TUNING THE GROWTH AND MECHANICAL PROPERTIES OF CALCITE USING IMPURITIES: INSIGHT FROM MOLECULAR SIMULATION

Alexander Broad[1,2]

[1] Department of Physics and Astronomy, University College London
[2] London Centre For Nanotechnology

Supervised by Prof Ian J. Ford[1,2]
and Prof. Dorothy M. Duffy[1,2]

A dissertation submitted for the degree of
Doctor of Philosophy of
University College London

DECLARATION

I, Alexander Broad confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Alexander Broad
September 2022

ABSTRACT

Over many millions of years, evolution has provided living organisms with the tools to control the growth and properties of materials from the molecular scale upward. One of the many ways this is achieved is through the introduction of impurities into the solution in which these materials grow. A long-term goal of materials scientists is to harness nature's control mechanisms and apply them in the world of engineering. However, these mechanisms of growth control are highly complex, and understanding them requires insight into physical processes at the molecular scale. While experiments are so-far unable to offer such a high resolution, computer simulations can be used to directly model these physical process with no limit on the resolution. Throughout this thesis, an array of computational methodologies is applied to calcite in an attempt to understand how impurities are able to drive the growth process, and ultimately alter the mechanical properties of the crystal.

A series of metadynamics simulations are applied to calcite kink sites, revealing a more complex growth mechanism in which kink-terminating ions do not initially occupy their crystal lattice sites, and only do so upon the adsorption of an additional solute.

A combination of metadynamics and Kinetic Monte Carlo simulations are used to examine the adsorption free energies and growth inhibiting properties of amino acids and polyamines, the results of which are compared directly to experiment. This offers a robust insight into the molecular mechanisms that underpin how organic molecules are able to tune the growth of calcite.

Simulations are also applied to two case studies of impure calcite. By examining lattice spacings, determining stress distributions and simulating a series of crack propagation events, insight into mechanisms through which biogenic crystals exhibit superior mechanical properties is found.

Finally, the nature of non-Markovianity when using reaction coordinates - such as those used in rare event methodologies applied throughout this thesis- are investigated. By introducing non-Markovianity into the system, barrier crossing rates in a coarse-grained system more closely resemble those in the original two-dimensional system. Furthermore, we study the breakdown in rare-events sampling when a poor reaction coordinate is used, and identify which rare-events sampling techniques are more appropriate for detecting poor reaction coordinate choices.

IMPACT STATEMENT

Crystallisation is a highly complex phenomenon which has a profound impact on the development of biological systems, and frequently appears in sectors ranging from the chemical industry to the environment, oil and gas, water and advanced materials. In the biological world, evolution has provided nature with the means to exhibit remarkable control over the mineralisation process, right down to the nanoscale. This provides them with hard structures, such as bones and teeth, tailored to their respective purposes. In the industrial world, it is a long-term goal of materials scientists to exploit this biomineralisation process to design nanomaterials with properties similar to those found in nature. Additionally, more detrimental forms of crystallisation are frequently observed in everyday life: an example of which is the build-up of limescale. In the oil industry, a build-up of scale occurs within the pipelines, ultimately hindering flow through the pipelines. The prevention of such crystallisation events is a major challenge to the industry. Identifying means through which the crystallisation process can be controlled therefore has a vast scope of academic and industrial applications.

One way in which the crystallisation process can be controlled is through the introduction of impurities which impose additional constraints on the growth process, leading to control over morphology and size. They are also able to incorporate into the crystal, resulting in altered mechanical properties. In this thesis, we use classical molecular dynamics simulations to offer insight into the molecular processes that underpin calcite growth, and demonstrate several means through which impurities can impact the growth and mechanical properties of calcite. We demonstrate that the calcite kink growth process is more complex than previously thought, consisting of a multistep mechanism in which adsorbates transition to an intermediate state, and only adsorb into the lattice site on the arrival of another adsorbate. We also demonstrate that amino acids, and several other impurities, adsorb to calcite through a combination of interactions involving amine and carboxyl groups. Their adsorption also typically results in a minimal disruption to the local water structure, and inhibits the crystal growth rate through mechanisms dependent on the kink type.

Having identified several mechanisms through which impurities can tune the crystal growth process, we use molecular simulation to examine the mechanical properties for two case studies of impure crystals. We identify several mechanisms through which impurities can tune the size and fracture toughness of the material. Our findings have provided significant insight on how impurities can tune the mineralisation process, from the growth of crystals to their final properties. We have provided a template for studying the interaction of impurities with calcite kink sites, and made progress towards developing a kinetic Monte Carlo models for modelling the effects of additive on crystal growth rates and morphologies. We have also demonstrated that

molecular simulation can be used to demonstrate the superior mechanical properties of biologically optimised crystalline materials.

LIST OF PUBLICATIONS

This dissertation includes work from the following publications, which were completed in the duration of this PhD or are currently in preparation. The publications are given in order of appearance in this thesis.

A. Broad, R. Darkins, D. M. Duffy, and I. J. Ford, "Calcite kinks grow via a multistep mechanism," *The Journal of Physical Chemistry C*, vol. 126, no. 37, pp. 15980–15985, 2022.

O. Nahi, **A. Broad**, A. N. Kulak, H. M. Freeman, S. Zhang, T. D. Turner, L. Roach, R. Darkins, I. J. Ford, and F. C. Meldrum, "Positively charged additives facilitate incorporation in inorganic single crystals," *Chemistry of Materials*, vol. 34, no. 11, pp. 4910–4923, 2022.

A. Broad, R. Darkins, D. M. Duffy, and I. J. Ford, "Adsorption of aspartate to calcite kinks and its effects on local water structure" *in preparation.*

Y.-Y. Kim, R. Darkins, **A. Broad**, A. N. Kulak, M. A. Holden, O. Nahi, S. P. Armes, C. C. Tang, R. F. Thompson, F. Marin, et al., "Hydroxyl-rich macromolecules enable the bio-inspired synthesis of single crystal nanocomposites," *Nature communications*, vol. 10, no. 1, pp. 1–15, 2019.

A. Broad, I. J. Ford, D. M. Duffy, and R. Darkins, "Magnesium-rich nanoprecipitates in calcite: atomistic mechanisms responsible for toughening in ophiocoma wendtii," Physical Chemistry Chemical Physics, vol. 22, no. 18, pp. 10056–10062, 2020.

ACKNOWLEDGEMENTS

First and foremost, I would like to offer my deepest thanks to Amy Megson for all her support, encouragement and belief in me over the years.

I also owe my thanks to my family. I would like to thank my parents, Kate and Simon Broad, for taking a genuine interest in all my work. I would also like to thank my brother, Jack Broad, and my grandparents, Daphne and William Broad.

My time at UCL over the last four years has been particularly marked by the friends I've made along the way. I would especially like to thank my A22 friends: Jack (and Kirshita) Baker, Shereif Mujahed, Anuradha Vibhakar and Jared Jeyaretnam. I would also like to offer my thanks to my A25 friends, particularly for their patient listening to my laborious discussions about this thesis.

None of this work would have been possible without the input and guidance of my mentors. I would therefore like to extend a huge thanks to my supervisors, Professor Dorothy Duffy and Professor Ian Ford, for all the time they've invested in me. I should additionally thank Ian for his excellent module on advanced statistical mechanics, which ultimately led me to this PhD. I also owe my sincere thanks to Robert Darkins for his encouragement of all my brilliant ideas (and his polite dismissal of all my rubbish ones), not to mention his replying to my vast number of emails.

Finally I'd like to thank Stevie the cat (real name Orlando), who often comes to visit me and Amy in our flat. Maybe one day his owners will read this and finally understand where their cat keeps going.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

1.1 Materials science: an outlook

Over millions of years, evolution has provided living organisms with minerals with bespoke properties which are crucial for their survival. Such materials include shells, [4] skeletons, [5] and teeth [6]. Growing these materials requires an enormous level of control over their structure and chemical composition from the nanoscale up.

Similarly, the fashioning of bespoke materials has been pivotal in the development of human civilisation. From the beginnings of the copper, bronze and iron ages to the modern day, our understanding of material production has determined humanity's development. Our material production has typically involved high temperatures and the mining of necessary minerals. On the other hand, biological systems are limited to the chemical composition and ambient conditions of their environment when forming biominerals. Nevertheless, evolution has provided living organisms with the means to produce a range of highly sophisticated materials with properties tailored to their respective purposes.

The key difference between materials found in the industrial world and those found in the biological world is the profound control of the crystallisation process, right down to the nanoscale, which living organisms possess. A long-term ambition of materials scientists is to exploit the mechanisms such as those found in nature in order to design nanomaterials (i.e. materials with structural and mechanical features tuned to a level of precision on the order of nanometers) with tailored structures and mechanical properties. This perhaps seems ambitious, but nature has already proved to us that it is certainly achievable.

1.2 Biomineralisation

Biomineralisation is the name given to the processes through which biological systems tune the crystallisation process to produce materials with bespoke structural and mechanical properties. As mentioned previously, bones and teeth are a real-world example of materials formed through biomineralisation.

There are several means through which the crystallisation process can be directed to produce biominerals. Examples of this include growing the crystal within a confined region, such as some kind of mould, or changing the chemical composition of the solution i.e. strategically adjusting the supersaturation of the solution or introducing impurities which either modify the

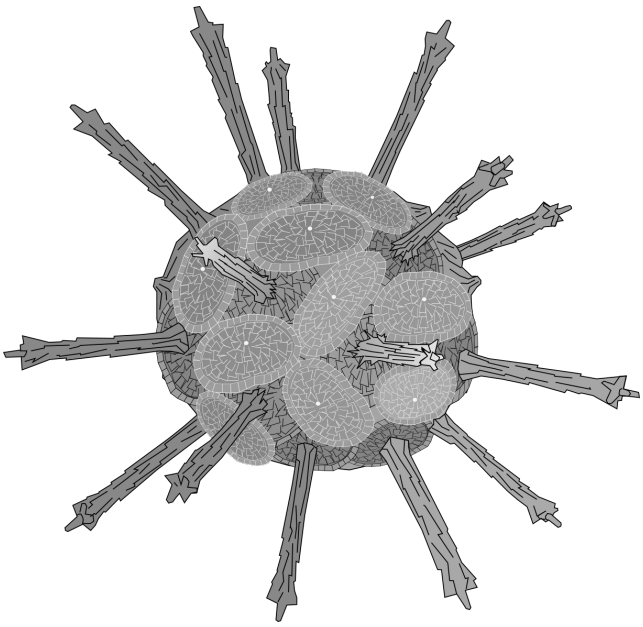

Figure 1.1: Depiction of the coccolithophore Rhabdosphaera clavigera. The plates coating the coccolithophore, as well as the protruding spines, entirely consist of calcium carbonate This provides an elegant, yet relatively simple, example of biomineralisation. Original image can be found in [1].

morphology of the crystal or change the chemical composition of the material itself. Impurities, and their effects on the crystallisation process, will be a major focus of this thesis.

There are many materials which can constitute biominerals. These include, but are not limited to, sulfides, sulfates, oxides, hydroxides, phosphates and carbonates [7]. The most abundant biomineral on Earth is calcium carbonate ($CaCO_3$). The most stable $CaCO_3$ polymorph is calcite, which is present in living organisms such as coccoliths [1] and sea urchins [8, 9]. One particularly elegant example is *Rhabdosphaera clavigera*, which is a species of coccolithophore coated in protective plates and protruding spines made of calcium carbonate. Calcium carbonate, particularly calcite, is the subject of the majority of this thesis, and will be discussed in detail in the following chapter.

While Rhabdosphaera clavigera provides an example of the extraordinary control biological systems exhibit on the crystallisation process, it is still a relatively straightforward example of biomineralisation. Elsewhere in the biological world, materials such as shells [10, 11] and bones [12] are specifically designed for their toughness and resilience to fracture. This is often achieved by introducing a series of complex, hierarchical layers of hard and soft materials [13]. It is no surprise that such materials have a far more complex chemical composition than those of Rhabdosphaera clavigera.

Impurities have a huge scope for altering crystal growth, morphologies, chemical composition and mechanical properties, and we are a long way off

from truly understanding their microscopic and macroscopic impacts. In the present day, we are typically limited to studying single crystals, grown in the presence of impurities. While single crystal biominerals containing impurities incorporated into their lattice are less commonly found in nature, there are certainly examples which can be studied, as we will discover in this thesis. Nevertheless, much remains to be understood even for these relatively simplistic settings.

As mentioned above, the vast majority of this thesis is dedicated to studying calcite. Recent advances in experimental techniques have allowed the documentation of certain impurities' effects on the growth process and morphological and mechanical properties of calcite. For example, impurities are already shown to produce impressive morphological impacts on calcite [14]. Experimental proceedings alone lack the resolution to examine the molecular processes behind these effects. However, this issue can be resolved using computer simulations to model these molecular processes. Significant advances in both experimental and computational techniques have been made in the last few decades, although there still exists a large gap in what is possible to measure in the realms of theory and experiment.

The objective of this thesis is to develop our understanding of the impacts of impurities on calcite growth from a theoretical perspective. Ultimately, we wish to develop a model which can predict observables such as morphology and mechanical properties, so that experimentation and computational modelling can be better applied in tandem in our pursuit of understanding biomineralisation. In order to achieve this, we must begin by finding greater consistency between computational predictions and experimental results. This is currently difficult to achieve, due to the large gap between accessible time and length scales of experimental procedures and computational modelling. However, recent advances in force field parametrisation, rare event sampling techniques and coarse-graining methods provides us with an opportunity to make thermodynamically consistent measurements at the atomic scale and expand upon them to make predictions of macroscopic observables which are directly comparable with experiment.

The objectives of this thesis can be split into several components. First of all, we wish to use molecular simulation to model the interactions of adsorbates, such as Ca and CO_3 units and certain impurities, with the surfaces of calcite most relevant to growth. Doing so will provide details on the molecular mechanisms behind calcite growth, as well the mechanisms through which certain functional groups (e.g. carboxyl and amino groups) bind to calcite to alter its growth. We also wish to use these results to build a coarse grained model of impurity-directed calcite growth which can be used to validate the values calculated at the molecular level, and can be directly compared against experimental results. Finally, we wish to use molecular simulation to model examples of impure calcite which have been shown to exhibit superior mechanical properties to pure calcite. In doing so, we wish to advance our understanding of biomineralisation by elucidating the molecular mechanisms

responsible for the superior properties of certain examples of biogenic calcite.

1.3 Thesis outline

To summarise this chapter, we currently have little understanding of how impurities are able to tune the crystallisation process of calcite, from their interaction with kink sites and effect on growth, to their resulting macroscopic effects on the mechanical properties on the host crystal. The study of the latter is hindered by a lack of understanding of how calcite growth itself occurs at the molecular scale.

Chapter 3 provides an overview of the theoretical tools applied throughout this thesis. Starting with the fundamental statistical mechanics of thermodynamic systems, this chapter introduces the theoretical premise of Molecular Dynamics, as well as the dynamics of coarse grained systems including the Langevin Process and the Kinetic Monte Carlo method.

Chapter 4 discusses the implementation of the analytical and computational techniques used throughout this thesis. As well as discussing the parameters used when implementing the methods discussed in Chapter 3, this chapter discusses the simulation cell set-ups and introduces several techniques such as calculating adsorption free energies, stress distributions and average lattice spacing.

Chapter 5 lays the groundwork for studying the interaction of impurities with calcite kink sites by examining the binding configurations and adsorption free energies of Ca and CO_3 units to kink sites. In doing so, we provide new insight into the stable configurations of propagating kinks, and address the question of whether cation dehydration is a rate-limiting process for calcite kink propagation.

Chapter 6 concerns the interaction of impurities with kink sites. We begin by studying the binding of amino groups to calcite CO_3-terminated kink sites. This work is carried out in order to complement experimental findings, which find that amine-terminated molecules are able to occlude in calcite at a high efficiency. We then proceed to investigate the binding of aspartate to calcite kink sites, something which has frequently been the subject of experimental study. By calculating the adsorption free energy to every kink site, we provide insight into the mechanisms through which amino acids bind to calcite and inhibit its growth. We also use our results to parametrise a Kinetic Monte Carlo model, which is used to directly replicate experimental measurements of growth inhibition.

Having examined the nature of impurities binding to a growing crystal, we proceed to examine the mechanical properties of impure calcite crystals in Chapter 7. We begin by examining the effect of weakly-interacting, hydroxyl-rich nanoparticles on the stress distribution and lattice spacing of bulk calcite crystals. We then proceed to examine the case study of Ophiocoma wendtii

by calculating stress distributions and strains of Mg-rich calcite. We additionally simulate a crack propagation event for pure and Mg-rich calcite, and compare the stress against strain of the crystal as it fractures. We identify three mechanisms in which Mg-rich calcite has an improved fracture toughness compared to pure calcite.

Chapter 8 presents a study of how a poor choice of reaction coordinate introduces non-Markovian kinetics into the system. We begin by examining the crossing rate of a two-dimensional saddle point, and compare the crossing rate of the coarse grained, one-dimensional barrier in both a Markovian and non-Markovian setting. We then proceed to run various rare-event sampling methods for the two-dimensional system, and discuss how to identify where a reaction coordinate is poorly chosen from the results of a simulation. We then apply this knowledge to the work carried out in previous chapters, and discuss the successes and shortcomings of the rare event sampling techniques and reaction cooredinates used throughout this thesis.

Finally, Chapter 9 provides a summary of the thesis and identifies future areas of research.

2 BACKGROUND

Chapter 1 provided an overview of the motivations behind this thesis. In this chapter, we provide a more in-depth discussion of the literature relevant to this thesis. We begin with a discussion on calcite, including its molecular structure and growth mechanisms. We proceed to discuss the various computational modelling techniques which have typically been applied to calcite, as well as other inorganic single crystals. We continue with a discussion on our current knowledge and understanding of the surfaces which facilitate growth, as well as the growth processes themselves. Impurities are also introduced in this chapter, with discussions on their growth-altering abilities, morphological impact, incorporation, and final mechanical impact. The computational modelling of impurities is also discussed here. This chapter concludes with a summary, detailing the current limitations of materials modelling and the current open questions which will be addressed throughout this thesis.

2.1 Calcite

As mentioned previously, materials which can constitute biominerals include sulfides, sulfates, oxides, hydroxides, phosphates and carbonates [7], the most abundant of which is Calcium Carbonate ($CaCO_3$). The most stable $CaCO_3$ polymorph is calcite, which is present in living organisms such as coccoliths [1] and sea urchins [8, 9]. The other polymorphs of $CaCO_3$ are aragonite and vaterite. Aragonite is a metastable polymorph of $CaCO_3$, which is stable at high pressures [15] and otherwise decays to calcite over geological timescales [16]. Vaterite is another metastable polymorph, which rapidly decays to calcite when exposed to water [17]. Calcite is the subject of the vast majority of this thesis, and will be introduced throughout this section.

2.1.1 Morphology and molecular structure

The equilibrium morphology of calcite in pure solution and under ambient conditions is the rhombohedron. The rhombohedral crystal structure of calcite can be described with hexagonal axes. This hexagonal unit cell has lattice parameters $a = b = 4.990$ Å, $c = 17.061$ Å, $\alpha = \beta = 90°$ and $\gamma = 120°$ [18]. Figure 2.1 shows the hexagonal unit cell of calcite. Using the nomenclature of hexagonal axes, the lowest energy surface of calcite is the $\{10.4\}$ surface.

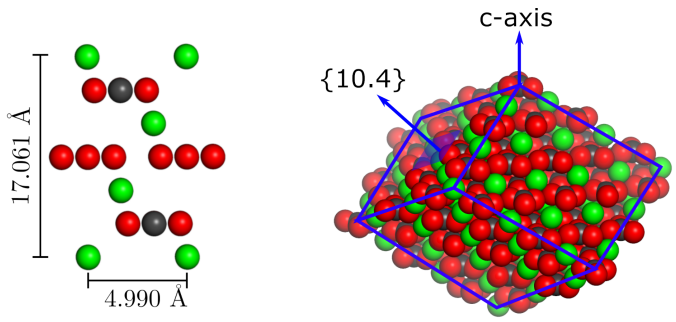

Figure 2.1: Hexagonal unit cell of calcite with lattice parameters a and c shown (left). Rhombohedral crystal structure of calcite showing the c-axis and the {10.4} plane (right). Here, Ca is shown in green, C in grey and O in red.

