $\Delta\Delta A_{hyd}^{dual} = \Delta\Delta A_{hyd}$, and thus represents a situation in which interaction energies arise in a rigorous context that is distinct from the single topology approach.

D. Invariance of Forward-Backward Free Energy Perturbation Differences with Respect to Potential Shifts by a Constant

It can be shown that adjusting the energy values within the ensemble average by a constant value is equivalent to adjusting the overall free energy difference by the same amount i.e.:

$$\Delta A_{0\to 1} + \alpha = \frac{1}{\beta} \ln \left\langle \exp\left[-\beta(U_1 - U_0 + \alpha)\right]\right\rangle_0 \tag{D.1}$$

As:

$$\begin{aligned}
\Delta A_{0\to 1} + \alpha &= \frac{1}{\beta} \ln \left\langle \exp\left[-\beta(U_1 - U_0)\right]\right\rangle_0 + \alpha && \text{(D.2)}\\
&= -\frac{1}{\beta} \ln \left\langle \exp\left[-\beta(U_1 - U_0)\right]\right\rangle_0 + \frac{1}{\beta}\beta\alpha && \text{(D.3)}\\
&= -\frac{1}{\beta} \ln \left\langle \exp\left[-\beta(U_1 - U_0)\right]\right\rangle_0 - \frac{1}{\beta}\ln \exp[-\beta\alpha] && \text{(D.4)}\\
&= -\frac{1}{\beta} \ln \left\langle \exp\left[-\beta(U_1 - U_0)\right] \exp[-\beta\alpha]\right\rangle_0 && \text{(D.5)}\\
\Delta A_{0\to 1} + \alpha &= -\frac{1}{\beta} \ln \left\langle \exp\left[-\beta(U_1 - U_0 + \alpha)\right]\right\rangle_0 && \text{(D.6)}
\end{aligned}$$

The discrepancy between the perturbation in either direction is given by γ:

$$\begin{aligned}
&\gamma = \Delta A_{MM\to QM} + \Delta A_{QM\to MM} && \text{(D.7)}\\
&\gamma = -\frac{1}{\beta} \ln \left\langle \exp[-\beta(U_{QM} - U_{MM})]\right\rangle_{MM} - \frac{1}{\beta} \ln \left\langle \exp[-\beta(U_{MM} - U_{QM})]\right\rangle_{QM} && \text{(D.8)}
\end{aligned}$$

The use of different values of α is equivalent to shifting the potential energies by different constants and hence adjusting the relative heights of the potential energy surfaces. As the relative heights of the surfaces can be represented with only a single scalar we shall, without loss of generality, consider the effect of shifting only the MM potential surface with the value α. This gives an adjusted potential energy surface

$U_{ref} = U_{MM} - \alpha$.

$$\gamma = -\frac{1}{\beta}\ln\left\langle \exp[-\beta(U_{QM} - U_{ref})]\right\rangle_{MM} - \frac{1}{\beta}\ln\left\langle \exp[-\beta(U_{ref} - U_{QM})]\right\rangle_{QM} \qquad \text{(D.9)}$$

Using (D.6) it is trivial to show that the value of γ is invariant with the value of α:

$$\gamma = -\frac{1}{\beta}\ln\left\langle \exp[-\beta(U_{QM} - (U_{MM} - \alpha))]\right\rangle_{MM} - \frac{1}{\beta}\ln\left\langle \exp[-\beta((U_{MM} - \alpha) - U_{QM})]\right\rangle_{QM} \qquad \text{(D.10)}$$

$$\gamma = -\frac{1}{\beta}\ln\left\langle \exp[-\beta(U_{QM} - U_{MM})]\right\rangle_{MM} + \alpha - \frac{1}{\beta}\ln\left\langle \exp[-\beta(U_{MM} - U_{QM})]\right\rangle_{QM} - \alpha \qquad \text{(D.11)}$$

$$\gamma = \Delta A_{\mathrm{MM}\rightarrow\mathrm{QM}} + \Delta A_{\mathrm{QM}\rightarrow\mathrm{MM}} \qquad \text{(D.12)}$$

E. Toy System Raw Data

We here present Monte Carlo results in corroboration of the model system calculations presented in section 4.3. All Monte Carlo results show the average and standard error from 5 independent repeats from runs of 10^6 MC moves. Parameters and equivalent figures can be found in section 4.3.