2.1.2 Growth mechanisms

The growth of calcite occurs through the adsorption of solutes, namely Ca and CO_3 ions, to the calcite surface. The sign and magnitude of the growth rate is determined by the *supersaturation* of the solution. Denoted by S, the supersaturation is given by:

$$S = \sqrt{\frac{a_{\mathrm{Ca}} a_{\mathrm{CO_3}}}{K_{sp}}} \tag{2.1}$$

where a_{Ca} and $a_{\mathrm{CO_3}}$ are the activities of Ca and CO_3, and K_{sp} is the solubility product, defined as the product of a_{Ca} and $a_{\mathrm{CO_3}}$ when the physical adsorption and dissolution reactions are at equilibrium, and equal to $10^{-8.48}$ for calcite [19]. If S is above one, the crystal will grow; if S is below one, the crystal will dissolve.

The mechanism of calcite growth is known to be described by the terrace-step-kink model [20], in which elevated islands of ions propagate along the surface [21]. This flat surface is known as the terrace. The propagation of the island along the (10.4) directions produces a step where the island terminates. The propagation of steps occurs through the nucleation of a row of ions to form a new step. The lattice sites where the row of ions terminate are known as kinks. An important parameter in kink nucleation is the *critical length* of the step, defined as the step length which maximises the free energy increase relative to the absence of the new step segment. Below the critical length, the step favours dissolution; beyond it, the step favours growth. The critical length is a function of the supersaturation of the solution: the greater the supersaturation, the smaller the critical length becomes.

Figure 2.2 shows a schematic detailing the mechanisms behind calcite growth. The dominant mechanism of growth actually depends on the supersaturation of the solution. At low supersaturation, growth primarily occurs through pre-existing steps. These steps are sourced through the existence of

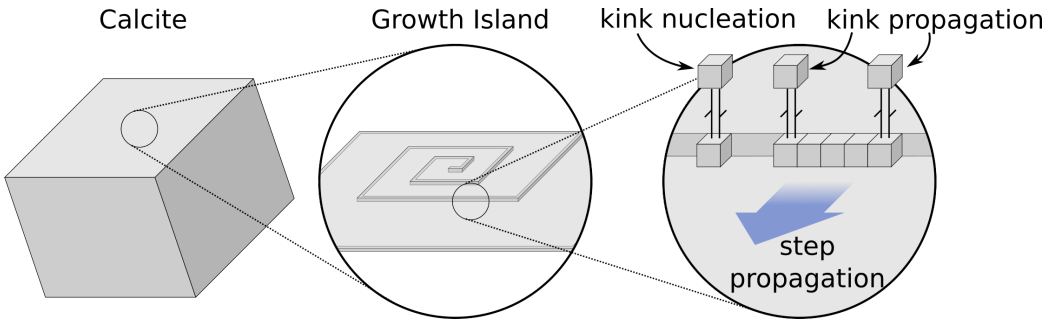

Figure 2.2: Schematic detailing the growth of calcite at low supersaturation. Here, the steps are continually sourced by the existence of screw dislocations. Steps propagate through kink nucleation, where ions adsorb to the step. The row of ions expands in a process known as kink propagation.

screw dislocations, [22] which are atomic-scale defects on the terrace which provide a local offset of the crystal lattice [23]. At high supersaturation, the two-dimensional nucleation of growth islands becomes sufficiently rapid such that propagation of growth islands becomes the dominant mechanism of calcite growth.

2.1.3 Terraces, steps and kinks

Due to the rhombohedral morphology of calcite, there are two unique geometries of step which can exist on the surface. These two steps are labelled acute and obtuse. A further consequence of this is that a total of 8 kink geometries exist. However, due to the glide plane symmetry of calcite, only four unique kink geometries exist. Since a kink can be terminated by either Ca or CO_3, the total number of unique kink types doubles. Finally, considering the two distinct orientations of CO_3 ions, the number of unique kink types doubles again to a total of 16. Throughout this study, we refer to the different kink geometries as a, b, c and d. The two CO_3 orientations we denote by (i) and (ii). If the kink is Ca-terminated, the adjacent terminating CO_3 ion determines whether the kink is labelled (i) or (ii). Figure 2.3 details the four unique kink geometries that arise from the morphology of calcite, as well as the four unique ion orientations which determine the kink type.

2.2 Modelling techniques

Computational materials modelling exists in many forms, each with their own resolution and accessible time- and length-scales. This section provides a brief description of the computational methods used to model crystallisation processes, and discusses its use and applicability in the field of crystallisation. For a more extensive discussion of the various models discussed here, Chapter 3 provides an overview of the theory behind the modelling technique used in

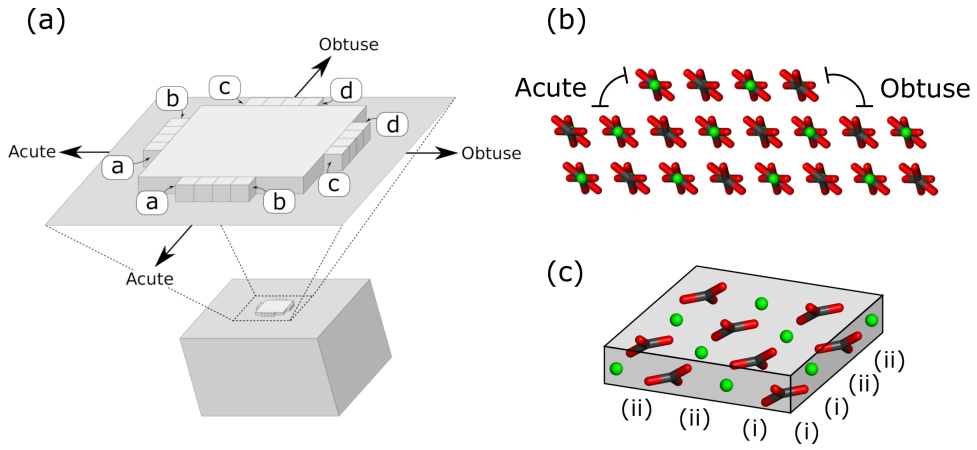

Figure 2.3: (a) Schematic of a calcite crystal with a magnified growth island, detailing all kink geometries and their labels. (b) Image detailing the geometry of acute and obtuse steps. (c) Schematic detailing the four subsets of kink type as described in the main text.

this thesis.

Ab initio methods, such as Density Functional Theory (DFT) are means of approximating the solution to the Schrödinger equation for many-body systems. Such methods have been applied in this field in the past. However, such methods based on quantum mechanics are computationally demanding, and simulations are typically limited to very small time-scales (up to a few 10s of ps). Simulating the aqueous surfaces which facilitate growth also becomes especially difficult with quantum mechanical models, which can typically only accommodate surfaces in vacuo. For this reason, computational studies of calcium carbonate applying quantum mechanics have been limited to isolated bulk materials and surfaces [24, 25].

Instead, molecular simulation, which approximates atoms as point charges obeying Newtonian mechanics, is far less computationally demanding and is a popular choice for modelling materials such as calcite. The principles of molecular simulation will be discussed in detail in Chapter 3. Classical molecular simulations are typically able to simulate the aqueous conditions in which calcite grows for time-scales of up to a few μs. Since crystal growth typically only involves the adsorption and dissolution of ions to growth sites (i.e. no chemical reactions), it is feasible that quantum mechanics may be neglected when modelling it. However, one disadvantage of neglecting chemical reactions is that the effects of the solution pH, for example the prevalence of bicarbonates, are not accounted for. Ab initio molecular dynamics are able to overcome such issues, [26] although this method is again very computationally demanding. Classical molecular simulations therefore constitute the majority of computational studies of calcium carbonate.

In practice, crystal growth consists of a vast number of physical reactions, and experimental observables such as morphology, impurity content etc. are

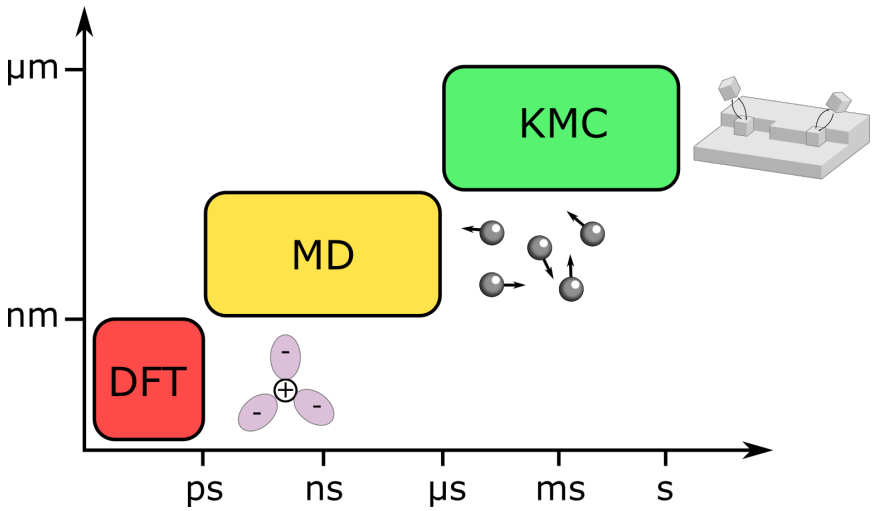

Figure 2.4: Computational methods typically applied to modelling crystallisation processes, as well as the time- and length-scales over which each method applies.

measured at length-scales beyond the reach of molecular simulation, as well as requiring time-scales far beyond the reach of molecular simulation to observe these growth processes. In these instances, modellers are forced to abandon atomistic techniques, and use either an analytical model or a coarse-grained computational model. Of the latter, a popular model in the field is the Kinetic Monte Carlo (KMC) model. KMC simulations are used to model the growth and dissolution of the crystal as a series of individual adsorption and dissolution processes, each with their own specified reaction rates (see Chapter 3). KMC simulations are an extremely useful technique, as they can produce results which can be directly compared with experimental reality. However, the results are highly dependent on the individual reaction rates used, which can be difficult to parametrise. Often, simplistic models are used to derive reaction rates. This will be elaborated on later in this chapter.

2.3 Growth and surfaces

A significant limitation of experimental studies of materials is the resolution that experiments are able to provide. The entire crystallisation process, from the nucleation and growth of crystals to their mechanical and morphological properties, are determined entirely by molecular-scale interactions and processes, which are either difficult or impossible to examine experimentally. However, computational modelling offers an invaluable insight into such molecular details. Unlike experimentation, computational modelling has no limit on its resolution.

2.3.1 Morphologies

As discussed in Section 2.1.1, the natural morphology in calcite is the rhombohedron. While this is well established for pure calcite from observation, morphology remains an important area of study in the field of crystal growth. This is for two reasons: first, crystal structures can be complex, and the prediction of morphologies using a computational model is often a good gauge of the success and applicability of the model; second, morphologies can be highly dependent on solution conditions, particularly when impurities are present in solution. The latter will be discussed in Section 2.4.

Assuming no significant kinetic effects are taking place, the morphology of a crystal is that which minimises the total Gibbs free energy of the surfaces, ΔG_{surf}, as follows,

$$\Delta G_{\text{surf}} = \sum_i \gamma_j \mathcal{A}_j \tag{2.2}$$

where γ_j is the free energy per unit area of crystal face j, and $\mathcal{A}_j$ is its surface area. In 1901, Wulff stated that the length of the vector drawn between the centre of a crystal to the crystal face, h_j, is proportional to γ_j [27]. This was originally stated without proof, although proof was later supplied by Herring [28]. With this condition, we arrive at the Wulff theorem:

$$\frac{\gamma_1}{h_1} = \frac{\gamma_2}{h_2} = \frac{\gamma_3}{h_3} = ... \tag{2.3}$$

The Wulff theorem can therefore be used to predict the crystal morphology by calculating γ_j for each surface. This process is known as Wulff construction, and is often used to predict crystal morphologies. In the case of calcite, the $\{10.4\}$ faces are expressed almost exclusively. One should therefore expect h_j to be the the smallest for $\{10.4\}$ faces, and thus γ_j is the lowest for the $\{10.4\}$ faces. In fact, early molecular simulations of ionic crystals often focussed on calculating energy differences or using energy minimisation techniques to predict surface morphologies [29, 30, 31]. Such techniques were first applied to calcite by de Leeuw and Parker, who calculated hydration energies for different calcite, aragonite and vaterite surface [32]. Their results showed the $\{10.4\}$ surface of calcite to have the lowest energy, in agreement with the natural morphology of calcite. A similar study calculated and compared the equilibrium morphologies of calcite and dolomite using Wulff construction [33].

In a later study, Piana et al. [34] used a combination of direct calculation of reaction rates using molecular simulation, and KMC to predict the equilibrium morphology of urea. The calculations agreed well with the observed morphology, heralding a success of molecular simulation being directly used to predict a macroscopic observable. However, this study was possible for urea because of its rapid growth rate. For calcite, the growth is far slower, and cannot be calculated directly using molecular simulation. A more generalised method for determining morphologies was later introduced by Li

and Doherty [35] to calculate steady-state morphologies of paracetamol in different solvents, using a mechanistic Spiral Growth model to parametrise individual attachment rates. They found general agreement with experimental observations, and noted the applicability of their method to other organic molecules. More recently, Anderson et al. [36] proposed a generalised method for predicting crystal morphologies based on a unified kinetic three-dimensional partition model. This method has been successfully used to predict morphologies of calcite, urea, L-cystine, zeolites and metal organic frameworks, and has recently been developed into the CrystalGrower software [37].

2.3.2 Crystal-water interfaces

The structure of water coordinated with calcite surface sites has been an area of interest over the past few decades, and it is widely believed that the water structure near surfaces plays an important role in crystal growth. A 2004 study by Kerisit and Parker computed free energies of water adsorption to the calcite {10.4} surface by calculating water density profiles [38]. Free energies offer a better insight into the stability of a configuration than energies, as they also include an entropic term (see Chapter 3). Indeed, it was found that the free energy of adsorption was greater (smaller in magnitude) than the energy of adsorption, indicating a loss of entropy upon adsorption onto the surface. In this study, the residence times of water molecules interacting directly with the surface were also calculated to be about 300 ps. A similar study also noted a strong hierarchical layering of water above the surface from the water density peaks near the {10.4} surface [39]. A first water layer exists, in which oxygen atoms in water directly bind to Ca atoms on the surface. The second layer exists through forming hydrogen bonds with the first water layer.

A similar study by Spagnoli et al. [40] extended this procedure to calcite steps and vacancies and observed disruptions in the two-dimensional water density profile which were highly dependent on the type of vacancy or step. A later study by Wolthers et al. [41] additionally considered the water structure near growth islands and etch pits. They noted subtle differences in hydrogen bonding between acute and obtuse edges. Another study by Wolthers et al. [42] calculated residence times for water molecules coordinated with terrace sites, step sites and corner sites. The residence times were found to be highly dependent on the site type, varying by up to two orders of magnitude. Water coordinated with the terrace was found to have residence times of 40 ps.

In a more recent study by De La Pierre et al. [43] the water structure and residence times were calculated for calcite terraces and steps. Calcite step Ca ions were able to host three bound water molecules, the residence times of some of which are significantly larger than those found on flat {10.4} surfaces (2 ns). It should be noted that the three calculated residence times for water

molecules at flat {10.4} terraces discussed in references [38], [42] and [43] vary by two orders of magnitude. This is due to the different interatomic potentials used in each study. The largest residence time is likely to be the most accurate, due to the superiority of the interatomic potentials used [44, 45].

The existence of strong water layers interacting with calcite surfaces has lead to the general opinion that surface cation dehydration is a rate-limiting process of calcite growth [46, 47, 48, 49]. In fact, this assumption has previously been applied to analytical models for calcite step propagation [50]. However, this claim has very little direct evidence, and recent calculations of water residence times have shown typically small water residence times for lone Ca ions [51, 52]. It should be reiterated, though, that water residence times do increase drastically for terrace and step Ca ions [53]. The extent of the rate-limiting effects of dehydration remains a somewhat open question.

2.3.3 Growth and dissolution

The adsorption of ions to terraces, steps and kinks is an important set of processes for understanding crystal growth. The growth and dissolution processes themselves are determined by the thermodynamics and kinetics of the adsorption and dissolution of individual solutes to calcite. Such processes have therefore received a great deal of attention in both experimental and computational studies. Experimentally, the formation and advance of both growth islands and etch pits have predominantly been studied using Atomic Force Microscopy (AFM) [54, 55, 56]. Computational studies have consisted of both molecular simulations of individual atomistic processes [57, 38, 40, 43], and KMC simulations [58, 59, 60, 61, 41, 56, 62, 37]. KMC methods are typically based on straightforward models describing individual adsorption and dissolution rates. One such model is the solid-on-solid dissolution model [63], which assumes the dissolution free energy linearly depends on the number of ionic bonds formed with neighbouring ions. A lone ion adsorbed to the terrace interacts with one neighbouring ion, a step ion interacts with two, and a kink-forming ion with three. An atom adsorbed into the step has four ionic bonds with nearest neighbours. Figure 2.5 shows a schematic of the solid-on-solid model.

The dissolution process of calcite at the nanoscale was first observed using AFM, in which it was found that it occurs through the formation of etch pits (inverted growth islands) which advance along the terrace [54, 55]. The results of this study, as well as the solid-on-solid dissolution model, were also used to parametrise a KMC model for calcite growth. While this study was successful, an AFM study alone offers little insight into the molecular mechanisms behind the complex dissolution process. Molecular simulation again provides a valuable insight into these mechanisms. While molecular simulation is currently unable to access the time- and length-scales needed

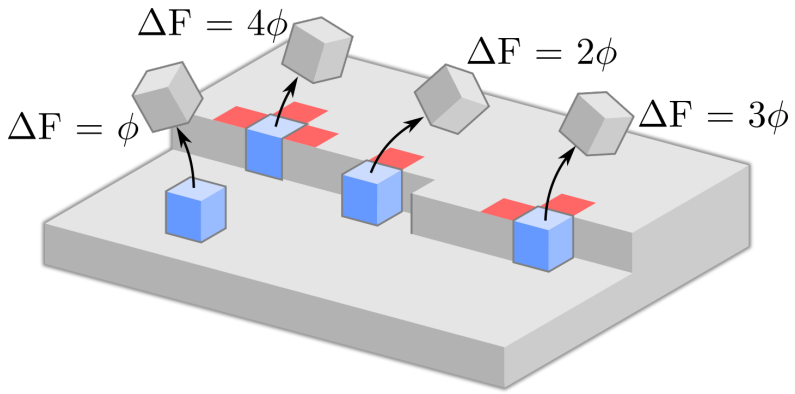

Figure 2.5: Schematic detailing the solid-on-solid model for a crystal. Here, the free energy cost of breaking a single ionic bond is given by ϕ. The free energy cost of removing an ion (shown in blue) is therefore given by ϕ multiplied by the number of ionic bonds broken. Schematic shows an ion adsorbed on the terrace, an ion adsorbed onto the step, a kink-forming ion and an ion adsorbed into the step.

to model the advancements of steps and etch pits, it is able to model the adsorption and dissolution of individual Ca and CO_3 ions to and from calcite surfaces. An early study by de Leeuw et al. [57] calculated dissolution energies for ions at calcite step sites, both in vacuo and in water. It was found that obtuse step dissolution in vacuo came at a smaller energy cost than acute step dissolution. The introduction of water was found to stabilise both steps, and the energy costs of step dissolutions were found to be similar between acute and obtuse steps. A similar study by Kristensen et al. [39] used energy minimisation techniques to calculate the kink pair defect energy, both for kink pairs along an infinitely long step, and when kinks were placed inside etch pits in a cleavage surface. It was found that kink pair formation energies were similar for different kink types. It should be noted, however, that these simulations were only carried out in vacuo, and therefore do not represent the aqueous conditions in which calcite grows.

As mentioned previously, thermodynamic averages and reaction rates are not determined solely by energy differences. Free energy differences, which contain an entropic term, determine the thermodynamics and kinetics of the system. In a study by Kerisit and Parker [38], free energy surfaces for adsorption of Ca, Mg and Sr ions to the calcite {10.4} surfaces were calculated, and it was noted that the water layers provided a large barrier for adsorption. A further study [40] determined the free energies for the adsorption of Ca and CO_3 ions to the calcite {10.4} surface, as well as acute and obtuse steps. It was found that the free energy of adsorption does not scale linearly with the number of ionic bonds formed with nearest neighbours. This threw into question the validity of the solid-on-solid model.

A later study by De Yoreo et al. [56] combined AFM and KMC to determine the step-edge structure and dynamics. Critically, it was found that

thermal fluctuations of steps and the kink density on growth islands were not sufficiently large such that they produce an abundance of kink sites. This led to the conclusion that the growth of calcite is limited by the nucleation of kinks, rather than their growth. This conclusion, however, is limited by the validity of the AFM method used. Several discrepancies have been noted between the results of AFM studies and theoretical models [64, 65, 19] A recent study by Darkins et al. [62] demonstrated that the application of AFM can hinder mass transport to the surface, thus producing incorrect critical lengths.

Since there exists a total of 16 unique calcite kink sites, as well as four unique step sites, it is likely that calcite step propagation is a complex process, determined by a large array of reaction rates. As discussed above, the solid-on-solid model is a simple model assuming that the dissolution rate is entirely given by the number of ionic bonds formed with nearest neighbours. This relies on the assumption that the dissolution rate is equal across all four step types and all sixteen kink sites. This model is known as the Kossel model [66].

Over the last several years, many AFM studies of calcite step velocities [58, 59, 60, 61] have found that varying the solution stoichiometry (ratio of Ca and CO_3 concentrations) at constant supersaturation produces step velocity functions which do not peak at $[Ca]=[CO_3]$. The acute step velocity peaks for $[Ca]<[CO_3]$ and the obtuse step velocity peaks for $[Ca]>[CO_3]$. This implies that acute kinks are more likely to be terminated by Ca ions, and obtuse steps are more likely to be terminated by CO_3 ions. A study by Andersson et al. [50] used these step velocity functions to parametrise an analytical model of calcite step growth. It was found that the stoichiometry curves could only be accurately replicated when the adsorption of ion pairs and ion chains onto steps was considered. It should be noted, however, that this study relied on the assumption that the step velocity is determined solely by the kink nucleation rate, a consequence of the results of the AFM studies by De Yoreo et al. mentioned above. Since the mass transport effects noted by Darkins et al. exist for calcite in AFM studies, it is likely that these assumptions are not valid. Furthermore, more recent KMC models have been able to replicate the stoichiometry curves by considering only single-ion attachment.

Crucial to understanding the molecular mechanisms behind calcite growth is the study of calcite growth sites: chiefly, the study of the step and kink sites which facilitate growth. Experimental techniques are currently unable to provide a level of precision beyond measuring step velocities. Studies of stoichiometry have demonstrated that calcite growth is a non-Kossel set of processes, although little has been known until recently about the thermodynamics and kinetics of individual steps or kinks. Some previous studies mentioned above have examined the dissolution of ions from steps and kinks, although these have often been limited to calculating energy (rather than free energy) differences, or are based on simulations in vacuo.

More recently, the development of rare-events sampling techniques (see Chapter 3), access to larger computing powers, and the development of thermodynamically consistent force fields has led to major advancements in the field. In 2010, Raiteri et al. [44] derived force fields able to reproduce the solubility of calcite. This was used to study the adsorption of Ca and CO_3 ions to the calcite {10.4} surface. It was found, contrary to previous studies using less thermodynamically consistent force fields, that Ca ions did not form bonds with the surface, and that CO_3 ions only interacted weakly with the surface. This was found to be due to the existence of water molecules strongly bound to the surface Ca ions, something also observed previously. In 2015, these force fields were extended to other alkaline Earth metals, including Mg [45].

In a subsequent study by De La Pierre et al. [43], the adsorption of Ca and CO_3 ions to calcite steps was examined using metadynamics. It was found that only CO_3 ions were able to form metastable bonds with the step. This was again a contrast with previous simulations, which found both ions to adsorb at steps. It was also found that the most stable binding site for a CO_3 ion was one in which it only formed bonds with the upper terrace. By tethering the CO_3 ion to its (meta)stable states, the free energy of kink pair formation was examined. The free energy of kink pair formation was found to be highly dependent on the kink type, again contradicting previous molecular simulations. It was also found that the largest kink pair formation free energy corresponds to a configuration in which the Ca ion does not reside in the lattice site. It is worth noting that a similar observation has been observed for barite, in which barium ions do not adsorb directly to step lattice sites, but instead prefer to reside above the kink site, forming two bonds with the upper terrace [67].

Over the last decade, the force fields of Raiteri et al have revealed the calcite growth process to be significantly more complex than previously thought: most notably, the weaker propensities for adsorption of Ca and CO_3 ions to all sites found compared with those of previous studies, and the more complex mechanism of kink nucleation. While this work has been extended to ion pair formation, kink nucleation and kink pair formation, very little work has been carried out on adsorption to and dissolution from calcite kink sites. Molecular studies of calcite kink sites do exist, although they have been limited to either one unique Ca-terminated kink site [52] or two CO_3-terminated kink sites [68]. These studies suggest that the dissolution rate for different kink sites is dependent on the kink type. This follows an ongoing trend of simulations demonstrating that simple models of calcite growth may not apply: from early simulations demonstrating the solid-on-solid model to be insufficient, to recent models finding kink nucleation mechanisms highly dependent on the kink pairs formed, and finding a complex set of molecular processes behind kink nucleation in which adsorbates do not reside in their lattice sites. Nevertheless, a great deal is missing from this picture, most notably the lack of molecular studies on calcite kink sites.

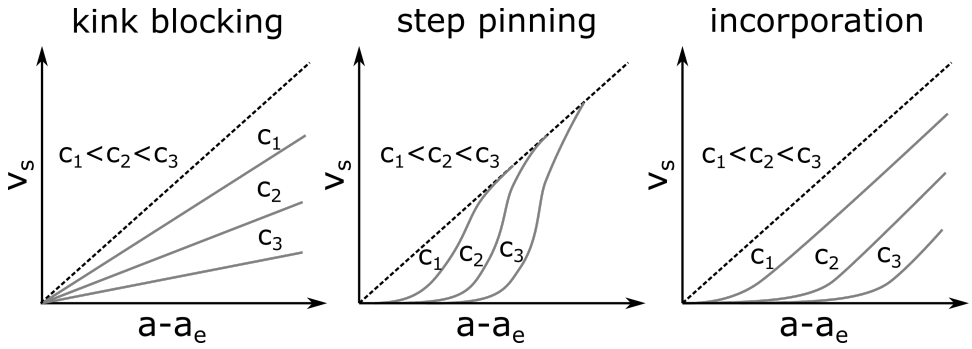

Figure 2.6: Schematics of the characteristic functions of step velocities (v_s) as a function of the difference between ion activities (a) and the equilibrium ion activity (a_e) for three separate impurity concentrations (C_1, C_2, C_3).

2.4 Impurities

Although single-crystal calcite is occasionally present in the biological world, impurities offer a far larger scope for the formation of materials with bespoke morphological and mechanical properties. The role of impurities in tuning the growth of calcite is the subject of extensive experimental and computational study and research on impurities and their influence on crystal growth has been carried out alongside research on pure calcite. There are two reasons for this: first, crystals seldom grow in pure solution, and will always contain impurities to some degree; secondly, and more importantly, impurities are one of the means through which biological systems are able to tune the growth and morphological and mechanical properties of biogenic crystals with exquisite control. This section provides a review of the mechanisms through which impurities can tune crystal growth and properties, as well as discussing the progress made in understanding them using experimental and computational techniques.

2.4.1 Effects on growth

Fundamentally, impurities can tune crystal growth at the molecular scale in several ways, each with their own characteristic effects on step morphologies and dependence of step velocity on supersaturation [20] as depicted in Figure 2.6. The strongest interaction between impurities and the crystal surface will typically occur at the kink sites, due to their potential to form ionic bonds with multiple ions. If the binding to kink sites is sufficiently small such that detachment occurs more rapidly than Ca or CO_3 units, the impurity hinders growth via a mechanism known as kink blocking [69]. A kink-blocking impurity inhibits the adsorption of either Ca or CO_3 units to the kink site. The result of this is a fractional inhibition of the step velocities independent of the solution conditions except for the concentration of the kink-blocking impurity. This relative reduction of the crystal growth rate is equal to the

fractional coverage of the kink sites, θ. The dependence of θ on the concentration of the impurity in the aqueous phase, C can be described analytically using several models, [70] the simplest of which is the Langmuir equation:

$$\theta = \frac{KC}{1 + KC} \tag{2.4}$$

where K is the affinity constant, related to the adsorption free energy, ΔG_{ads} as follows:

$$\Delta G_{ads} = -k_B T \log(KC_0) \tag{2.5}$$

where $k_B T$ is the Boltzmann factor and C_0 is the base concentration unit (which we set as the molar concentration throughout this thesis).

Adsorption free energies are experimentally calculable for calcite (although the same method can be applied to other crystals) by injecting Ca and CO_3 into the solution in order to maintain a constant supersaturation as a calcite crystal grows. The rate of injection is related to the growth rate of the crystal. This can be repeated for different kink-blocking impurity concentrations and fit to equation 2.4. This process has been applied to aspartate (asp), glycine (gly) [71] and poly-aspartate chains [72]. Asp, for example, is found to obey the Langmuir equation, indicating a straightforward, kink-blocking interaction between asp and growth sites. For glycine and poly-aspartate chains, the growth effects are more complex. The adsorption free energy is, in practice, a complex quantity which depends on an array of interactions between molecules and the calcite surface. Adsorption free energies will be elaborated on in Chapter 4, and will frequently feature in Chapters 5 and 6.

In some instances, the binding between an impurity and the crystal surface is sufficiently strong that its rate of dissolution from the surface is far lower than that of an adsorbed crystal unit (Ca or CO_3 ion in the case of calcite). Under this limit, the step is not able to grow or dissolve while the impurity is adsorbed, and the only way in which the step can advance is by growing around the impurity. This process is known as step pinning [73]. If concentration of the impurity is sufficiently small and the supersaturation sufficiently high, the step can continue its advance unimpeded. However, if the supersaturation is lowered, there is a strong non-linear effect on the step velocities, which ultimately drop to zero even though the supersaturation is high enough for growth. Another mechanism of impurity-driven growth inhibition is incorporation, in which the impurities are able to become part of the growing crystal. Impurities which are able to incorporate are believed to cause small distortions to the crystal lattice, thus increasing the internal energy of the crystal. This in turn makes the crystal more soluble, leading to an increase in dissolution and therefore a lower effective supersaturation [74]. This explains the step velocity function seen in Figure 2.6. An example of such an impurity is Mg [75].

As well as the well-characterised mechanisms of calcite growth inhibition, some interesting novel effects have been documented from AFM studies. Most notable is the anomalous growth-enhancing effects of certain impurities. This was first observed for strontium (Sr) by Wasylenki et al. [76], who used AFM to measure calcite step velocities for different Sr activities. It was found that the step velocities started to increase for smaller Sr concentrations, before reaching a peak and rapidly dropping off after a certain activity. Although the mechanisms behind these results were unclear, the increase in step velocity was initially attributed to the incorporation of the Sr ions increasing the entropy of the crystal [77] without reducing the enthalpy as in the case of Mg. However, an identical effect was later observed for poly-aspartate chains and various other peptides for obtuse steps [78]. These molecules are far less likely to be able to occlude in calcite without reducing the enthalpy of the crystal, as they are larger than both Ca and CO_3 ions and less likely to be accommodated in the lattice. However, a correlation was found between the peptide net charge and the increase in step velocity. A correlation was also found between the peptide hydrophilicity and the step velocity increase. It was proposed that these peptides were promoting growth by dehydrating the surface when they bind. This potential disruption to the water layer was theorised to improve mass transport to the surface, as other adsorbates are no longer required to cross a large free energy barrier. This mechanism differs from that proposed for the step velocity increase from Sr. It is also worth mentioning that, while the obtuse step velocities may increase, this doesn't necessarily mean the rate of crystal growth will increase. A notable example is asp_5 (a aspartic acid polymer with 5 chains), which is found to increase the obtuse step velocity [78] but not the crystal growth rate [71].

2.4.2 Morphological impacts

When crystals grow in the presence of impurities, the propagation of different kink sites is inhibited (or possibly enhanced) by the impurities. In the case of a purely kink-blocking impurity, the inhibition of the propagation of each kink type is determined by the strength and nature of the interaction between the impurity and each kink type. When the growth inhibitions are dependent on the kink type, the varying growth rates have a morphological impact on the growth island, and ultimately the entire crystal. The morphological changes that impurities enact on calcite are dependent on a vast array of thermodynamic and kinetic effects of the impurities (with multiple functional groups in the case of amino acids or peptides) and 16 unique kink sites, or 32 if the impurities have a chirality. Furthermore, the morphological changes to growing crystals occur over time- and length-scales far out of reach of molecular simulations. Studies of impurity-determined calcite morphologies therefore remain almost exclusively experimental, consisting of direct micro-scale observations of laboratory-grown crystals, Scanning Elec-

tron Microscopy (SEM) images of crystal morphologies, or AFM studies of step morphologies.

An early study by Davis et al. [75] used AFM to calculate step velocities for, and examine the morphologies of, growth islands of calcite crystals grown in the presence of Mg ions. As Mg is often present in most naturally found solutions in which calcite grows, understanding its effect on calcite growth is important. Davis et al. noted a rounding of the acute steps, and an increase in spacing between steps for the crystal growth islands. They also determined that Mg was lowering the effective supersaturation of the solution by increasing the solubility of the host crystal, something typical of an impurity that incorporates into the crystal. A similar study by Meldrum et al. [79] used SEM images of entire crystals grown on the presence of Mg in a combination with citric acid or malic acid. The morphologies produced were highly sensitive to the combinations and concentrations of each molecule, with morphologies varying between having dumbbell shapes, spherical shapes and clusters of inter-grown crystals. This study highlighted the extreme sensitivity that crystal morphologies have on impurities.

Perhaps one of the most well-known and widely studied impurities is aspartic acid (asp). Amino acids such as asp, being the building block of proteins, have received a great deal of attention over the past few decades regarding their effects on calcite growth. One study by Orme et al. [80] used AFM to determine the morphology of calcite growth islands grown in the presence of both chiralities of asp. It was found that asp caused a rounding of the acute steps, while having little to no effect on the morphology of the obtuse steps. By switching between chiralities of asp, it was possible to produce a mirror image of the morphology, where the rounding of the steps was greater on one side of the glide plane than the other. SEM images of the resulting crystals were also taken, revealing crystals with an elongation along the c-axis, and the formation of pseudo-faces (higher-energy faces which only exist because of the impurities) related to the morphology of the growth islands. The rounding of the acute steps on growth islands (in this case, when both chiralities of asp are introduced to the solution) is depicted in Figure 2.7. One additional note here is that the rounding of the step does not correspond to a set of distinct changes to the expressed surfaces, but rather a continuous change in the expressed pseudo-faces. The use of Wulff construction fails in this instance, as the morphology cannot be broken down into a series of expressed surfaces.

Intriguingly, the lack of effect on obtuse step morphology is observed regardless of the concentration of asp. This has typically been attributed to a lack of interaction between asp molecules and obtuse kinks: the consensus is that asp doesn't bind to obtuse steps and kinks. This theory, however, is rather flawed. First of all, there exists no evidence for this, other than the lack of morphological effects. Furthermore, AFM studies demonstrate that obtuse step morphologies are inhibited by asp [78], indicating that an interaction must take place between asp and obtuse kinks. Additionally, AFM has also

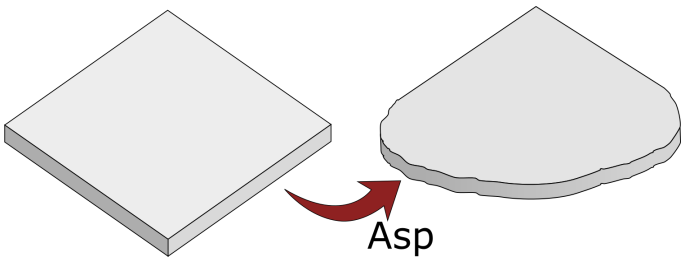

Figure 2.7: Schematic of calcite growth island grown with and without asp included in the solution. When asp is included, a rounding of the acute steps is observed. However, the obtuse step remains unaffected by asp. For this reason, asp is often assumed to not bind to obtuse steps.

revealed growth island morphologies for larger polyaspartate change, and found a complete reversal i.e. the obtuse morphologies are affected and the acute morphologies are affected less so [81].

It is clear that, while the morphological impacts of impurities are typically well documented through experiment, our understanding of the mechanisms behind them is typically lacking. Although replicating or predicting experimental morphologies with molecular simulation is far too expensive to be feasible, there is a potential scope for predicting adsorption free energies for impurities adsorbing to individual kink sites, and applying them in a KMC scheme. Previous theoretical and computational attempts to model morphology changes will be discussed in Section 2.4.5.

2.4.3 Incorporation

As discussed above, some impurities are able to become part of a growing crystal, in a process defined as incorporation or occlusion. Impure crystals are rarely more thermodynamically stable that pure crystals. Moreover, impurities typically incorporate by becoming kinetically trapped within the crystal. This occurs when the crystal grows around the impurity before it is able to escape from the surface. Although this fundamental principle is straightforward, incorporation in practice can be a complex and often counter-intuitive set of processes. Similar to morphology changes, occlusion is a process which is too computationally demanding to model using molecular simulation, although KMC and analytical models offer a means of modelling incorporation. Experimentally, incorporation is typically studied using methods such as Transmission Electron Microscopy (TEM), fluorescence spectroscopy, high-resolution synchrotron powder diffraction and direct observation of dyed crystals.

The motivation for studying incorporation comes from the fact that many biogenic crystals containing impurities exist. This has been known for many decades. Five decades ago, Towe and Thompson [82] used TEM to demonstrate that impurities exist within aragonite tablets of the nacre of the *Mytilus*

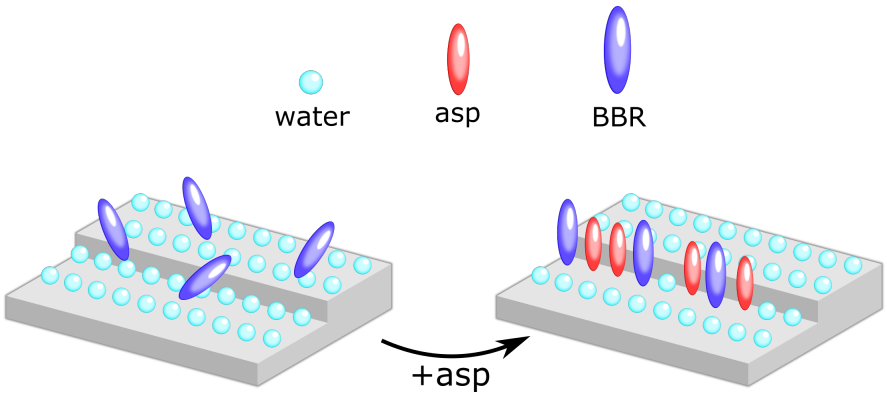

Figure 2.8: Schematic depicting the chaperoning of BBR to the surface (in this instance, the step) by asp disrupting the otherwise strongly-bound water layer.

bivalve shell. Later, Berman et al. [83] showed that proteins extracted from biogenic single crystals were able to incorporate into growing calcite crystals. Until recently, little thought was given to the mechanisms behind the incorporation process. However, 10 years ago, Borukhin et al. [84] used high-resolution synchrotron powder diffraction to calculate the incorporation efficiencies of numerous amino acids into calcite, as well as measure lattice distortions on the impure crystals. The efficiency of incorporation for each amino acid was found to vary with charge, size and pKa numbers of carboxyl and amino groups. One notable amino acid studied was asp, which was found to occlude with a greater efficiency than any other amino acid under most conditions.

Based on intuition, incorporation efficiencies are assumed to depend partially (though not entirely) on the strength of interaction between functional groups and the crystal surface. One caveat to this, however, is that the strength of interaction must not be so great that it inhibits the growth of the crystal by too large an amount [85]. Since asp contains two carboxyl groups, it has been assumed that asp has a strong interaction with calcite due to its carboxyl groups. This makes sense intuitively, as the deprotonated carboxyl group has a similar structure to a CO_3 ion, the only difference being the higher net charge (-1 compared to -2 for CO_3) and the one fewer oxygen atom. Negatively charged functional groups have often been assumed to be the primary driver of adsorption of proteins to calcite. However, recent studies have shown that occlusion mechanisms can often be rather counter-intuitive. One recent study by Kim et al. [86] showed that protein-functionalised nanoparticles are able to occlude in calcite at very high levels, despite the protein functional groups having a neutral charge and a weak interaction with the calcite surface. Another recent study by Nahi et al. [87] found that positively-charged polyamine groups are able to incorporate in calcite at high levels, throwing doubt on the consensus that negatively-charged functional groups are the primary drivers of occlusion in calcite. The study by Nahi et al. will feature extensively in Chapter 6 and the study by Kim et al. in Chapter 7.