ϵ^B	$\Delta\Delta A_{hyd}$	$\Delta\Delta A_{hyd}^{inter}$	$\Delta\Delta A_{hyd}$ MC	$\Delta\Delta A_{hyd}^{inter}$ MC
0.000	-0.630	-0.630	-0.637 ± 0.001	-0.637 ± 0.001
0.100	-0.639	-0.639	-0.647 ± 0.001	-0.647 ± 0.001
0.200	-0.648	-0.648	-0.649 ± 0.001	-0.650 ± 0.001
0.500	-0.672	-0.675	-0.676 ± 0.001	-0.678 ± 0.001
0.700	-0.687	-0.691	-0.693 ± 0.001	-0.698 ± 0.001
1.000	-0.707	-0.714	-0.712 ± 0.001	-0.719 ± 0.001
1.500	-0.732	-0.745	-0.734 ± 0.001	-0.747 ± 0.001
2.000	-0.747	-0.770	-0.749 ± 0.001	-0.771 ± 0.000
2.500	-0.753	-0.787	-0.756 ± 0.001	-0.789 ± 0.001
3.000	-0.750	-0.799	-0.751 ± 0.000	-0.800 ± 0.000
3.500	-0.739	-0.805	-0.740 ± 0.000	-0.805 ± 0.000
4.000	-0.721	-0.807	-0.722 ± 0.001	-0.807 ± 0.000
5.000	-0.669	-0.801	-0.670 ± 0.000	-0.801 ± 0.000

(a) $\sigma^B = 1.7$ Å

ϵ^B	$\Delta\Delta A_{hyd}$	$\Delta\Delta A_{hyd}^{inter}$	$\Delta\Delta A_{hyd}$ MC	$\Delta\Delta A_{hyd}^{inter}$ MC
0.000	-0.630	-0.630	-0.634 ± 0.001	-0.634 ± 0.001
0.100	-0.644	-0.642	-0.645 ± 0.001	-0.644 ± 0.001
0.200	-0.657	-0.653	-0.663 ± 0.001	-0.659 ± 0.001
0.500	-0.695	-0.686	-0.696 ± 0.001	-0.687 ± 0.001
0.700	-0.717	-0.705	-0.722 ± 0.001	-0.709 ± 0.001
1.000	-0.746	-0.730	-0.750 ± 0.000	-0.734 ± 0.000
1.500	-0.786	-0.764	-0.787 ± 0.000	-0.764 ± 0.000
2.000	-0.817	-0.790	-0.817 ± 0.000	-0.790 ± 0.000
2.500	-0.841	-0.809	-0.840 ± 0.000	-0.808 ± 0.000
3.000	-0.860	-0.824	-0.859 ± 0.000	-0.824 ± 0.000
3.500	-0.874	-0.835	-0.876 ± 0.000	-0.836 ± 0.000
4.000	-0.886	-0.844	-0.887 ± 0.000	-0.843 ± 0.000
5.000	-0.904	-0.856	-0.903 ± 0.000	-0.856 ± 0.000

(b) $\sigma^B = 2.0$ Å

ϵ^B	$\Delta\Delta A_{hyd}$	$\Delta\Delta A_{hyd}^{inter}$	$\Delta\Delta A_{hyd}$ MC	$\Delta\Delta A_{hyd}^{inter}$ MC
0.000	-0.630	-0.630	-0.636 ± 0.001	-0.636 ± 0.001
0.100	-0.606	-0.598	-0.607 ± 0.001	-0.600 ± 0.001
0.200	-0.581	-0.574	-0.583 ± 0.001	-0.576 ± 0.001
0.500	-0.521	-0.525	-0.523 ± 0.001	-0.528 ± 0.001
0.700	-0.495	-0.506	-0.501 ± 0.001	-0.512 ± 0.001
1.000	-0.468	-0.487	-0.471 ± 0.001	-0.490 ± 0.000
1.500	-0.442	-0.470	-0.443 ± 0.000	-0.472 ± 0.000
2.000	-0.428	-0.462	-0.431 ± 0.001	-0.465 ± 0.001
2.500	-0.420	-0.457	-0.419 ± 0.001	-0.456 ± 0.001
3.000	-0.415	-0.455	-0.414 ± 0.001	-0.455 ± 0.000
3.500	-0.411	-0.453	-0.412 ± 0.001	-0.453 ± 0.000
4.000	-0.409	-0.452	-0.409 ± 0.000	-0.452 ± 0.000
5.000	-0.405	-0.450	-0.405 ± 0.000	-0.450 ± 0.000

(c) $\sigma^B = 2.3$ Å

Table E.1.: Monte Carlo results corresponding to figure 4.3.

k_m	$\Delta\Delta A_{hyd}$	$\Delta\Delta A_{hyd}^{inter}$	$\Delta\Delta A_{hyd}$ MC	$\Delta\Delta A_{hyd}^{inter}$ MC
0.000	-0.750	-0.799	-0.784 ± 0.001	-0.795 ± 0.000
10.000	-0.760	-0.798	-0.761 ± 0.001	-0.799 ± 0.001
20.000	-0.767	-0.797	-0.769 ± 0.001	-0.798 ± 0.001
50.000	-0.778	-0.795	-0.779 ± 0.000	-0.796 ± 0.000
100.000	-0.785	-0.794	-0.786 ± 0.001	-0.794 ± 0.001
150.000	-0.788	-0.793	-0.789 ± 0.001	-0.794 ± 0.001
200.000	-0.790	-0.793	-0.789 ± 0.001	-0.792 ± 0.001
250.000	-0.790	-0.793	-0.791 ± 0.001	-0.793 ± 0.001
300.000	-0.791	-0.793	-0.793 ± 0.001	-0.794 ± 0.001
400.000	-0.791	-0.793	-0.793 ± 0.001	-0.794 ± 0.001
500.000	-0.792	-0.792	-0.792 ± 0.001	-0.793 ± 0.001
600.000	-0.792	-0.792	-0.791 ± 0.001	-0.791 ± 0.001
700.000	-0.792	-0.792	-0.791 ± 0.001	-0.792 ± 0.001
800.000	-0.792	-0.792	-0.793 ± 0.000	-0.794 ± 0.000
900.000	-0.792	-0.792	-0.793 ± 0.001	-0.793 ± 0.001

Table E.2.: Monte Carlo results corresponding to figure 4.4.

b_Q^0	$\Delta\Delta A_{hyd}$	$\Delta\Delta A_{hyd}^{inter}$	$\Delta\Delta A_{hyd}$ MC	$\Delta\Delta A_{hyd}^{inter}$ MC
1.000	-0.630	-0.630	-0.635 ± 0.001	-0.635 ± 0.001
1.010	-0.629	-0.630	-0.631 ± 0.001	-0.632 ± 0.001
1.020	-0.628	-0.630	-0.629 ± 0.002	-0.631 ± 0.002
1.030	-0.627	-0.630	-0.631 ± 0.001	-0.633 ± 0.001
1.040	-0.626	-0.630	-0.631 ± 0.001	-0.637 ± 0.002
1.050	-0.624	-0.630	-0.630 ± 0.002	-0.636 ± 0.001
1.060	-0.623	-0.630	-0.632 ± 0.004	-0.637 ± 0.002
1.070	-0.622	-0.630	-0.627 ± 0.003	-0.634 ± 0.002
1.080	-0.621	-0.630	-0.622 ± 0.003	-0.631 ± 0.003
1.090	-0.620	-0.630	-0.624 ± 0.006	-0.631 ± 0.002
1.100	-0.619	-0.630	-0.623 ± 0.004	-0.636 ± 0.002

Table E.3.: Monte Carlo results corresponding to figure 4.5.

F. Invariance of Interaction Energy Error With Respect to State 3

As dicussed fully in section 4.6 the difference between the free energy of hydration calculated by total and interaction energies is given by:

$$\Delta\Delta A_{hyd} - \Delta\Delta A_{hyd}^{inter} = -\frac{1}{\beta} \ln \frac{\dfrac{\int \exp\left[-\beta(U_{\mathrm{Q}}^{l} + U^{s} + U_{4}^{inter})\right] d\mathbf{r}d\mathbf{q}}{\int \exp\left[-\beta(U_{\mathrm{P}}^{l} + U^{s} + U_{4}^{inter})\right] d\mathbf{r}d\mathbf{q}}}{\dfrac{\int \exp\left[-\beta U_{\mathrm{Q}}^{l}\right] d\mathbf{r}^{l}}{\int \exp\left[-\beta U_{\mathrm{P}}^{l}\right] d\mathbf{r}^{l}}} \tag{F.1}$$

Interestingly, this expression does not contain the partition function of state 3 of the free energy of hydration cycle (figure 4.1). Instead it contains the partition functions of states 1, 2 and 4, as well as an apparent hybrid state composed of the intramolecular terms of state 3 and the intermolecular terms of state 4. This expression can be shown to approach zero in the limit of integral separability, however separability of this expression is independent of the properties of state 3, i.e. the state sampled through application of the Zwanzig equation.

This surprising property is demonstrated here through calculations with the model system. Two sets of calculations were completed with the parameter sets given by tables F.1 and F.2. Figure F.1 demonstrates that $\Delta\Delta A_{hyd} = \Delta\Delta A_{hyd}^{inter}$ so long as the state 4 remains completely separable. Meanwhile figure F.2 shows that impairing the separability of state 4, leads to an increasing degree of interaction energy error.