As well as the incorporation of single impurities, combinations of impurities can produce unexpected results when introduced to the solution. In one study by Marzek et al. [88], it was observed that the dye molecule Brilliant Blue (BBR) does not occlude in calcite on its own. However, when asp was also introduced into the solution, BBR was able to occlude into the calcite to the extent of being visible under an optical microscope. Several other amino acids were also examined, and it was found that the impurities with the highest incorporation efficiencies tended to produce the largest effect. Furthermore, it was found that the amount of BBR incorporated was strongly correlated with the amount of asp incorporated. Several potential mechanisms were identified for this novel effect, the simplest being that asp and BBR complex in solution and incorporate together. This mechanism was ultimately discounted due to molecular dynamics simulations finding no association between the asp and BBR. Instead, the favoured mechanism is one where the adsorption of asp to the surface disrupts the strong water layers and effectively dehydrates the surface. This supposedly allows the BBR molecule to bypass the water layer and adsorb to the surface. This mechanism is the same mechanism discussed in Section 2.4.1, which was used to explain the increased growth rate of calcite when peptides are included in solution. It is worth noting, however, that asp was not one of the molecules studied to have this growth-enhancing effect. Recently, Nahi et al. [89] carried out a similar study, except ethanol was used as a solvent instead of including aspartate in solution. By creating a solution of 10% ethanol and 90% water, the occlusion of BBR was increased significantly compared with a pure water solution. The increased occlusion was again ascribed to a disruption of the water layers and an effective dehydration of the surface. This dehydrating mechanism has no direct experimental evidence, and remains untested by molecular simulation, largely because the processes involved are likely to exceed the time-scales accessible to molecular simulation.

2.4.4 Mechanical impacts

The means through which impurities can influence material properties is of great interest to materials scientists. Being able to tune the mechanical properties of materials commonly found in nature is a long-term ambition. Rarely in the biological world are hard structures found to contain single crystal materials. In reality, materials such bones and teeth have a far more complex chemical composition [90]. For understanding the mechanisms behind the mechanical optimisation of materials, we are often forced to turn to simple examples of single crystals containing impurities. One example is Mg, which has been found to improve hardness in calcite [91]. Amino acids have also been studied for their effect on calcite material properties. One such study by Kim et al. [92] examined the hardness of calcite grown in solutions containing either asp or glycine (gly). It was found that hardness increased when either asp or gly was introduced to the solution. The hardness of the

impure calcite increased with the amount occluded in the crystal for both impurities. This study provided insight into the mechanisms through which impurities can alter the mechanical properties of the host crystal. However, single crystals containing impurities are rarely found in the biological world, making the study of mechanisms behind mechanical optimisation difficult from both an experimental and computational perspective. However, exceptions to this do exist, and we will examine such a case study extensively in Chapter 7.

2.4.5 Computational modelling

As far as molecular simulation is concerned, impure crystals have been the subject of study as long as pure crystals have. However, due to the vast scope of interactions between impurities and the calcite surface, as well as the large time- and length-scales over which impurities adsorb, hinder growth and incorporate, simulations are often restricted to determining binding configurations and free energies for impurities on calcite surfaces. Computational studies of impurities have so-far largely been a removed area of study from experimental studies of macroscopic effects, and tenuous links are often made between the results of molecular simulation and experimental results.

Simulations of impurities and their effect on calcite date back to the early 1990's. An early computational study by Titiloye et al. [93] calculated surface energies and segregation energies (the energy difference from swapping an impurity in the bulk with a Ca or CO_3 ion on the surface) for calcite containing impurities such as Mg^{2+}, Li^{+} and HPO_4^{2-}. They found general qualitative agreement with experimental findings, although the calculated energies were not found to be accurate on a quantitative level. This was attributed to the limitations of fitting force fields: the force fields for calcite were originally fitted to elastic and vibrational data by Pavese et al. [94], and the force fields for the impurities were fitted using Gilbert's approximation method [95]. While these were validated by calculating elasticity tensors to generally good agreement, they are not designed to produce accurate surface energies. This highlights a general caveat to using force fields, especially for impurities: we cannot expect calculations to produce quantitative agreement with experiment.

In some instances, computational studies have been used to complement experiment. One such example is the studies of Elhadj et al. [81] who used AFM to determine the morphology of calcite growth islands in the presence of different polyaspartate chains. They additionally included semiempirical quantum mechanics optimisations of various binding configurations for the polyaspartate chains studied. It was found that acute steps were preferred by smaller polyaspartate chains, and obtuse steps were preferred by larger chains. This was attributed to acute steps requiring more energy to be dehydrated compared to the obtuse steps. The results were in agreement with

experimental findings, although only steps were considered in simulations, rather than kinks. Furthermore, in this study, only interaction energies were calculated, which neglected thermodynamic ensemble averaging and entropy. These results, therefore, are not considered to be quantitatively reliable.

More recently, Aschauer et al. [96] used molecular simulation to compare the binding energy of polyaspartic acid (p-asp) with polyacrylic acid (PAA) to an acute calcite step. In accordance with experimentally-observed growth inhibitions [96], it was found that p-asp had a greater strength of interaction with the surface. Intriguingly, this was found to be due to the amino group binding with the terrace. Note that this is contrary to the conventional view that carboxyl groups drive the incorporation of amino acids discussed in Section 2.4.3. In a similar study by Nada [97], the stable configurations of asp interacting with the acute and obtuse steps, and some kinks, were explored using molecular simulation. They found that asp was able to bind directly to the acute step. For the obtuse steps, as well as the terrace, asp preferred to sit above the water layer without interaction directly with the crystal. The differences between the binding configurations were attributed to the water structure itself. It was also found that asp bound directly to both acute and obtuse kinks. In all cases, the binding was dominated by carboxyl groups interacting with Ca ions. Note the difference between these findings and the results of Aschauer et al. who found that binding is dominated by the amino group for polyaspartic acid. The difference between the results of these two studies is due to the different force fields used in each study. This highlights the extreme sensitivity of the results from simulations on the choice of force field used, as the results of these two studies are contradicting, not only quantitatively, but qualitatively.

In the last few years, modellers have made progress from simply calculating binding energies and determining binding configurations, to actually calculating adsorption free energies for certain impurities. A recent study by Stepic et al. [72] made an attempt to directly compare experimental adsorption free energies for aspartate derivatives with the results of molecular dynamics simulations. However, although they found relatively good agreement between experiment and simulation, the calculations neglected the interactions with the kink sites to which kink-blocking impurities such as asp typically bind. Furthermore, they do not consider the entropic freedom of the impurity in solution when determining adsorption free energies (the adsorption free energy is a complex quantity which will be elaborated on in Chapter 4).

Recent advances in force field parametrisation has lead to an improved consistency between experimentally and computationally derived thermodynamic quantities. Most notable is the fitting of force fields to solvation free energies by Raiteri et al. discussed earlier in this chapter. More recently still, this process has been applied to some organic molecules. A very recent study by Aufort et al. [98] used a similar method to that of Raiteri et al. to derive both polarisable and rigid-ion force fields (see Figure 2.9) for methy-

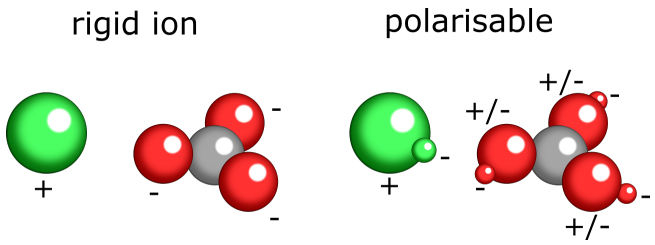

Figure 2.9: Example of rigid-ion and polarisable force fields for $CaCO_3$. An additional weightless, negatively-charged ion is included for many (or all) atom types. The weightless ion is attached to its host atom, usually with a harmonic spring. After every time step, their energies are minimised by updating their positions. This allows the existence of a spontaneous dipole.

lammonium, acetate, and zwitterionic glycine. The free energies of binding for these molecules to calcite steps was then calculated. It was typically found that the polarisable force fields produced a stronger binding, in some instances completely changing the configuration corresponding with the thermodynamic minimum. Intriguingly, both sets of force fields produced a weak binding to step sites, with binding free energies not exceeding 13 kJ/mol in magnitude.

Regarding the mechanical properties of impure crystals, computational studies of the molecular mechanisms behind impurity-directed optimisation of single crystals is relatively easy, assuming the mechanisms do not take place over length-scales significantly above the nanoscale. For example, in the study of impurity-driven hardness in calcite by Kim et al. [92], molecular dynamics simulations of asp embedded in calcite were also carried out. It was found both molecules created minor distortions on the lattice, explaining the increased hardness observed. Another study by Cote et al. [99] employed molecular dynamics simulations to examine the stress on calcite containing either asp or Mg, as a function of the strain imposed on the crystal. It was found that both impurities increased the slope of the stress-strain plot. In this area, at least, computational studies are able to provide significant insight into the molecular mechanisms that underpin the experimentally-observed effects.

Overall, computational studies and experimental studies involving impurities have largely been removed from one another, with a few notable exceptions. Of all experimental parameters, only the adsorption free energies have been directly compared between experimental and computational studies. A long term ambition in this field would be to use the results of molecular dynamics simulations to parametrise a KMC model. Such a model could be used to determine growth inhibition, morphology changes and occlusion efficiencies, which can be validated with experimental findings. However, making such progress requires a great deal of study regarding the kink sites which facilitate growth, something which has so-far been lacking in this field. Such a model must also rely on thermodynamically consistent force fields de-

scribing the interactions of impurities with calcite and water, something not available until recently.

2.5 Summary

Computational techniques for materials modelling have vastly improved in recent years. Where early molecular dynamics studies of crystalline materials such as calcite often relied on simple energy calculations or simulating materials such as calcite in vacuo, modern techniques allow molecular dynamics to be used to simulate complex aqueous surfaces and calculate free energy differences. Due to recent advances in rare-events sampling, molecular dynamics simulations are no-longer limited by the time-scales over which molecular processes occur, but rather how many collective variables are required to represent such molecular processes. Advances in force field parametrisation have allowed reliable studies of the calcite-water interface, as well as the adsorption and dissolution of ions. For processes that are too large-scale for molecular simulation, Kinetic Monte Carlo simulations allow the modelling of growth processes for time-scales of up to seconds. On the other hand, experimental proceedings have also advanced to such an extent that they can examine processes almost down to the nanoscale. For example, Atomic Force Microscopy allows measurements to be made on the propagation of growth islands on the surface, something which overlaps with Kinetic Monte Carlo Simulations in length- and time-scale.

Impurities and their effect on the growth and properties of calcite is a major area of interest within the field of crystallisation. Not only are impurities always present in the typical growth conditions for calcite, but their effect on the growth and assembly have been harnessed by nature to an extraordinary degree of precision. The effects of many impurities on the growth and morphology, as well as their occlusion and mechanical impact, are well documented in experiments. However, a large gap remains between what is observable in experiments and what is calculable in simulations. This is largely due to the significantly larger time- and length-scales over which experimentally observable properties, such as morphology changes etc, are visible. With the current available computational power and resources, the most feasible way around this would be to employ some Kinetic Monte Carlo scheme with reaction rates derived from molecular dynamics calculations. However, this has so-far been beyond the capabilities of molecular simulation. This comes down to a number of reasons. First of all, deriving thermodynamically consistent force fields comes at a large cost in time and computational power. Repeating this process for any impurity of interest is an exhausting process. Furthermore, virtually no attempts have been made to study the interactions between impurities and the kink sites which primarily facilitate growth. Indeed, studies of calcite kink sites themselves are extremely lacking, save for a single study which neglected water, as well as the entropy

difference in binding. This is further hindered by the large residence times of water interacting with step and kink sites, as well as the numerous potential binding configurations that can exist between impurities with multiple functional groups and multiple binding sites on the surface.

Additionally, much work is still required in determining the molecular mechanisms behind the advantageous mechanical properties of biomaterials. Real-world examples such as bones and teeth are far too complex to model with current capabilities. Much progress has been made in recent years regarding both the experimental and theoretical study of impure calcite, although we are still restricted to simple cases of single crystals containing impurities, and we are yet to model any systems corresponding with actual case studies of biogenic crystals with optimised mechanical properties.

Throughout this chapter, we have reviewed the structural and morphological properties of calcite, and discussed our current understanding of its growth. We have also discussed the significance of impurities and how we know they are able to tune the growth phase of crystallisation and alter final shape and mechanical properties. Having reviewed our current knowledge relevant to this thesis, we turn our attention in the next chapter to the theory applied throughout this thesis.

3 THEORY

This chapter provides an overview of the theory required to understand computational techniques for deriving thermodynamic properties of physical systems and modelling physical processes. The theory outlined in this section provides a basis for the study carried out throughout this thesis.

Section 3.1 provides an introduction to statistical mechanics, starting from the Boltzmann equation for entropy. Concepts such as free energy, the equipartition theorem and the potential of mean force are introduced.

Following on from the introduction to statistical mechanics, Section 3.2 provides an overview of modern techniques used in molecular dynamics. Beginning with a discussion on how the many-body electron system is reduced to a set of force fields for atomic interactions, this section introduces methods for dealing with long range electrostatics, integrating the classical equations of motion, and regulating the system temperature and pressure.

Section 3.4 discusses various stochastic processes such as Langevin motion, in which a particle's motion is dictated by a stochastic force. This section also covers the Kinetic Monte Carlo method, a technique used to replicate large-scale physical processes through the introduction of random events selection.

3.1 Statistical mechanics

The principles of statistical mechanics are that a connection can be made between the collective motion of particles that form a system, and the thermodynamic properties of the system. A system containing one mole of particles contains a total of $\sim 6 \times 10^{23}$ particles. Solving for the microscopic dynamics of such a system is futile, either analytically or computationally. Instead, we must consider the *ensemble averages* of the system by making a connection between the microscopic dynamics of the system and the corresponding macroscopic thermodynamics of the system.

Before proceeding, it is crucial to introduce the concept of a microstate. In a classical setting, a microstate is a single collection of particle positions and momenta. For a quantum system the microstate is a set of wave functions which solve the Schrödinger equation. However, we proceed with our discussion assuming a classical limit. Consider a system of moving particles. Recording the individual positions and momenta of each position and momentum and a given point in time will return a single microstate. Throughout this chapter, microstates are denoted by Γ. Γ represents the spatial coordinates and momentum of every particle in the system. For a three-

dimensional system, Γ is a set of $6N$ unique coordinates. Throughout this chapter, we also make reference to integrals over *phase space*. The collection of particle positions and momenta can be recorded as a single point in a $6N$-dimensional space. The evolution of the system may then be represented by the propagation of this point through the $6N$-dimensional space, which we call phase space. An infinitesimal volume of the phase space is represented by $\mathrm{d}\Gamma$, which is given by:

$$\mathrm{d}\Gamma = \prod_{i=1}^{3N} \mathrm{d}P_i \mathrm{d}R_i \tag{3.1}$$

where R_i and P_i are position and momentum components respectively. It is of further necessity to define the Hamiltonian of the system. Denoted by $\hat{\mathcal{H}}$, the system's Hamiltonian is defined as the total energy of a single microstate as follows,

$$\hat{\mathcal{H}}(\Gamma) = \mathcal{V}(\Gamma) + \sum_{i=1}^{3N} \mathcal{K}_i \tag{3.2}$$

where $\mathcal{V}(\Gamma)$ the potential energy of the system, and $\mathcal{K}_i$ is a component of the kinetic energy. For a classical, non-relativistic system, $\mathcal{K}_i$ is related to the momentum, P_i by $\mathcal{K}_i = P_i^2/2m$, where m is the mass of the particle. Throughout this thesis, only classical, non-relativistic systems are considered, so the above definition of $\mathcal{K}_i$ is assumed throughout.

3.1.1 Entropy and the microcanonical ensemble

To begin, consider an isolated box containing N moving, non-interacting particles with a total energy of E. A barrier exists which bisects the box and constrains every particle to one half of the box, as depicted by configuration A in Figure 3.1. Each particle is therefore free to occupy any point within the constrained region. Now, consider removing the barrier. Every particle is now free to occupy any point within the entire box, as is depicted by configuration B in Figure 3.1.

Now, supposing we wish to return from configuration B back to configuration A. The re-insertion of the barrier will not allow this unless every particle is residing on the left-hand-side. The probability of this occurring is easily calculated: since every particle has a $1/2$ probability of residing on the left-hand side of the middle of the box, the probability of every particle doing so becomes $(1/2)^N$. For a single mole of particles, the probability thus becomes $\sim (1/2)^{10^{23}}$. Clearly, this astronomically small number demonstrates that the re-insertion of the barrier will not allow a return from configuration B back to configuration A. Such a process may therefore be considered irreversible. It is possible to return to state A from state B by reinserting the barrier on

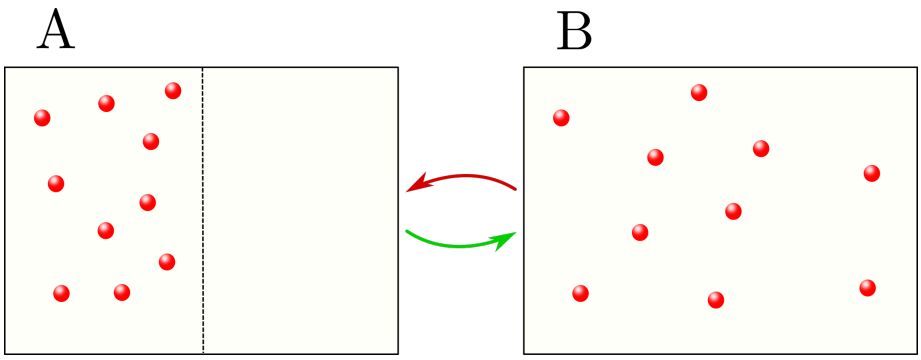

Figure 3.1: Schematic depicting a system of non-interacting particles. On the left-hand side, the particles are constrained within a region of half the total box size by a barrier. When the barrier is removed, the particles will occupy the entire box. Although there has been no change in the system potential energy, it will require an external input of energy to return the system to its original state. This process is therefore irreversible.

the right-hand-side of the box, and dragging it to the centre of the box, thus compressing the gas. However, this requires an external force, and therefore an external input of energy in the form of work done on the system.

The only physical difference between states A and B is the additional freedom of movement that the removal of the barrier allows. It is convenient to define Ω as the number of available microstates of a system. Ω is more easily understood in terms of quantum mechanics, since the degeneracy of a quantum system is always finite. For a classical system, the number of available states becomes infinite. Nevertheless, this caveat may be ignored through the underlying assumption that any classical system has an underlying quantum nature, which allows the number of available states to be finite. It is clear that an increase in Ω is associated with a change to the system which is only reversible via an external input of work being forced onto the system. This is reminiscent of a change in system entropy. It turn out that the system entropy, S, is given by a natural logarithm of Ω as follows:

$$S = k_B \log \Omega \tag{3.3}$$

where k_B is the Boltzmann constant. The natural logarithm is taken because is allows the entropy to become extensive. For example, consider two systems with Ω_1 and Ω_2 available states, respectively. Combining the two systems therefore produces a total of $\Omega_1\Omega_2$ available states. The total entropy, S_{tot}, therefore becomes

$$\begin{aligned} S_{tot} &= k_B \log(\Omega_1\Omega_2) \\ &= k_B \log \Omega_1 + k_B \log \Omega_2 \\ &= S_1 + S_2 \end{aligned} \tag{3.4}$$

Furthermore, as discussed above, an system may only increase its number of available states in the absence of external work. Since the logarithm is monotonically increasing, S may only increase unless external work is applied. The above entropy definition is therefore both extensive, and may only increase unless external work is applied. These properties are consistent with the thermodynamic definition of entropy. Equation 3.4 therefore relates the properties of individual configurations with the macroscopic property entropy.

3.1.2 Ergodicity and equal a priori probabilities

The above discussion provides a relation between microscopic dynamics to macroscopic properties. However, multiple caveats exist for this line of reasoning. First, the definition of Ω is dependent on whether the system is able to visit all available microstates. Consider a simple case of a single particle moving horizontally in a square, two-dimensional box. The particle will simply move back and forth, retracing its path. The region of accessible phase space is therefore not given by the volume of the box. Consider now a one-dimensional harmonic oscillator. Its trajectory will encompass every region of phase space as allowed by its total energy. The latter example is said to be *ergodic*. Ergodicity is defined as a system's ability to occupy every accessible region in phase space. Although a direct proof does not exist, it is widely assumed that all systems of physical interest are ergodic. Ergodicity also implies that averages over time will equal averages over microstates. A property, f, averaged over a trajectory $x(t)$, will obey the following relationship:

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t f(t)\mathrm{d}t = \int f(x(t))p(x)\mathrm{d}x \tag{3.5}$$

Ergodicity is an important principle when considering molecular simulation, which will feature in Section 3.2. Another important principle in statistical mechanics is the postulate of equal a priori probabilities. This postulate states that an isolated system in equilibrium will have an equal probability of existing in any microstate. In other words, the probability of a system existing in microstate Γ is given by:

$$P(\Gamma) \equiv \frac{1}{\Omega} \tag{3.6}$$

The postulate of equal a priori probabilities is important in using equation 3.4 to define the entropy of an isolated system. Without it, weighting must be given to the relative probabilities of different microstates when considering the total entropy of a system. This will be further discussed in Section 3.1.4. Ergodicity and equal a priori probabilities also allow us to express the total

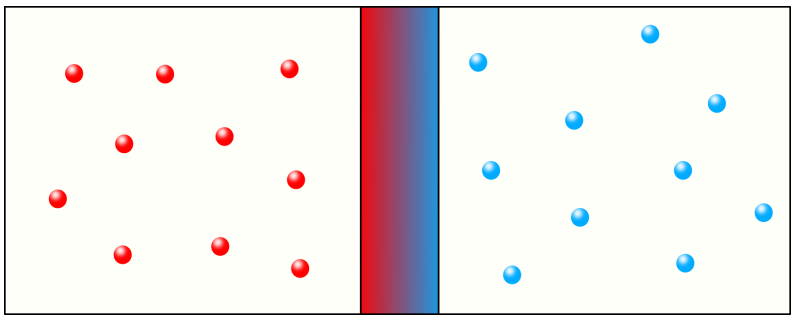

Figure 3.2: Schematic of two separate systems which are able to exchange energy with one another. The direction of the flow of energy is determined by the temperature of each system.

number of available microstates as a mathematical function of the system energy E as follows,

$$\begin{aligned}\Omega &= \left(\frac{1}{\lambda}\right)^{3N} \int \mathrm{d}\Gamma \delta\left(\hat{\mathcal{H}}(\Gamma) - E\right) \\ &= \left(\frac{1}{\lambda}\right)^{3N} \int \prod_{i=1}^{3N} \mathrm{d}P_i \mathrm{d}R_i \delta\left(\mathcal{V}(\Gamma) + \sum_{i=1}^{3N} \mathcal{K}_i - E\right)\end{aligned} \tag{3.7}$$

where δ is the Dirac delta function, and $1/\lambda$ is a fundamental phase space density. The definition of Ω becomes ambiguous in the limit of classical dynamics, where phase space is considered to be continuous. Under this limit, the number of available microstates becomes infinite. This creates a major issue for statistical mechanics, which assumes that the number of available microstates must be finite. Fortunately, quantum mechanics allows us to bypass this issue, as the degeneracy for any quantum system is always finite. For simple systems, it can be shown that $\lambda = h$ (Planck's constant). However, the value of λ is irrelevant for calculating most macroscopic properties, so we make little mention of it throughout this thesis.

Due to the constraint on the Hamiltonian, equation 3.7 is typically very difficult to solve, and analytical solutions only exist for very straightforward systems.

3.1.3 Temperature

Consider two systems of particles with total energies E_1 and E_2, as depicted in Figure 3.2. The two systems are able to exchange energy with each other, and they are otherwise isolated. As the second law of thermodynamics states, the total entropy of the two systems, ΔS_{tot}, can only increase over time. This can be written as:

$$\Delta S_{tot} = \left(\frac{\partial S_{\text{tot}}}{\partial E_1} - \frac{\partial S_{\text{tot}}}{\partial E_2} \right) \Delta E_1 \geq 0 \tag{3.8}$$

where the partial differentials have been simplified using the conservation of energy $\Delta E_1 = -\Delta E_2$. It therefore follows that:

$$\begin{aligned} \text{if } \frac{\partial S_{\text{tot}}}{\partial E_1} > \frac{\partial S_{\text{tot}}}{\partial E_2} \Rightarrow \Delta E_1 > 0 \\ \text{if } \frac{\partial S_{\text{tot}}}{\partial E_1} < \frac{\partial S_{\text{tot}}}{\partial E_2} \Rightarrow \Delta E_1 < 0 \end{aligned} \tag{3.9}$$

In other words, energy will be transferred from a system with a higher $\partial S/\partial E$ to a system with a lower $\partial S/\partial E$. It is clear, therefore, that the temperature of the system is connected with this quantity. The thermodynamic relation,

$$\frac{1}{T} = \left(\frac{\partial S}{\partial E} \right)_{V,N} \tag{3.10}$$

also demonstrates that our definition of entropy is able to reproduce properties which are consistent with the thermodynamic representation of temperature.

3.1.4 Canonical ensemble

Consider now an example of a box of particles able to exchange energy with a much larger box of particles as depicted in Figure 3.3. The two boxes are in thermal equilibrium with one another. The smaller box is labelled as the system, and the larger box is labelled as the heat reservoir. The energies of the boxes are labelled as E_1 and E_2 for the system and heat reservoir respectively. The postulate of equal a priori probabilities gives the probability of the system existing in microstate Γ_1 as being equal to the number of microstates of the joint system where the heat bath has energy E_2, $\Omega_2(E_2)$, divided by the unconstrained number of microstates, $\Omega(E)$:

$$P(\Gamma_1) = \frac{\Omega_2(E_2)}{\Omega(E)} \tag{3.11}$$

Combining the above two equations, taking the logarithms of both sides and using equation 3.4 gives:

$$k_B \log(P(\Gamma_1)) = S_2(E_2) - S(E) \tag{3.12}$$

In the limit of $E_1 \ll E$, $S_2(E_2)$ can be Taylor expanded to give:

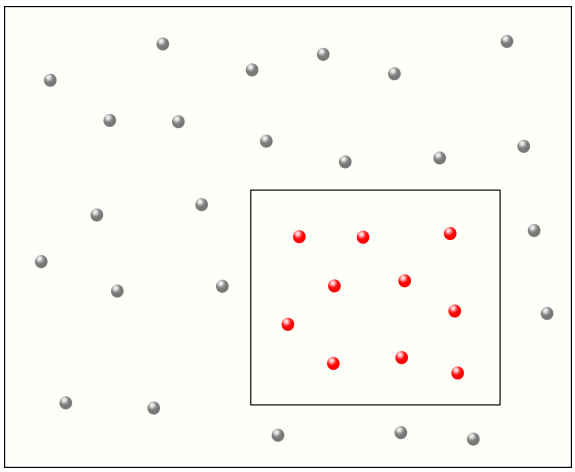

Figure 3.3: Schematic of a system able to exchange energy with a much larger system. As this other system is much larger, it can be treated as a heat bath with temperature T.

$$S_2(E_2) = S_2(E) + (E_2 - E)\left[\frac{\partial S_2}{\partial E_2}\right]_{E_2=E} \tag{3.13}$$

Noting that $E_2 - E = -E_1$ and using the thermodynamic relation $T^{-1} = \partial S/\partial E$, equation 3.12 becomes:

$$k_B \log(P(\Gamma_1)) = S(E) - S_2(E) - \frac{E_1}{T} \tag{3.14}$$

Finally, rearranging this gives:

$$P(\Gamma_1) = \frac{\Omega(E)}{\Omega_2(E)} \exp\left(-\frac{E_1}{k_B T}\right) \propto \exp\left(-\frac{E_1}{k_B T}\right) \tag{3.15}$$

The prefactor in the above equation is dependent only on the heat reservoir, and can be replaced with an integral over phase space as follows.

$$P(\Gamma) = \frac{1}{Z} \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \qquad Z = \int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma \tag{3.16}$$

Here, the $_1$ notation is dropped for convenience. The system energy E_1 has also been replaced by the Hamiltonian of the system, $\hat{\mathcal{H}}(\Gamma)$. This is because $\hat{\mathcal{H}}(\Gamma)$ will always be equal to E_1, and is calculable for a physical system. It is important to note that the integral no-longer contains the delta function as per the microcanonical ensemble. This is because the system energy is not conserved in an open system.

3.1.5 Ensemble averages

Ensemble averages refer to averages over configurations. As discussed in Section 3.1.2, ergodicity implies that ensemble averages will equal averages over time in the limit of $t \to \infty$. Since averages over time are the observable quantities in practice, ensemble averages are of importance in equating real-world observations with theory. The value of an ensemble average, denoted by $\langle A \rangle$, is given by:

$$\langle A \rangle = \int A(\Gamma) P(\Gamma) \mathrm{d}\Gamma \tag{3.17}$$

For a canonical system, the ensemble average becomes:

$$\langle A \rangle = \frac{\displaystyle\int A(\Gamma) \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma}{\displaystyle\int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma} \tag{3.18}$$

3.1.6 Gibbs entropy

Unlike a microcanonical system, a canonical system does not have a defined entropy, since energy may be transferred into and out of the system. Furthermore, entropy is an observable of the system, rather than a microstate-dependent observable. Ensemble averaging therefore cannot be used to determine the system entropy. Instead, it is possible to consider the entropy of the ensemble. The reservoir depicted in Figure 3.2 may be instead depicted as a series of M copies of the system. Each copy is allowed to exchange energy with its neighbours. The average entropy of the system may then be given as the entropy of the entire ensemble divided by the number of copies. In other words:

$$\langle S \rangle = \frac{S_M}{M} \tag{3.19}$$

where S_M is the entropy of the entire ensemble. As with the microcanonical ensemble, the following discussions are paradoxical if phase space is not discretised. It is therefore considerably more straightforward to proceed under the assumption that that phase space is discrete. It is also assumed that $M \gg 1$ such that there are a total of M_i boxes that exist in microstate Γ_i. The probability of a copy existing in microstate Γ_i is given by $P_i = M_i/M$. The total degeneracy of the ensemble, Ω_M, is given by the total number of ways in which the total system can be arranged by microstate. This is given by:

$$\Omega_M = \frac{M!}{\prod_i M_i} \tag{3.20}$$

Taking the logarithm and using Stirling's approximation, $\log(X!) = X\log(X) - X$, Ω_M can be simplified to:

$$\begin{aligned}
\log(\Omega_M) &= M\log M - M - \sum_i (M_i \log M_i - M_i) \\
&= M\log M - \sum_i (M_i \log M_i) \\
&= \sum_i M_i (\log M - \log M_i) \\
&= -\sum_i M_i \log\left(\frac{M_i}{M}\right) \\
&= -M\sum_i P_i \log P_i
\end{aligned} \tag{3.21}$$

Therefore, $\langle S\rangle$ may be given by:

$$\langle S\rangle = -k_B \sum_i P_i \log P_i \tag{3.22}$$

This is the definition of Gibbs entropy in discrete phase space. In the limit of $M \to \infty$, this may be rewritten as an integral over phase space:

$$\langle S\rangle = -k_B \int P(\Gamma)\log[P(\Gamma)]\mathrm{d}\Gamma \tag{3.23}$$

where $P(\Gamma)$ is now a probability density. It must not be forgotten, however, that the underlying dynamics must correspond to a discrete phase space. It is worth noting that, in the microcanonical limit, $P(\Gamma) = 1/\Omega$, and the Gibbs entropy reduces to $\langle S\rangle = k_B \log\Omega$, which is the microcanonical definition of entropy.

3.1.7 Helmholtz free energy

The thermodynamic definition of Helmholtz free energy, F, is as follows:

$$F = E - TS \tag{3.24}$$

The concept of the *Free* energy is that it represents a system's capability of performing mechanical work. As discussed throughout this chapter, a

system's entropy as well as its internal energy, determines the system's capability of releasing heat energy. From the definition of F, the following thermodynamic relations hold:

$$S = -\left(\frac{\partial F}{\partial T}\right)_V \qquad P = -\left(\frac{\partial F}{\partial V}\right)_T \tag{3.25}$$

In sections 3.1.5 and 3.1.6, a framework is given for calculating both a system's energy and entropy. For a canonical system, the energy is given by an ensemble average, and the entropy is given by Gibbs entropy. F is therefore given by:

$$\begin{aligned} F = \langle E \rangle - T\langle S \rangle &= \frac{1}{Z}\int \hat{\mathcal{H}}(\Gamma) \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma \\ &\quad + k_B T \frac{1}{Z}\int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \log\left[\frac{1}{Z}\exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right)\right] \mathrm{d}\Gamma \\ &= -k_B T \log Z \frac{1}{Z}\int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma \end{aligned} \tag{3.26}$$

The above integral is equal to Z by definition, and therefore cancels out. The Helmholtz free energy is therefore given by:

$$F = -k_B T \log Z \tag{3.27}$$

This tells us, much like energy in a microcanonical system, the free energy of a canonical system is minimised in equilibrium. It also provides a framework through which the thermodynamic properties of a canonical system may be deduced from its microscopic dynamics.

3.1.8 Equipartition of kinetic energy

The equipartition theorem states that a system's energy in equilibrium is equally partitioned between all degrees of freedom. While general proofs exist for the equipartition theorem, this thesis covers only a derivation for the equipartition of kinetic energy. Consider a single component of the kinetic energy of a single classical, non-relativistic particle, $K_j = P_j^2/2m$. Its average value is given by its ensemble average:

$$\begin{aligned}
\langle K_j \rangle &= \frac{\int \frac{P_j^2}{2m} \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma}{\int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma} \\
&= \frac{\int \prod_{i=0}^{3N} \mathrm{d}P_i \mathrm{d}R_i \frac{P_j^2}{2m} \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right)}{\int \prod_{i=0}^{3N} \mathrm{d}P_i \mathrm{d}R_i \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right)} \\
&= \frac{\int \prod_{i=0}^{3N} \mathrm{d}P_i \mathrm{d}R_i \frac{P_j^2}{2m} \exp\left(-\frac{1}{k_B T}\left[\sum_{i=0}^{3N} \frac{P_i^2}{2m} + \mathcal{V}(r_0, r_1, r_2, ...)\right]\right)}{\int \prod_{i=0}^{3N} \mathrm{d}P_i \mathrm{d}R_i \exp\left(-\frac{1}{k_B T}\left[\sum_{i=0}^{3N} \frac{P_i^2}{2m} + \mathcal{V}(r_0, r_1, r_2, ...)\right]\right)}
\end{aligned} \tag{3.28}$$

where $\mathcal{V}(r_0, r_1, r_2, ...)$ is the potential energy of the system. Here, we have made the assumption that position and momentum components are perpendicular to one another. This is a reasonable assumption for a classical system, and is applicable to the work carried out throughout this thesis. Under the assumption that the Hamiltonian is separable, the velocity integrals cancel each other out, and we are left with:

$$\begin{aligned}
\langle K_j \rangle &= \frac{\int \prod_{i=0}^{3N} \mathrm{d}P_i \frac{P_j^2}{2m} \exp\left(-\frac{1}{k_B T}\left[\sum_{i=0}^{3N} \frac{P_i^2}{2m}\right]\right)}{\int \prod_{i=0}^{3N} \mathrm{d}P_i \exp\left(-\frac{1}{k_B T}\left[\sum_{i=0}^{3N} \frac{P_i^2}{2m}\right]\right)} \\
&= \frac{\int \mathrm{d}P_j \frac{P_j^2}{2m} \exp\left(-\frac{1}{k_B T}\left[\frac{P_j^2}{2m}\right]\right)}{\int \mathrm{d}P_j \exp\left(-\frac{1}{k_B T}\left[\frac{P_j^2}{2m}\right]\right)}
\end{aligned} \tag{3.29}$$

both the numerator and denominator are Gaussian integrals which simplifies the average to:

$$\langle K_j \rangle = \frac{1}{2} k_B T \tag{3.30}$$

Equation 3.30 demonstrates that, no-matter which kinetic energy component is taken, its average will always be equal to $k_B T/2$. This is an important result which will be elaborated on in Section 3.2.9.

3.1.9 Isobaric-isothermal ensemble

In the canonical ensemble, the system is able to exchange energy with a heat reservoir at temperature T. In the isobaric-isothermal ensemble, also known as the NPT ensemble, the volume of the system is also allowed to fluctuate. At thermal equilibrium between the system and reservoir, the pressure of the system is determined by the pressure of the reservoir, P. Repeating the derivation of the canonical ensemble in Section 3.1.4, the probability, $P(\Gamma_1)$, of the system existing in microstate Γ_1 is given by:

$$P(\Gamma_1) = \frac{\Omega_2(E_2, V_2)}{\Omega(E, V)} \tag{3.31}$$

therefore:

$$k_B \log(P(\Gamma_1)) = S_2(E_2, V_2) - S(E, V) \tag{3.32}$$

same as with the canonical ensemble, $S_2(E_2, V_2)$ can be Taylor expanded under the assumption that $E_1, V_1 \ll E_2, V_2$:

$$\begin{aligned} S_2(E_2, V_2) \approx S_2(E, V) &+ (E_2 - E)\left[\frac{\partial S_2(E_2, V_2)}{\partial E_2}\right]_{E_2=E} \\ &+ (V_2 - V)\left[\frac{\partial S_2(E_2, V_2)}{\partial V_2}\right]_{V_2=V} \end{aligned} \tag{3.33}$$

Similar to the canonical ensemble derivation, we note that $V_2 - V = -V_1$ and $E_2 - E = -E_1$. We also note the thermodynamic relations $T^{-1} = [\partial S_2/\partial E_2]$ and $P = T[\partial S_2/\partial V_2]$. We finally arrive at:

$$P(\Gamma) \propto \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma) + PV(\Gamma)}{k_B T}\right) \tag{3.34}$$

where the system energy E_1 is again replaced with the Hamiltonian $\hat{\mathcal{H}}(\Gamma)$. Equation 3.34 is analogous to equation 3.34, where the probabilities are weighted by an exponential.

3.1.10 Potential of mean force

As discussed earlier, a three-dimensional system containing N particles has a total of $3N$ position coordinates and $3N$ momentum coordinates. In most systems of interest, and throughout this thesis, the momentum coordinates can be integrated out, since the potential energy of the system often doesn't depend on individual momenta. Nevertheless, the evolution of a microstate

over a $3N$-dimensional phase space is incredibly complex and difficult to interpret. For this reason, the concept of the reaction coordinate is introduced. A reaction coordinate, which is denoted by $s(\Gamma)$, is a function which maps the $3N$ positional coordinates onto a single value. Intuitively, $s(\Gamma)$ can be any function of the $3N$ positional coordinates. Examples include the position of a single particle, the centre-of-mass of several particles, or the dihedral angle between four particles.