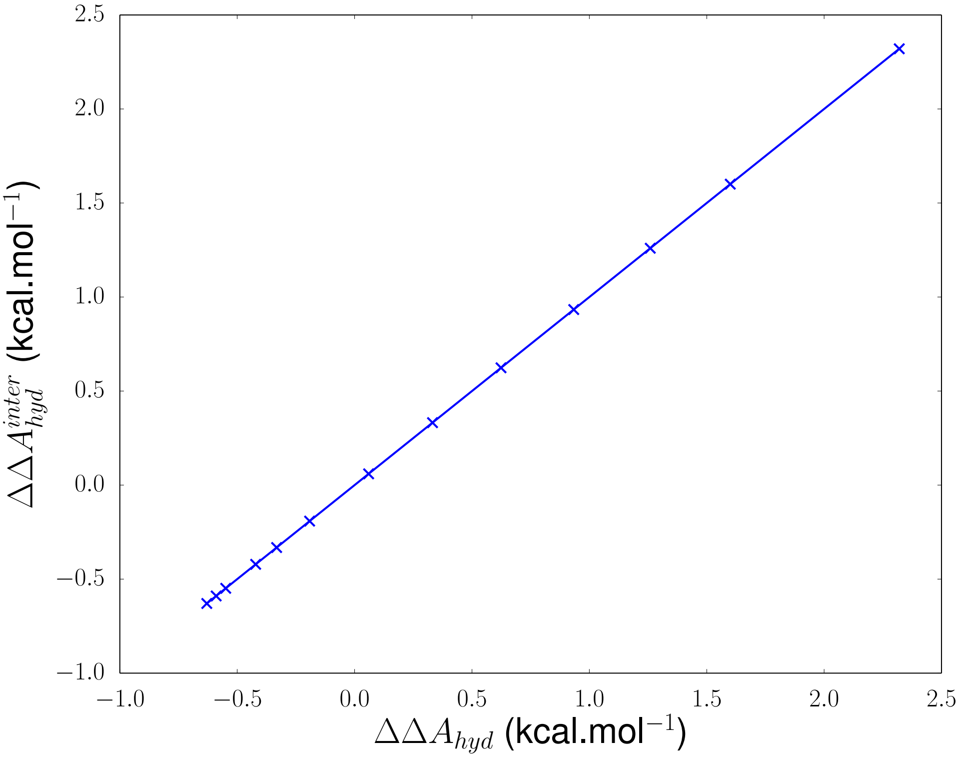

Figure F.1.: Model system calculations with variable separability of state 3 and perfect separability of state 4 i.e. ϵ_3^B is varied whilst $\epsilon_4^B = 0$.

Parameter	Value
k_{P}	50 kcal.mol^{-1}.Å^{-2}
k_{Q}	100 kcal.mol^{-1}.Å^{-2}
b^0_{P},b^0_{Q}	1.0 Å
ϵ_3^A	1.0 kcal.mol^{-1}
ϵ_4^A	2.0 kcal.mol^{-1}
ϵ_3^B	var kcal.mol^{-1}
ϵ_4^B	0.0 kcal.mol^{-1}
σ_3^A,σ_4^A	1.0 Å
σ_3^B,σ_4^B	var Å
L	3.0

Table F.1.: Parameter sets used in the calculations with the model system for figure F.1.