For convenience, we begin by considering only one reaction coordinate. Consider a single reaction coordinate, $s(\Gamma)$, and its value, s. The mean force along the reaction coordinate is given by:

$$\langle f(s)\rangle = \frac{1}{Z}\int_{\Omega}\delta(s(\Gamma)-s)\left(-\frac{\mathrm{d}\hat{\mathcal{H}}(\Gamma)}{\mathrm{d}s(\Gamma)}\right)\exp\left[-\frac{\hat{\mathcal{H}}(\Gamma)}{k_BT}\right]\mathrm{d}\Gamma \tag{3.35}$$

This is simply a canonical average of the the force, $f(s) = \mathrm{d}\hat{\mathcal{H}}(\Gamma)/\mathrm{d}s$ over all coordinates perpendicular to the reaction coordinate, hence the Dirac delta function, $\delta(s(\Gamma) - s)$. It is convenient to define the configurational integral, $Q(s)$, as:

$$Q(s) = \int_{\Omega}\delta(s(\Gamma)-s)\left(-\frac{\mathrm{d}\hat{\mathcal{H}}(\Gamma)}{\mathrm{d}s(\Gamma)}\right)\exp\left[-\frac{\hat{\mathcal{H}}(\Gamma)}{k_BT}\right]\mathrm{d}\Gamma \tag{3.36}$$

such that:

$$\langle f(s)\rangle = \frac{1}{Q}\left(k_BT\frac{\mathrm{d}}{\mathrm{d}s}\right)\int\delta(s(\Gamma)-s)\exp\left[\frac{-\hat{\mathcal{H}}(\Gamma)}{k_BT}\right]\mathrm{d}\Gamma = k_BT\frac{\mathrm{d}\log Q}{\mathrm{d}s(\Gamma)} \tag{3.37}$$

The Potential of Mean Force (PMF), which is labelled $\Phi(r)$, can therefore be defined as:

$$\Phi(s) = -k_BT\log Q(s) \qquad \text{such that} \qquad \langle f(s)\rangle = -\frac{\mathrm{d}\Phi(s)}{\mathrm{d}s(\Gamma)} \tag{3.38}$$

Now, consider a canonical probability density as a function of s such that its integral over s is one. This probability density function (PDF), labelled $\rho(s(\Gamma) = s)$ can be given by:

$$\rho(s) = \frac{1}{Z}\int\delta\left(s\left(\Gamma\right)-s\right)\exp\{-\hat{\mathcal{H}}(\Gamma)/k_BT\}\mathrm{d}\Gamma \propto Q(s) \tag{3.39}$$

Equation 3.39 tells us two things. First, $\rho(s)$ is proportional to $\exp(-\Phi(s)/k_BT)$ providing a link between the PMF and the PDF. Second, equation 3.39 can

be simplified using equation 3.27, to a function which we label as $F(s)$. $F(s)$ describes the free energy of the system with the reaction coordinate constrained at $s(\Gamma) = s$. Using equation 3.27, we have:

$$F(s) = \Phi(s) + C = k_B T \log(\rho(s)) + C' \tag{3.40}$$

where C and C' are constants. Their values are fairly arbitrary in practice, since they only have physical meaning when normalising probability densities, which may be done by integrating over s.

$$\rho(s) = \frac{\exp\left(-\frac{F(s)}{k_B T}\right)}{\int_{s'} \exp\left(-\frac{F(s')}{k_B T}\right) \mathrm{d}s'} = \frac{\exp\left(-\frac{\Phi(s)}{k_B T}\right)}{\int_{s'} \exp\left(-\frac{\Phi(s')}{k_B T}\right) \mathrm{d}s'} \tag{3.41}$$

$F(s)$ is often referred to as the Free Energy Surface, or Free Energy Function. Its dependence on s can sometimes differ subtly from the PMF in practice. This will be discussed later in this chapter. Equation 3.40 provides a powerful framework for rare event sampling techniques, which will be discussed later in this chapter. Although it has only been derived for one-dimensional reaction coordinates, equation 3.40 may be arbitrarily extended to define multidimensional free energy surfaces which are functions of multiple reaction coordinates.

3.2 Molecular dynamics

It comes as little surprise to learn that the canonical partition function is highly non-trivial to solve analytically. Analytical solutions to equation 3.27 are limited to simple examples of little physical interest. It makes sense, therefore, to focus on computational approximations for calculating ensemble averages. There are two ways in which ensemble averages can be calculated approximately: Molecular Dynamics simulations, and Monte Carlo simulations. The former, in which ensemble averages are calculated through the approximation of physical trajectories, is the more popular choice of computational method, and will be discussed throughout this section. Here, we discuss atomistic molecular dynamics, in which atoms are approximated as point masses.

3.2.1 From quantum to classical potentials

In reality, the interaction between a configuration consists of a large number of inter- and intra-atomic interactions between negatively charged electrons and positively charged nuclei. The quantum states and their trajectories will obey the time-dependent Schrödinger equation:

$$i\hbar\frac{\mathrm{d}}{\mathrm{d}t}\psi(r_i, R_i, t) = \hat{H}\psi(r_i, R_i, t) \tag{3.42}$$

where $\hbar$ is the reduced Planck constant, ψ is the wave function, and r_i and R_i are the electronic and nuclear coordinates. $\hat{H}$ is the non-relativistic Hamiltonian:

$$\begin{aligned}\hat{\mathcal{H}} = &-\sum_i \frac{\hbar^2}{2M_i}\nabla^2_{R_i} - \sum_i \frac{\hbar^2}{2m_e}\nabla^2_{r_i} \\ &+ \frac{e^2}{4\pi\epsilon_0}\sum_{i<j}\frac{1}{|r_i - r_j|} + \frac{e^2}{4\pi\epsilon_0}\sum_{i<j}\frac{Z_i Z_j}{|R_i - R_j|} \\ &- \frac{e^2}{4\pi\epsilon_0}\sum_{i,j}\frac{Z_i}{|R_i - r_j|}\end{aligned} \tag{3.43}$$

where M_i are the nuclear masses, m_e is the electron mass, e is the electron charge and ϵ_0 is the permittivity of free space. The first two terms represent the kinetic energy of the nuclei and electrons, and the remaining terms deal with the electrostatic interactions. Solving this equation for a many-body wave function is extremely computational expensive, and therefore unsuitable for any system of interest. Approximations exist which simplify equation 3.43 and reduce the computational expense of solving equation 3.42. Such approximations form the basis of Density Functional Theory (DFT), a widely used method for solving the many-body Schrödinger equation. However, this method is still far too computationally expensive to simulate the physical processes discussed throughout this thesis. Instead, we make the reasonable assumption that the physical processes observed do not constitute any chemical reactions. This allows the neglect of quantum mechanics and the coarse-graining of a many-body atom to a single point charge. Instead of the many-body Schrödinger equation, Newton's equations of motion are solved for a collection of atoms approximated as point charges.

$$\frac{\mathrm{d}P_i}{\mathrm{d}t} = -\Delta_{R_i}U(\{R_i\}) \tag{3.44}$$

where R_i are the atomic positions and P_i are their momenta. U is the potential energy component of the Hamiltonian, dependent of the total set of atomic positions, $\{R_i\}$. Δ_{R_i} is the derivative with respect to R_i. By neglecting quantum mechanics, far larger timescales are accessible in a simulation. DFT is able to simulate quantum processes over timescales reaching 10s of picoseconds, whereas Molecular Dynamics simulations (i.e. solving equation 3.44) may typically simulate classical processes over timescales roughly spanning picoseconds to microseconds.

3.2.2 Force fields and partial charges

Approximating atoms as point masses comes with a major caveat. Where methods such as DFT solve equations based on first principals, classical methods such as Molecular Dynamics must coarse-grain the complex set of interactions between between nuclei and electrons, to a single set of classical interactions between atoms. One way this can be achieved is through the introduction of partial charges. Covalently bound molecules exist through the overlap of electron densities between two or more atoms. Approximating this classically becomes difficult in this instance: it is clear that the atoms over which electrons are shared do not have average charges equal to that of an integer multiple of the electron charge. In order to accommodate this, the concept of *partial charges* is introduced. Partial charges are an average charge of a covalently bonded atom, which arise through the overlap of electron densities. Returning to equation 3.44, $U(\{R_i\})$ may be given by:

$$U(\{R_i\}) = \mathcal{V}(\{R_i\}) + \frac{1}{4\pi\epsilon_0} \sum_{i<j} \frac{q_i q_j}{|R_i - R_j|} \tag{3.45}$$

where q_i are the atomic partial charges, and $\mathcal{V}(R_i)$ is a potential which encompasses Van der Waals Attraction, Pauli Repulsion and intramolecular bonds. $\mathcal{V}(\{R_i\})$ may be given as a set of potentials:

$$\begin{aligned} \mathcal{V}(\{R_i\}) = &\sum_{i,j} \mathcal{V}_2^{ij}(R_i, R_j) \\ &+ \sum_{i,j,k} \mathcal{V}_3^{ijk}(R_i, R_j, R_k) \\ &+ \sum_{i,j,k,l} \mathcal{V}_4^{ijkl}(R_i, R_j, R_k, R_l) \end{aligned} \tag{3.46}$$

where V_2^{ij}, V_3^{ijk} and V_4^{iklj} represent two-, three- and four-body interactions respectively. Such potentials may be optimised through first principles calculations such as DFT, or they may be parametrised to replicate experimental data.

3.2.3 Van der Waals attraction

Van der Waals attractions arise from the electronic structure of two interacting atoms. When two atoms interact with one another, the electronic structure of both atoms shifts in order to minimise the total electric potential. This creates an effect, where the electronic structure of both atoms both form a dipole. Since electrons move at a much faster rate than nuclei, the process of dipole arrangement can be considered spontaneous compared to Molecular Dynamics time-scales. Van der Waals attraction consists of

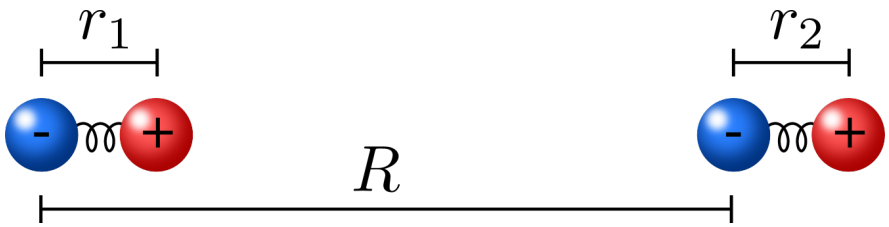

Figure 3.4: diagram of the dipole arrangement used to derive the van der Waals attraction force.

two subsets: London dispersion forces, in which two non-dipole atoms both spontaneously induce a dipole in one-another, and Debye forces, in which a pre-existing dipole induces a dipole in a non-dipole atom. Both dependences on atomic distance are similar, and both forces can be derived through the same method.

The dependence of van der Waals forces on atomic separation may be derived using a simple model consisting of a pair of dipoles, each separated by a distance of r_1 and r_2, and with charges c and $-c$. The dipoles are separated from each other by a distance of R. Figure 3.4 depicts the dipole arrangement that arises through the interaction of two nearby atoms. The positive and negative charges interact with one another through a harmonic potential with spring constant k. The total potential energy of the system is therefore given by:

$$V = \frac{c^2}{4\pi\epsilon_0 R}\left[1 + \frac{1}{1+\frac{r_2-r_1}{R}} - \frac{1}{1-\frac{r_1}{R}} - \frac{1}{1+\frac{r_2}{R}}\right] + \frac{1}{2}k(r_1^2+r_2^2) \tag{3.47}$$

By taking the limit $r \ll R$, the above can be Taylor expanded to give:

$$\begin{aligned} V &= \frac{c^2}{4\pi\epsilon_0 R}\left[1 + \left(1 - \frac{r_2-r_1}{R} + \frac{(r_2^2+2r_1r_2+r_1^2)}{R^2}\right)\right. \\ &\qquad \left. - \left(1 - \frac{r_1}{R} + \frac{r^2}{R^2}\right) - \left(1 + \frac{r_2}{R} + \frac{r^2}{R^2}\right)\right] + \frac{1}{2}k(r_1^2+r_2^2) \\ &= \frac{1}{2}k(r_1^2+r_2^2) - \frac{c^2}{2\pi\epsilon_0 R^3}(r_1r_2) \\ &= \frac{1}{2}k(r_1^2+r_2^2) - \frac{C}{R^3}(r_1r_2) \end{aligned} \tag{3.48}$$

where the constant C is defined for convenience. The Schrödinger equation with this potential is difficult to solve, due to the potential energy containing both r_1 and r_2. However, the potential energy may be recast by defining two new constants, k_s and k_a, as:

$$\begin{aligned} k_s &= k - \frac{C}{R^3} \\ k_a &= k + \frac{C}{R^3} \end{aligned} \tag{3.49}$$

The total potential is therefore given by:

$$V = \frac{1}{2}k_s r_s^2 + \frac{1}{2}k_a r_a^2 \tag{3.50}$$

where $r_s = r_1 + r_2$ and $r_a = r_1 - r_2$. The Schrödinger equation for the system may therefore given by two coupled equations:

$$\begin{aligned} -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial r_s^2}\psi + \frac{1}{2}k_s r_s^2\psi &= E_s\psi \\ -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial r_a^2}\psi + \frac{1}{2}k_a r_a^2\psi &= E_a\psi \end{aligned} \tag{3.51}$$

The above equation is the Schrödinger equation for two quantum harmonic oscillators. The ground state energies, $E_{s,0}$ and $E_{a,0}$, are given through the well known Ladder Operators by:

$$\begin{aligned} E_{s,0} &= \frac{\hbar}{\sqrt{2}}\sqrt{k + \frac{C}{R^3}} \\ E_{a,0} &= \frac{\hbar}{\sqrt{2}}\sqrt{k - \frac{C}{R^3}} \end{aligned} \tag{3.52}$$

The energy dependence on R, $U(R)$ can be calculated by calculating the difference between the sum of E_s and E_a, and the ground state energy of two independent harmonic oscillators, both with an energy given by $\hbar/2\sqrt{k/m}$.

$$U(R) = \frac{\hbar k}{\sqrt{2}}\left[\sqrt{1 + \frac{C}{kR^3}} + \sqrt{1 - \frac{C}{kR^3}} - 2\right] \tag{3.53}$$

by assuming that $C/kR^3 \ll 1$, equation 3.53 can be further Taylor expanded to give:

$$\begin{aligned} U(R) = \frac{\hbar k}{\sqrt{2}}\Bigg[&\left(1 + \frac{1}{2}\frac{C}{kR^3} - \frac{1}{8}\left(\frac{C}{kR^3}\right)^2\right) + \\ &\left(1 - \frac{1}{2}\frac{C}{kR^3} - \frac{1}{8}\left(\frac{C}{kR^3}\right)^2\right) - 2\Bigg] \\ = -\frac{\hbar k}{\sqrt{2}}&\left(\frac{1}{4}\frac{C^2}{k^2}\right)\frac{1}{R^6} = -\frac{C'}{R^6} \end{aligned} \tag{3.54}$$

The van der Waals attraction therefore decays with $1/R^6$. Equation 3.54 reduces the complex interaction between many-body quantum systems to a simple decaying function, which can be easily implemented in molecular dynamics. In practise, the charge arrangements are typically far more complex than a pair of dipoles. The pre-factor C' is therefore typically not derived from theory. Instead, it is often either parametrised from first principles calculations (DFT) or from experimental data. Van der Waals forces are weak compared with electrostatic interactions, and their significance in driving properties such as crystal structures tends to vary between materials.

3.2.4 Pauli repulsion

Pauli repulsion is a force which arises from the Pauli exclusion principle. Consider two identical atoms, the electrons of which occupy their ground state. When separated by a large distance, there will be no significant overlap between electron wave functions between the two atoms, and therefore each electron can occupy the same quantum states. However, as the two atoms approach one another, the overlap between electron wave functions increases, and the electrons are no longer able to occupy identical quantum states as stated by the Pauli exclusion principle. This forces electrons into higher energy states, thus generating a repulsive force.

Pauli repulsion does not have a dependence on atomic distance which can be analytically derived as van der Waals attraction does. Instead, it is common practice to introduce an additional interaction with a radial dependence that wins out over van der Waals attraction at short distances. A popular choice of potential that encompasses both Pauli repulsion and van der Waals attraction is the Lennard Jones potential,

$$V_{\mathrm{LJ}}(r) = 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right] \tag{3.55}$$

where ϵ and σ are parametrisable constants. The $1/r^{12}$ term represents Pauli repulsion in this instance. Another popular choice is the Buckingham potential,

$$V_{\mathrm{Buck}}(r) = A\exp(-Br) - \frac{C}{r^6} \tag{3.56}$$

where A, B, and C are parametrisable constants. Here, the Pauli repulsion is represented through an exponential.

3.2.5 Bonds, angles and dihedrals

Molecules that exist through strong covalent bonds can be assumed to suffer minimal displacement from their equilibrium configurations. Under this

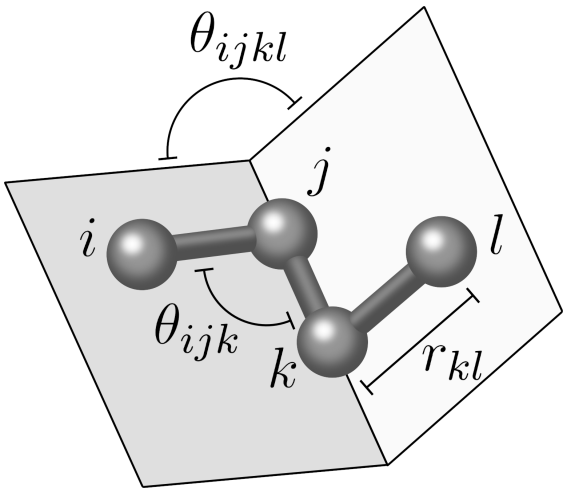

Figure 3.5: Schematic depicting a molecule consisting of atoms i, j, k and l, detailing an atomic distance r_{kl}, an angle θ_{ijk} and a dihedral angle θ_{ijkl}.

assumption harmonic potentials may be used to model their intra-molecular bonds. The complex electronic structure of a molecule, and its effect on the molecular geometry, may therefore be coarse-grained to a series of harmonic radial, angular and dihedral angular bonds. Figure 3.5 shows a schematic describing each component. The values of each are calculated as follows.

$$\begin{aligned}
\vec{r}_{ij} &= \vec{r}_i - \vec{r}_j \\
\vec{\theta}_{ijk} &= \arccos\left(\frac{\vec{r}_{ij}.\vec{r}_{kj}}{|\vec{r}_{ij}||\vec{r}_{kj}|}\right) \\
\vec{\theta}_{ijkl} &= \arccos\left(\frac{(\vec{r}_{ij}\times\vec{r}_{jk}).(\vec{r}_{jk}\times\vec{r}_{kl})}{|\vec{r}_{ij}\times\vec{r}_{jk}||\vec{r}_{jk}\times\vec{r}_{kl}|}\right)
\end{aligned} \tag{3.57}$$

The entire molecular interaction may therefore be represented by a series of harmonic forces acting on r_{ij}, θ_{ijk} and θ_{ijkl}. The spring constants of each must be carefully parametrised to reproduce the configurational freedom of the molecule.

3.2.6 Cut-off radii and periodic boundaries

The number of particles in a simulation cell determines the timescales accessible in a simulation. Consider a system of N particles. Each particle interacts with $N-1$ other particles. Molecular Dynamics therefore requires a total of $N(N-1) \sim N^2$ force calculations. A simulation time that scales with N^2 rapidly becomes a problem when attempting to simulate macroscopic systems. Fortunately, approximations may be made which reduce the computational expense of integrating the equations of motion. The first approximation is to introduce a cut-off atomic distance, beyond which the forces are not calculated. This method works well for van der Waals forces,

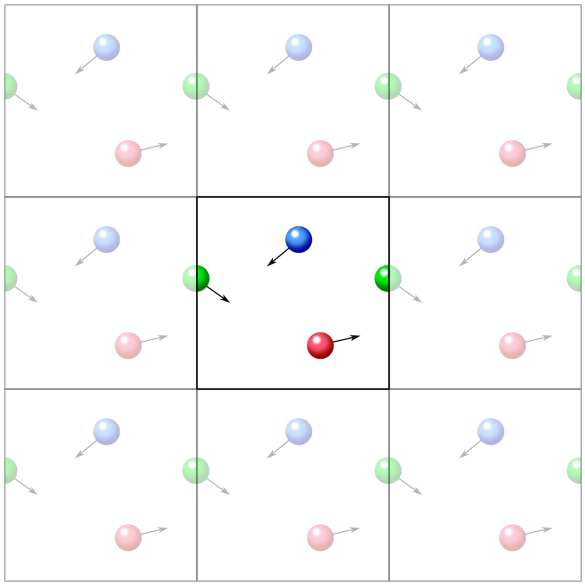

Figure 3.6: Schematic detailing a two-dimensional simulation box with periodic boundary conditions. The simulation cell is repeated infinitely in all dimensions. The particles within the cell interact with all other particles within the cell, as well as all periodic images. If a particle displaces out of the simulation cell, it reappears at the other side.

which scale as $1/r^6$ and therefore decay quickly. With this approximation, the number of forces calculated scales with N, rather than N^2.

The second approximation is to assume that macroscopic systems such as bulk crystalline materials and surfaces can be approximated as an infinite repetition of a single cell. This can be introduced to molecular simulation through the introduction of periodic boundaries. As detailed in Figure 3.6, a particle exiting the simulation cell under periodic boundary conditions will reappear at the opposite side of the cell. Analogous to equations 3.46 and 3.45, the total potential energy of the simulation therefore becomes:

$$\begin{aligned} V(R_i) = &\sum_{C_1} \sum_{i,j} V_2^{ij}(R_i, R_j + C_1) \\ &+ \sum_{C_1,C_2} \sum_{i,j,k} V_3^{ijk}(R_i, R_j + C_1, R_k + C_2) \\ &+ \sum_{C_1,C_2,C_3} \sum_{i,j,k,l} V_4^{ijkl}(R_i, R_j + C_1, R_k + C_2, R_l + C_3) \\ &+ \frac{1}{4\pi\epsilon_0} \sum_{i<j} \frac{q_i q_j}{|R_i - R_j - C_1|} \\ &+ \frac{1}{4\pi\epsilon_0} \sum_{C_1!=0} \sum_{i,j} \frac{q_i q_j}{|R_i - R_j - C_1|} \end{aligned} \tag{3.58}$$

3.2.7 Long-range electrostatics

The first three terms of equation 3.58 may be easily calculated across periodic boundaries due to the cut-off distance. The cut-off may be implemented because the interaction with the shortest range are the van der Waals forces, which decay with $1/r^6$. However, the electrostatic energy term decays as $1/r$, and thus requires a far larger cut-off distance. Instead, the standard method for handling these long-range electrostatics over periodic boundaries is the Ewald Summation Method [100]. Consider the charge density as a function of distance from atom i, $\rho_i(r)$. Under periodic boundary conditions, this is given by:

$$\rho_i(r) = \sum_n \sum_j{}' q_i q_j \delta(r - R_i - n) \tag{3.59}$$

where δ is the Dirac delta function and the prime denotes a sum over j for $n \neq 0$ and a sum over $j \neq i$ when $n = 0$. The potential field generated by this charge distribution about atom i, $\phi_i(r)$, is given by the solution to Poisson's equation,

$$\nabla^2(\phi_i(r)) = -\frac{\rho_i(r)}{\epsilon_0} \tag{3.60}$$

which becomes:

$$\phi_i(r) = \frac{1}{4\pi\epsilon_0} \int \frac{\rho(r')}{|r - r'|} \mathrm{d}^3 r' \tag{3.61}$$

When equation 3.59 is applied to equation 3.61 and summed over i, the final two terms of equation 3.58 are reproduced. However, the Ewald Summation Method involves splitting ρ_i into two functions as follows:

$$\rho_i = \rho_i^S + \rho_i^L \tag{3.62}$$

where:

$$\begin{aligned} \rho_i^S(r) &= \sum_n \sum_j{}' q_i q_j \delta(r - R_i - n) - \sum_n \sum_j{}' q_i q_j G(r - R_i - n) \\ \rho_i^L(r) &= \sum_n \sum_j{}' q_i q_j G(r - R_i - n) \end{aligned} \tag{3.63}$$

where G is a normalised Gaussian function given by:

$$G(r) = \left(\frac{\alpha}{\pi}\right)^{\frac{3}{2}} \exp\left(-\alpha r^2\right) \tag{3.64}$$

Here, α is a parametrisable standard deviation. It should be noted that $\lim_{\alpha\to 0} G_\alpha(r) = \delta(r)$. Using equation 3.61 $\phi_i(r)$ becomes:

$$\phi_i(r) = \frac{1}{4\pi\epsilon_0}\int \frac{\sum_n \sum_j{}' q_i q_j \delta(r' - R_i - n) - \sum_n \sum_j{}' q_i q_j G(r' - R_i - n)}{|r - r'|}\mathrm{d}^3 r' + \frac{1}{4\pi\epsilon_0}\int \frac{\sum_n \sum_j{}' q_i q_j G(r' - R_i - n)}{|r - r'|}\mathrm{d}^3 r' \tag{3.65}$$

It is worth noting here that the first of these two terms will become screened at large distances, and can be solved through application of Poisson's equation. The second term is long range. However, this issue can be resolved by applying Poisson's equation in Fourier space. Doing so removes the requirement to sum over periodic images. Equation 3.65 therefore reduces to:

$$\phi_i(r) = \frac{1}{4\pi\epsilon_0}\sum_n \sum_j{}' \frac{q_i q_j}{|r - R_j - n|}\mathrm{erfc}(|r - R_j - n|\sqrt{\alpha}) + \frac{1}{V\epsilon_0}\sum_{k\neq 0}\sum_j \frac{q_i q_j}{k^2}\exp(ik.(r - R_j))\exp\left(-\frac{k^2}{4\alpha}\right) - \frac{1}{4\pi\epsilon_0}\sqrt{\frac{\alpha}{\pi}}q_i^2 \tag{3.66}$$

where V is the volume of the simulation cell, and $\mathrm{erfc}(x)$ is an extension of the error function,

$$\mathrm{erfc}(x) = 1 - \frac{2}{\pi}\int_0^x e^{-y^2}\mathrm{d}y \tag{3.67}$$

which decays to zero as $x \to \infty$. The first term in equation 3.66 is short-range, due to the erfc convergence. A cut-off distance term may therefore be imposed on this term. The second term is also short-range in k-space due to the exponent, and may therefore be assigned a cut-off distance in k-space. Throughout this thesis, a method known as the Particle-Particle-Particle Mesh (PPPM) method [101] is applied to calculate the second term of equation 3.66. The PPPM method divides the simulation cell into a discrete grid in which the charge density is calculated via discrete Fourier transforms. The fast Fourier Transform (FFT) method allows calculations in k-space which scales with N. It is worth noting, at this point, that the PPPM method, in addition to introducing cut-off distances on short-range forces, allows the simulation time to scale with N, rather than N^2, allowing for a vastly improved computational efficiency.

3.2.8 Integrating the equations of motion

The previous sections deal with calculating the system potential energy. Once both the particle velocities and potential energies are known, they may be

used to find an approximate solution to the equations of motion. As previously stated, the purpose of molecular dynamics is to calculate ensemble averages by integrating the equations of motion. The most straightforward means of doing so is Euler's approach, which is derived through a straightforward Taylor expansion,

$$R_i(t+\Delta t) = R_i(t) + \frac{\mathrm{d}R_i}{\mathrm{d}t}\Delta t + \frac{F_i}{m}\Delta t^2 + \mathcal{O}\Delta t^3 \tag{3.68}$$

where Δt is the simulation timestep and F_i is the force acting on atom i given by the gradient of the potential energy:

$$F_i = \nabla - \mathcal{V}(R_i(t)) \tag{3.69}$$

However, Euler's approach becomes unsuitable for molecular dynamics. Not only is the error of the calculation high, but it lacks the time-reversal symmetry of Newton's equation. Instead, a more appropriate approach is to use the Verlet integration scheme, which can be derived as follows:

$$\begin{aligned}
R_i(t+\Delta t) &= R_i(t) + \frac{\mathrm{d}R_i}{\mathrm{d}t}\Delta t + \frac{F_i}{2m}\Delta t^2 + \frac{1}{3!}\frac{\mathrm{d}^3R_i}{\mathrm{d}t^3}\Delta t^3 + \mathcal{O}(\Delta t^4) \\
R_i(t-\Delta t) &= R_i(t) - \frac{\mathrm{d}R_i}{\mathrm{d}t}\Delta t + \frac{F_i}{2m}\Delta t^2 - \frac{1}{3!}\frac{\mathrm{d}^3R_i}{\mathrm{d}t^3}\Delta t^3 + \mathcal{O}(\Delta t^4) \\
R_i(t+\Delta t) &= R_i(t+\Delta t) + R_i(t-\Delta t) - R_i(t-\Delta t) \\
&= 2R_i(t) - R_i(t-\Delta t) + \frac{F_i}{m}\Delta t^2 + \mathcal{O}(\Delta t^4)
\end{aligned} \tag{3.70}$$

Note that $R_i(t-\Delta t)$ can be stored with minimal computational expense. The Verlet algorithm is advantageous over the Euler method due to its $\mathcal{O}(\Delta t^4)$ error and its time-reversal symmetry. The latter advantage also renders the Verlet algorithm advantageous over more sophisticated integration schemes such as the 4^{th}-order Runge Kutta method. However, one disadvantage of the Verlet scheme is that it cannot be initialised with a starting configuration, as a previous configuration is also required. Another issue with the Verlet scheme is that the instantaneous definition of the velocity,

$$V_i(t) = \frac{R_i(t+\Delta t) - R_i(t-\Delta t)}{2\Delta t} + \mathcal{O}(\Delta t^2) \tag{3.71}$$

requires $R_i(t+\Delta t)$ to be calculated. This presents a difficulty when modifying the system velocities, something which will be further discussed in Section 3.2.9. An alternative is the Verlet algorithm is the Velocity Verlet algorithm, which is implemented in three stages.

1. Integrate the velocities by half a timestep:

$$V_i\left(t+\frac{\Delta t}{2}\right) = V_i(t) + \frac{F_i(t)}{2m}\Delta t \tag{3.72}$$

2. Update positions by one timestep:

$$R_i(t+\Delta t) = R_i(t) + V_i\left(t+\frac{\Delta t}{2}\right)\Delta t \tag{3.73}$$

3. update velocities by one timestep:

$$V_i(t+\Delta t) = V_i\left(t+\frac{\Delta t}{2}\right) + \frac{F_i(t)}{2m}\Delta t^2 \tag{3.74}$$

The Velocity Verlet algorithm can be shown to reproduce a trajectory identical to that produced by the Verlet algorithm. This can be demonstrated as follows. Equations 3.72 and 3.73 give:

$$\begin{aligned} R_i(t+\Delta t) &= R_i(t) + V_i(t)\Delta t + \frac{F_i(t)}{2m}\Delta t^2 \\ R_i(t-\Delta t) &= R_i(t) - V_i(t)\Delta t + \frac{F_i(t)}{2m}\Delta t^2 \end{aligned} \tag{3.75}$$

Adding the two above equations reproduces equation 3.70. This alone does not prove the trajectories are identical between Verlet and Velocity Verlet algorithms. To show this, $R_i(t+2\Delta t)$ must be calculated. Using equations 3.72, 3.73 and 3.74, this can be done as follows:

$$\begin{aligned} R_i(t+2\Delta t) &= R_i(t+\Delta t) + V_i\left(t+\frac{3\Delta t}{2}\right) \\ &= R_i(t+\Delta t) + V_i(t+\Delta t)\Delta t + \frac{F_i(t+\Delta t)}{2m}\Delta t^2 \\ &= R_i(t+\Delta t) + V_i\left(t+\frac{\Delta t}{2}\right)\Delta t + \frac{F_i(t+\Delta t)}{m}\Delta t^2 \\ &= R_i(t+\Delta t) + V_i(t)\Delta t + \frac{F_i(t)}{2m}\Delta t^2 + \frac{F_i(t+\Delta t)}{m}\Delta t^2 \end{aligned} \tag{3.76}$$

Equation 3.75, can be used to substitute for $R_i(t)$ to give:

$$R_i(t+2\Delta t) = 2R_i(t+\Delta t) - R_i(t) + \frac{F_i(t+\Delta t)}{m}\Delta t^2 \tag{3.77}$$

This demonstrates that the Velocity Verlet algorithm reproduces the exact trajectory of the standard Verlet algorithm. The Velocity Verlet algorithm is advantageous over the Verlet algorithm as it allows the initialisation of a simulation without reference to a past configuration. It also allows for the external modification of the velocities. Throughout this thesis, the Velocity Verlet algorithm is used to integrate the equations of motion.

3.2.9 Thermostatting and barostatting

Solving the equations of motion for a simulation cell will fundamentally find a solution for a microcanonical system, as the total system energy is not modified by any external input. Microcanonical systems are of very little interest, as few physical systems of interest are isolated. Instead, we wish to simulate a canonical system. In order to do so, the system energy must be allowed to fluctuate about an average value such that it has a defined system temperature.

In Section 3.1.8, it was shown that each kinetic energy mode has an average value equal to $k_B T/2$. The average total kinetic energy of the system, $\langle K \rangle$ is therefore given by:

$$\langle K \rangle \equiv \sum_{i=1}^{3N} \left\langle \frac{P_i^2}{2m_i} \right\rangle = \frac{3}{2} N k_B T \tag{3.78}$$

this provides an instantaneous definition of system temperature from the easily calculated kinetic energies of the system. The kinetic energies of the system should therefore fluctuate around a total value given by equation 3.78. The most straightforward means of implementing this is to periodically rescale the system velocities so that the instantaneous system temperature becomes equal to the input temperature. However, this approach fails to replicate the natural temperature fluctuations of the system. It also produces a curious effect known as the *Flying Ice Cube Effect*, in which the equipartition theorem is violated and the system energy is transferred from high-frequency modes to low-frequency modes [102].

A more suitable means of thermostatting is the Nosé-Hoover thermostat, [103] in which an additional degree of freedom is added to the system. This additional degree of freedom manifests as a frictional force:

$$m_i \frac{\mathrm{d}^2 r_i}{\mathrm{d}t^2} = F_i - \zeta(t) m_i v_i \tag{3.79}$$

where the friction coefficient $\zeta(t)$ has a time dependence given by:

$$\frac{\mathrm{d}\zeta(t)}{\mathrm{d}t} = \frac{1}{Q} \left[\sum_{i=1}^{N} \frac{1}{2} m_i v_i^2 - \frac{1}{2}(3N+1) k_B T \right] \tag{3.80}$$

where Q is the effective mass associated with the additional degree of freedom, which determines the speed at which the system temperature relaxes back to the equilibrium temperature. It is worth noting that $3N$ becomes $3N+1$ in equation 3.81 due to the additional degree of freedom that the frictional component brings. The Nosé-Hoover thermostat is deterministic, and is found to allow the system to sample a canonical distribution.

In order to understand the physical representation of Q, equation 3.81 can be recast as follows:

$$\frac{\mathrm{d}\zeta(t)}{\mathrm{d}t} = \frac{1}{\tau_{\mathrm{damp}}}\left[\frac{\sum_{i=1}^{N}\frac{1}{2}m_i v_i^2}{(3N+1)k_B T} - 1\right] \tag{3.81}$$

where:

$$\tau_{\mathrm{damp}} = \frac{2Q}{(3N+1)k_B T} \tag{3.82}$$

τ_{damp} is interpreted as a relaxation time-scale over which the Nosé-Hoover thermostat acts. If τ_{damp} is assigned a value too high, a large simulation time will be required to calculate ensemble averages. If, on the other hand, τ_{damp} is assigned a value too low, the temperature of the simulation will fluctuate wildly.

As demonstrated in Section 3.1.9, constant pressure conditions produce a canonical-like distribution of microstates. In order to replicate this, a modification of the Nosé-Hoover algorithm is implemented to include a barostat along with a thermostat [104]. The barostat behaves similarly to the thermostat, by including an additional degree of freedom and acts as a frictional force to the system volume. Similar to the Nosé-Hoover thermostat, the barostat has an associated τ_{damp} which must be assigned a value with care.

3.3 Rare event sampling

As discussed throughout Section 3.2, Molecular Dynamics is a powerful tool for calculating ensemble averages by exploiting ergodicity. Free Energy Surfaces are an example of such ensemble averages. Free energy Surfaces are useful for determining reaction pathways and have a large scope for calculating binding free energies and reaction rates. However, simulations are limited by computational expense, and ergodicity is often hindered by the timescales accessible for a molecular dynamics simulation. Rare event sampling is a tool designed to overcome such hindrances through exerting external, non-physical dynamics on the system.

3.3.1 Bias potentials

By running classical molecular dynamics simulations, free energy surfaces can be determined for any reaction coordinate, $s(\Gamma)$, by constructing a histogram for the collective variable and using equation 3.40. However, when two or more regions are separated by a region of high free energy, $F(s(\Gamma))$ (low $\rho(s(\Gamma))$), then ergodicity is hindered by the low crossing rate of such a region. Running extremely long simulations to account for this is undesirable.

Instead, an artificial bias potential as a function of the reaction coordinate, $V(s)$ can be introduced to sample regions of space which are rarely occupied otherwise. $V(s)$ is not to be confused with the system potential energy $\mathcal{V}$. Such a bias potential shifts the total system Hamiltonian to $\hat{\mathcal{H}}_{tot}(\Gamma)$, which is given by $\hat{\mathcal{H}}_{tot}(\Gamma) = \hat{\mathcal{H}}_0(\Gamma) + V(s(\Gamma))$, where $\hat{\mathcal{H}}_0(\Gamma)$ is the system energy without the bias potential. Defining $\mathcal{Z}_{\text{tot}}$ and $\mathcal{Z}_0$ as the corresponding canonical partition functions, the following relations hold:

$$\mathcal{Z}_{\text{tot}} = \int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma) + V(s(\Gamma))}{k_B T}\right) \mathrm{d}\Gamma \qquad \mathcal{Z}_0 = \int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right) \mathrm{d}\Gamma \tag{3.83}$$

from which it follows that:

$$\begin{aligned} \mathcal{Z}_{tot} &= \mathcal{Z}_0 \frac{\mathcal{Z}_{tot}}{\mathcal{Z}_0} \\ &= \mathcal{Z}_0 \int \left(\frac{\exp\left(-\frac{\hat{\mathcal{H}}(\Gamma)}{k_B T}\right)}{\int \exp\left(-\frac{\hat{\mathcal{H}}(\Gamma')}{k_B T}\right) \mathrm{d}\Gamma'} \exp\left(-\frac{V(s(\Gamma))}{k_B T}\right) \right) \mathrm{d}\Gamma \\ &= \mathcal{Z}_0 \left\langle \exp\left(-\frac{V(s(\Gamma))}{k_B T}\right) \right\rangle \end{aligned} \tag{3.84}$$

Now, two separate probability densities can be defined; $\rho_b(s)$ represents the PDF for the biased system with $s(\Gamma) = s$, where $\rho_0(s)$ gives the corresponding PDF in the absence of the bias potential. Repeating the derivation of equation 3.41 but with a bias potential, the probability density is now given by:

$$\rho_b(s) = \frac{1}{\mathcal{Z}_{tot}} \int \delta(s - s(\Gamma)) \exp\left(-\frac{\hat{\mathcal{H}}_0(\Gamma)}{k_B T}\right) \exp\left(-\frac{V(s(\Gamma))}{k_B T}\right) \mathrm{d}\Gamma \tag{3.85}$$

Substitution of equation 3.84 into this expression results in equation 3.86.

$$\begin{aligned} \rho_b(s) &= \frac{1}{\mathcal{Z}_0 \langle \exp\left(-\frac{V(s(\Gamma))}{k_B T}\right) \rangle} \int \delta(s - s(\Gamma)) \exp\left(-\frac{\hat{\mathcal{H}}_0(\Gamma)}{k_B T}\right) \exp\left(-\frac{V(s(\Gamma))}{k_B T}\right) \mathrm{d}\Gamma \\ &= \left\langle \exp\left(-\frac{V(s(\Gamma))}{k_B T}\right) \right\rangle^{-1} \exp\left(-\frac{V(s)}{k_B T}\right) \rho_0(s) \end{aligned} \tag{3.86}$$

Note that $\rho_0(s)$ is related to the free energy of the system through equation 3.41. Therefore, taking the logarithm of both sides of equation 3.86 and multiplying by k_BT yields equation 3.87:

$$\begin{aligned} F(s) = &- k_BT \log\{\rho_b(s)\} - V(s) \\ &- k_BT \log \left\langle \exp\left(-\frac{V(s(\Gamma))}{k_BT}\right)\right\rangle - k_BT \log(\mathcal{Z}_0) \end{aligned} \tag{3.87}$$

The final two terms in equation 3.87 are merely constants of the system, and can be considered unimportant. Therefore, the free energy profile is given in equation 3.88:

$$F(s) = -k_BT \log\{\rho_b(s)\} - V(s) + C \tag{3.88}$$

This result demonstrates that the inclusion of a non-physical bias potential still allows a calculation of the free energy surface. Equation 3.88 provides a useful tool for recovering free energy surfaces using an external potential. A bias potential may be used to force the reaction coordinate into otherwise under-sampled regions, thus allowing ergodicity within accessible simulation timescales.

3.3.2 Multiple reaction coordinates

Before proceeding, it is important to provide a generalisation for a multidimensional bias that depends on multiple reactions coordinates. Consider a set of d collective variables. Equation 3.88 in this case becomes:

$$F(s_1, ..., s_d) = -k_BT \log \rho(s_1, s_2, ..., s_d) - V(s_1, s_2, ..., s_d) + C \tag{3.89}$$

Given a probability density of multiple variables, $\rho(s_1, s_2, ..., s_d)$, the reduced probability density for only one variable, s_1, is given by:

$$\rho(s_1) = \int \rho(s_1, s_2..., s_d) \prod_{i=2}^{d} \mathrm{d}s_i \tag{3.90}$$

Assuming the d-dimensional free energy surface is already known, equation 3.41 can be used to simplify equation 3.90 to equation 3.91.

$$F(s_1) = -k_BT \log \left(\exp\left(-\frac{F(s_1, s_2, ..., s_d)}{k_BT}\right) \prod_{i=2}^{d} ds_i \right) + C \tag{3.91}$$

Therefore, assuming the free energy surface for multiple collective variables is known, it is straightforward to derive the free energy surface for one collective variable.