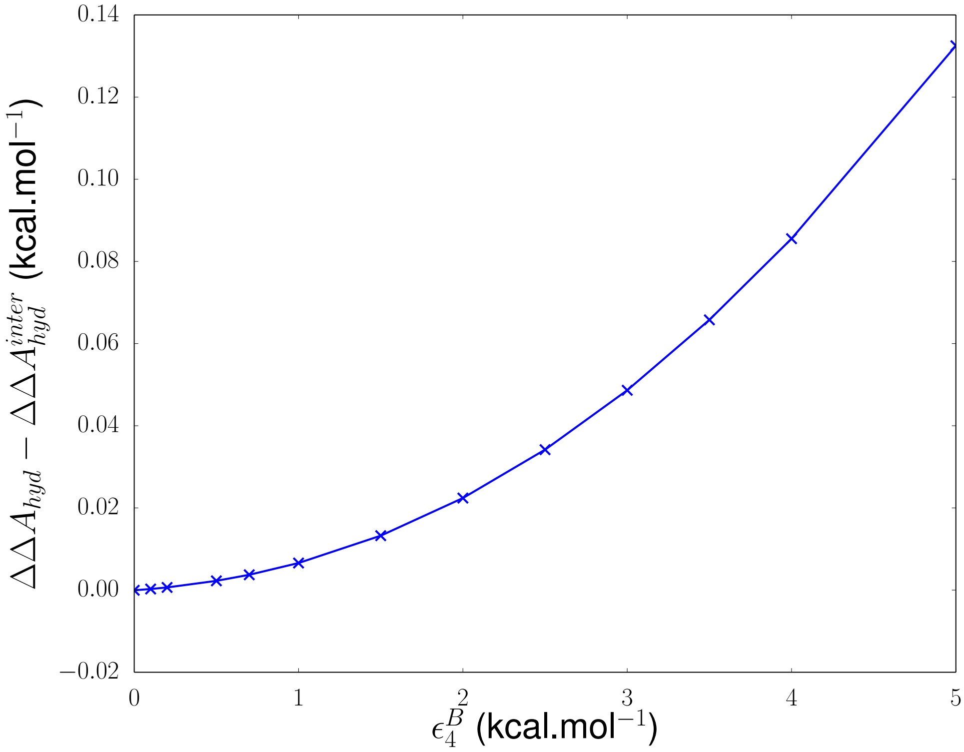

Figure F.2.: Model system calculations with perfect separability of state 3 and variable separability of state 4 i.e. ϵ_4^B is varied whilst $\epsilon_3^B = 0$

Parameter	Value
k_{P}	50 kcal.mol^{-1}.Å^{-2}
k_{Q}	100 kcal.mol^{-1}.Å^{-2}
b^0_{P},b^0_{Q}	1.0 Å
ϵ_3^A	1.0 kcal.mol^{-1}
ϵ_4^A	2.0 kcal.mol^{-1}
ϵ_3^B	0.0 kcal.mol^{-1}
ϵ_4^B	var kcal.mol^{-1}
σ_3^A,σ_4^A	1.0 Å
σ_3^B,σ_4^B	var Å
L	3.0

Table F.2.: Parameter sets used in the calculations with the model system for figure F.2.

G. Comparison of Protonation Software Packages for CD1c Structure Setup

The resolution of the electron density from protein crystallography is rarely sufficient to determine the position of hydrogen atoms. This is particularly challenging for histidine residues as this group has 3 possible protonation states. The imidazole can be protonated on either the delta nitrogen, the epsilon nitrogen or both. Determining the appropriate protonation state for each histidine is therefore a required step in the preparation of protein structures for MD simulation work.

Additionally for certain residues the electron density can be ambiguous with respect to the correct rotamer. These include histidine, glutamine, asparagine. For each of these residues it is important to consider potential flipped conformations that may provide superior hydrogen bonding opportunities, or avoid steric clashes. Histidine presents a particular problem due its combination of multiple protonation states with its flat ring structure compatible with flips.

For this reason multiple software packages have been developed to automate this process. In tables G.1 and G.2 here give a brief comparison of three different pieces of software developed for this purpose, WHATIF[120], REDUCE[121] and the PROTONATE3D module of MOE[113].

Chain	Res No	Res Type	WHATIF	MOE	REDUCE
CD1c	7	HIS	✓	✓	
	11	HIS		✓	
	14	GLN		✓	
	21	GLN			✓
	27	GLN	✓		
	38	HIS	✓	✓	✓
	57	ASN		✓	✓
	82	GLN	✓	✓	✓
	84	HIS	✓	✓	✓
	87	GLN	✓		
	128	ASN	✓		
	142	GLN		✓	✓
	146	HIS		✓	
	150	HIS		✓	
	161	ASN		✓	✓
	183	HIS		✓	✓
	185	GLN		✓	
	204	GLN		✓	
	228	GLN	✓		
	229	GLN	✓	✓	
	232	GLN		✓	✓
	266	HIS	✓	✓	✓
	272	GLN		✓	
β2-M	3	GLN			✓

Table G.1.: Residues flipped with respect to the crystal structure. A tick mark denotes a flip has been carried out by the corresponding software package.

Chain	Res No	WHATIF	MOE	REDUCE
CD1c	7	HID	HIP	HIE
	11	HIE	HIP	HIE
	38	HIP	HIP	HIE
	51	HIE	HIP	HIE
	84	HID	HIP	HID
	105	HIE	HIP	HID
	146	HIE	HIE	HIE
	150	HIE	HIP	HIE
	183	HIE	HID	HID
	208	HID	HIP	HIE
	266	HIE	HIP	HIE
β2-M	14	HID	HIP	HID
	32	HIE	HIE	HIE
	52	HID	HIP	HID
	85	HIE	HIE	HIE

Table G.2.: Protonation state of all histidines within CD1c crystal structure. HIE denotes protonation of the epsilon nitrogen, HID the delta and HIP both.