3.3.3 Thermodynamic integration

Thermodynamic integration is arguably the most straightforward rare event sampling technique. Equation 3.40 tells us that the integral of the mean force will recover the Free Energy Surface. By constraining a reaction coordinate to a given value, a series of positions along the reaction coordinate s_i can be constructed. By calculating the mean force along s at each s_i, $\langle f(s_i)\rangle$, the Free Energy Surface can be estimated.

$$F(s_{i+1}) - F(s_i) \approx \langle f(s_i)\rangle(s_{i+1} - s_i) \tag{3.92}$$

3.3.4 Umbrella sampling

Umbrella sampling is a similar technique to thermodynamic integration, in which a harmonic bias potential, $V(s, i)$, is used to constrain a reaction coordinate to construct a series of windows.

$$V(s, i) = \frac{1}{2}k(s - s_i)^2 \tag{3.93}$$

With this potential, a probability histogram, $\rho_i(s)$, can be constructed. The local free energy surface can be recovered from the histogram:

$$F(s \approx s_i) = -k_B T ln\{\rho_i(s)\} - k(s - s_i)^2 + C \tag{3.94}$$

Using equation 3.94, a series of local free energy surfaces are constructed for each window. Assuming sufficient overlap these local free energy functions can be pieced together to create a free energy surface. This process of combining local free energy functions is complex. A popular method for doing so is the Weighted Histogram Analysis Method (WHAM) [105]. Where umbrella sampling is used throughout this thesis, the WHAM method implemented by Grossfield [106] is used.

As well as being more efficient than Thermodynamic Integration, Umbrella sampling has the advantage of requiring few parameters: namely, the spring constant, k, and the distance between windows, $s_{i+1} - s_i$. These parameters must be adjusted to give sufficient overlap between histograms such that the WHAM method is implemented with accuracy. A disadvantage with umbrella sampling is its reduced ability to determine reaction pathways, due to the constraints at each window.

One additional use of umbrella sampling is its ability to calculate a sampling error, using a method known as *Monte Carlo bootstrapping*. For a data set of N points along the reaction coordinate, Monte Carlo bootstrapping selects N data points randomly from the data set, allowing for repetitions. Subsequently, WHAM is used to calculate a free energy surface from the data subset. This process is repeated a number of times with a different random

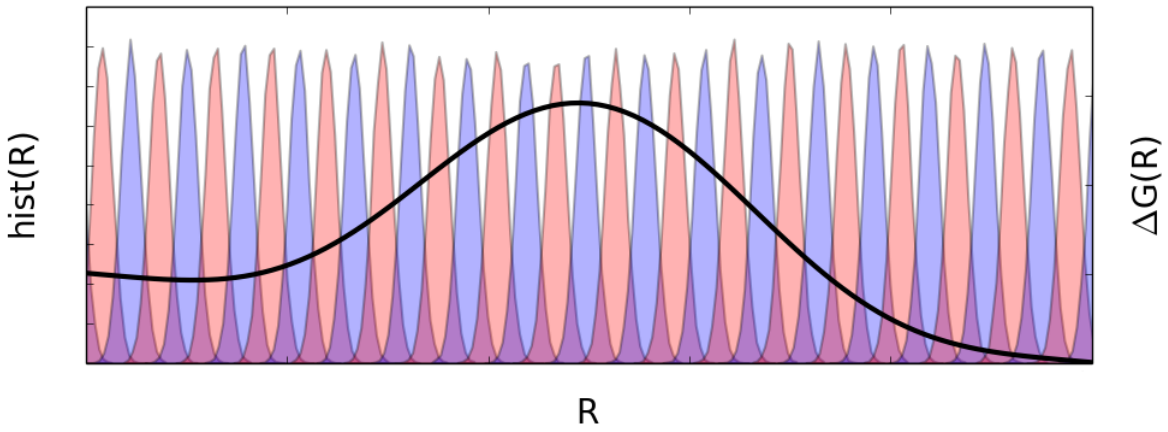

Figure 3.7: Depiction of local histograms (umbrellas) obtained by constraining at different regions along the reaction coordinate. Provided sufficient overlap between histograms, the free energy surface can be obtained using the Weighted Histogram Analysis Method.

number seed, and the free energy surfaces are used to obtain a sampling error.

3.3.5 Metadynamics

As well as having a spatial dependence, a bias potential may also be time-dependent, $V(s,t)$. In metadynamics, a time dependent bias potential is implemented through the periodic creation of Gaussian functions. Gaussian potentials of standard deviation σ and height W are deposited in increments of time τ.

$$V(s,t) = \sum_{k\tau<t} W \exp\left(-\frac{(s-s_{k\tau})^2}{2\sigma^2}\right) \tag{3.95}$$

where k is an integer and $s_{k\tau}$ is the value of s at time $k\tau$. Over time, the Gaussian functions accumulate, leaving a bias potential dependent on the history of the system. Should the system become kinetically trapped in a free energy well, the bias will accumulate until the system is forced out of the well. This allows for improved ergodicity of the system, which is otherwise hindered by the free energy landscape.

In theory, the bias will continue to accumulate until the total forces along the reaction coordinate are zero, and the Probability Density Function becomes constant. Under this limit, equation 3.40 tells us that:

$$F(s) = -\lim_{t\to\infty} V(s,t) \tag{3.96}$$

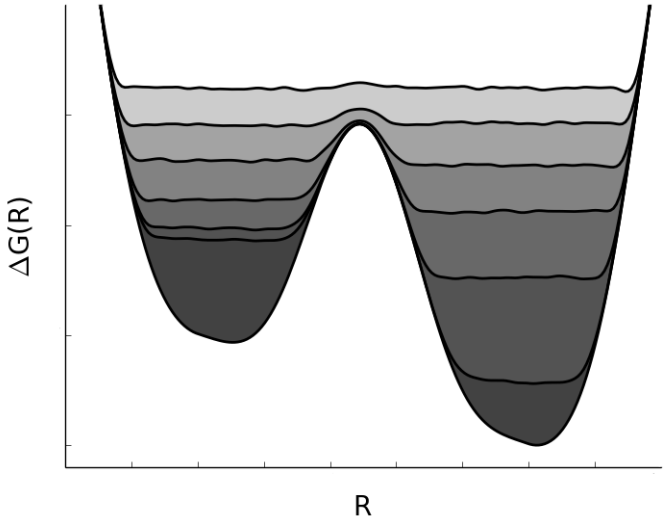

Figure 3.8: Schematic of the derivation of a free energy landscape using metadynamics. The true (unknown) free energy map is denoted by the black line. The Gaussians accumulate around the left free energy well, forcing the collective variable into the right well. The second well is subsequently filled with Gaussians until both wells are accounted for, after which the sum of the Gaussians fluctuate about the true free energy surface modulo a constant.

This process is illustrated in Figure 3.8. Metadynamics has the advantage of being straightforward to implement compared to umbrella sampling, especially when multiple reaction coordinates are used. It also allows the system to follow a semi-physical trajectory, allowing reaction pathways to be examined, unlike with umbrella sampling. However, parametrising W and σ for maximum efficiency requires some prior knowledge on the free energy surface. Metadynamics also tends to converge slower than umbrella sampling.

3.3.6 Well-tempered metadynamics

In practice, classical metadynamics fares poorly when determining the final Free Energy Surface. The Gaussian functions accumulate indefinitely, meaning that the final output always has a resolution equal to W. The resolution may be improved by reducing W, but this comes at the expense of larger convergence times. Well-Tempered Metadynamics is a technique used to overcome this issue through adding a spatial dependence and time-dependence to the Gaussian height, $W(s, k\tau)$. The value of $W(s, k\tau)$ is determined by equation 3.97, where ΔT is a parametrisable variable with units of temperature.

$$W(s, k\tau) = W_0 \exp\left(-\frac{V(s_k, k\tau)}{k_B \Delta T}\right) \tag{3.97}$$

The height of the Gaussian is inversely proportional to the exponential of the height of the bias at time $k\tau$. In reality, $V(s, k\tau)$ changes in a succession of discrete processes, but in the limit of $k \gg 1$ the change in V can be considered a continuous process and a simulation time $t = \kappa\tau$ can be defined. The rate of change in height of the barrier, $\dot{V}(s,t) = dV/dt$, is given by $W(s,t)$, and the frequency at which Gaussians are deposited is given by $\rho_b(s,t)$. The change in potential is therefore given by:

$$\dot{V} \propto W_0 \exp\left(-\frac{V(s,t)}{k_B \Delta T}\right) \rho_b(s) \tag{3.98}$$

After large simulation times, it can be assumed that the Gaussian heights are small enough such that the potential evolves very slowly in time. At this point, it can be assumed that the potential has converged to a point such that will vary at the same rate regardless of s. Therefore, substituting equation 3.88 for $\rho_b(s)$ into equation 3.98, the following equation holds:

$$\exp\left(-\frac{V(s, t\to\infty)}{k_B \Delta T}\right) \exp\left(\frac{-(F(s) + V(s, t\to\infty))}{k_B T}\right) = C \tag{3.99}$$

from this, it follows that $\frac{-V(s,t\to\infty)}{k_B\Delta T} \frac{-(F(s)+V(s,t\to\infty))}{k_B T} = C'$. This simplifies to equation 3.100.

$$F(s) = -\frac{T + \Delta T}{\Delta T} V(s, t\to\infty) + C'' \tag{3.100}$$

Consider the limits of ΔT. As $\Delta T \to 0$, the bias potential vanishes, and the probability density reduces to the PDF in the absence of a bias potential. As $\Delta T \to \infty$, equation 3.100 becomes $F(s) = V(s, t\to\infty) + C'''$; namely, the classical metadynamics algorithm is reproduced. It is conventional to define the bias factor, γ, as:

$$\gamma = \frac{T + \Delta T}{T} \tag{3.101}$$

Equation 3.100 can now be rearranged to give $F(s) + V(s, t\to\infty) = \frac{1}{\gamma} F(s)$. Finally, dividing both sides of this expression by $k_B T$, taking the logarithm and subbing equation 3.88 results in equation 3.102.

$$\rho(s, t\to\infty) \propto \exp\left(-\frac{F(s)}{k_B T} \frac{1}{\gamma}\right) \tag{3.102}$$

Note that the limits of the bias factor are 1 and ∞.

It is clear, therefore, that the choice of bias factor is important: too low and it will not allow sufficient exploration of collective variable space; too high and the Gaussian height will decrease too slowly and convergence of the

bias potential will require longer simulations. The bias factor must therefore be parametrised to allow both an effective bias potential and a reasonable final resolution of the calculated $F(s)$. An ideal choice of bias factor is obtained by dividing the largest barrier height, $\Delta F_{\max}$ by $k_B T$. Consider the relative maximum and minimum values of the PDF, $\rho_{\max}$ and $\rho_{\min}$. Equation 3.102 can be used to give $\rho_{min}/\rho_{max} = \exp(-(F_{max}-F_{min})/\gamma k_B T)$. If the bias factor is set to $\gamma = (F_{max}-F_{min})/k_B T$, then the ratio of probabilities becomes $\rho_{min}/\rho_{max} = \exp(-1)$. The minimum and maximum values of the PDF are of the same order of magnitude, implying that the barrier crossing rate will be sufficiently frequent, while minimising the simulation time required to generate a sufficiently high-resolution Free Energy Surface. This process can be used to significantly reduce the convergence time, although its validity again requires some prior knowledge of the free energy landscape.

3.3.7 Mean force integration

Mean Force Integration (MFI) is a more recently proposed technique [107] MFI is an extension of metadynamics which uses the frequent output of the collective variables to calculate the mean force, $dF_t(s)/ds$, from the time dependent bias potential and probability density $P_t(s)$ using the relationship:

$$\left\langle \frac{dF_t(s)}{ds} \right\rangle_t = \frac{\sum_{t'=0}^{t} P_{t'}(s) \frac{dF_{t'}(s)}{ds}}{\sum_{t'=0}^{t} P_{t'}(s)} \tag{3.103}$$

Combining equation 3.103 with equation 3.88 gives:

$$\left\langle \frac{dF(s)}{ds} \right\rangle_t = \frac{1}{\sum_k P_{k\tau}(s)} \left[-k_B T \sum_k \frac{dP_{k\tau}(s)}{ds} - \sum_k P_{k\tau}(s) \frac{dV_{k\tau}(s)}{ds} \right] \tag{3.104}$$

where k is an integer and τ is a specified time period. Equation 3.104 can therefore be used to calculate $F(s)$ by integrating. $P_t(s)$ can be calculated from the output of the value of the collective variables during simulations. However, in order for $P_t(s)$ to be differentiable, a simple histogram will not suffice. Instead, $P_t(s)$ is calculated using a kernel approach, which gives a continuous, differentiable function.

$$P_t(s) = \frac{1}{n_\tau h \sqrt{2\pi}} \sum_{t'=t}^{t+\tau} \exp\left[-\frac{(s - s_{t'})^2}{2h^2} \right] \tag{3.105}$$

where h is a tunable kernel width and n_τ is the number of kernels deposited per time tau. The main advantage of MFI is that it allows for faster convergence than well-tempered metadynamics. An additional advantage of MFI is that its standard errors can be calculated by block averaging $dF_t(s)/ds$ and propagating the errors when calculating $F(s)$. It is also worth noting

that MFI does not need to be implemented on the fly, and only requires post-processing the output of metadynamics simulations as long as the value of the collective variables is also outputted.

3.3.8 The Jarzynski approach

The Jarzynski equality [108, 109] relates an equilibrium free energy difference between two states, ΔF, to the output of non-equilibrium simulations as follows

$$\exp\left(-\frac{\Delta F}{k_B T}\right) = \left\langle \exp\left(-\frac{\Delta W}{kBT}\right)\right\rangle \tag{3.106}$$

Where ΔW is the work required in transitioning between the two states. This can be extended to calculating the free energy difference along a reaction coordinate as follows:

$$\exp\left(-\frac{\Delta F(x_1 \to x_2)}{k_B T}\right) = \left\langle \exp\left(-\frac{\Delta W(x_1 \to x_2)}{k_B T}\right)\right\rangle \tag{3.107}$$

where $\Delta W(x_1 \to x_2)$ is the work done in moving along the reaction coordinate from x_1 to x_2. This equation holds regardless of the time taken for the transition from x_1 to x_2. The Jarzynski approach is unique in the sense that it uses nonequilibrium simulations to recover equilibrium free energy differences. In practice, this method can be implemented by tethering a collective variable to a certain value using a moving harmonic potential. An advantage of the Jarzynski approach is that it can be applied to processes which are not easily reversible in simulations, for example the disassembly of molecular clusters [110]. Another advantage is that, since the free energy function is just an average of many repetitions of the same simulation, errors are easily calculated.

3.3.9 Minimum Free Energy Pathways

In practice, physical processes will often require multiple reaction coordinates in order to correctly sample the phase space of the system with rare event sampling. In this case, reactant and product states are separated by a multi-dimensional free energy barrier. Transitions from reactant to product states therefore involves the translation along many possible pathways over the multi-dimensional free energy surface. However, it has been shown [111] that the reaction pathways are dominated by a single trajectory, one which minimises the free energy along its path. The Pathway is therefore known as the Minimum Free Energy Pathway (MFEP). Here, we denote ξ as the reaction coordinate that produces the MFEP. The MFEP is therefore given

by $F(\xi)$, i.e. a function of ξ. An important property of the MFEP is that mean force normal to the pathway, $\langle f_\perp(\xi)\rangle$ is always zero,

$$\langle f_\perp(\xi)\rangle = -\nabla_\xi F(\xi) + (\nabla_\xi F(\xi).t(\xi))t(\xi) = 0 \tag{3.108}$$

where $t(\xi)$ is the normalised tangent vector to the MFEP.

3.3.10 String method

Calculating the MFEP can be achieved through the use of equation 3.108. With the knowledge that normal forces are always zero along the MFEP, the mean forces can be used to parametrise the MFEP. The String Method [112] involves a linear series of equidistant points, λ_i, along a set of reaction coordinates, known as *images*. This collection of images forms a *string*, hence the name String Method. Starting with an initial guess of the MFEP, the mean forces can be calculated and used to reposition each image using the steepest descent algorithm:

$$\lambda_i \rightarrow \lambda_i - h\nabla_\lambda F(\lambda_i) \tag{3.109}$$

where h is a small step which controls the sensitivity of the iteration with respect to the forces. It should be noted that h should be sufficiently small to allow convergence, similar to how a simulation time step must be small. After each point has been repositioned, the points are re-parametrised such that they are equidistant along the MFEP estimation. After many iterations of this process, the string will converge on the MFEP. Figure 3.9 demonstrates the process of an initial guess converging to the MFEP after a finite number of iterations. It is worth noting that the string method can be applied either in a simulation, where mean forces are calculated on the fly, or to a pre-existing multidimensional free energy surface.

3.3.11 Nudged elastic band method

The Nudged Elastic Band Method, [113] is an extension of the string method in which the images are connected by a series of harmonic springs with spring constants k_i. Uniformity between spacings is maintained by neglecting the system mean forces tangent to the springs, and neglecting the spring forces normal to the springs. The total force on image i, $f(\lambda_i)$ therefore becomes:

$$\begin{aligned} f(\lambda_i) = \left\langle f_\parallel(\lambda_i)\right\rangle + f_{\perp,i} = &-\nabla_\lambda F(\lambda_i) + ((\nabla_\lambda F(\lambda_i)).t_i)t_i \\ &+ ([k_{i+1}(\lambda_{i+1}-\lambda_i) - k_i(\lambda_i - \lambda_{i-1})].t_i)t_i \end{aligned} \tag{3.110}$$

where t_i is the tangent vector at λ_i. The string method is therefore the non-extensible limit of the Nudged elastic band method. In this thesis, we use the Nudged elastic Band method, due to its relative ease of implementation.

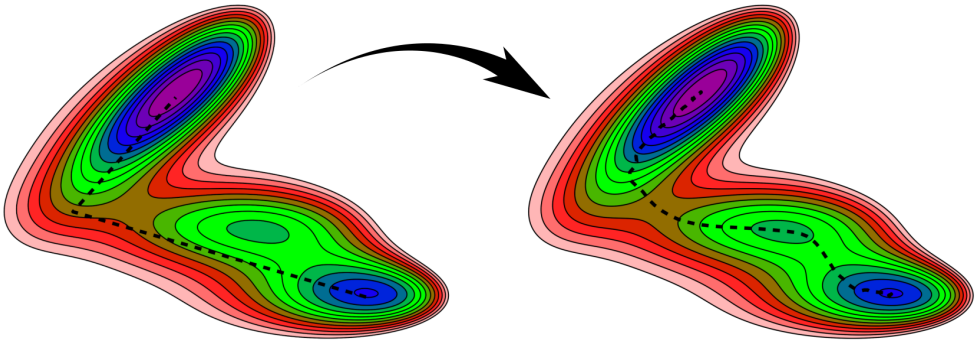

Figure 3.9: Demonstration of the String Method or Nudged Elastic Band method applied to the Müller-Brown potential [2]. In initial guess at the MFEP (dashed line in left-hand image) is made and, after a certain number of iterations, the string converges on the MFEP (right-hand image).

3.3.12 Selecting reaction coordinates

One crucial caveat of rare event sampling techniques is that the set of reaction coordinates must be chosen such that the bias potential allows the rapid exploration of all relevant regions of phase space. In other words, there must be no further rare events hindering ergodicity other than those which can be accurately represented by the reaction coordinates. Any free energy barriers hindering ergodicity that are perpendicular to all reaction coordinates will cause any free energy sampling method to reproduce an inaccurate free energy surface.

To date, there are no rigorous methods available for selecting reaction coordinates. Instead, we are typically required to use intuition in selecting reaction coordinates. An initial guess involving multiple reaction coordinates is often required. However, convergence times approximately scale exponentially with the number of reaction coordinates used, and any rare event sampling method involving more than three reaction coordinates will typically converge far too slowly to be of use. It is therefore highly beneficial to limit the dimensionality of the reaction coordinate space as much as possible. A reaction coordinate can be removed if, for example, no significant free energy barriers exist along it, and diffusion along this reaction coordinate is rapid. Alternatively, the Committor Function [114, 115] can be used to test the suitability of the reaction coordinate.

3.4 Dynamics of stochastic systems

It is often beneficial to neglect certain physical characteristics of a system for the purpose of computational efficiency. Coarse-graining is a method in which the degrees of freedom in a simulation are reduced for greater versatility of the simulation. Molecular Dynamics is itself a method of coarse-graining, in which the electronic structure of each atom is replaced by a set of force fields. Here, we explore a form of coarse-graining which represents the diffusion along a reaction coordinate as a particle moving under the influence of a frictional coefficient and a stochastic noise. We also explore a method of coarse-graining in which a system is represented by a series of discrete random processes, each with their own reaction rates.

3.4.1 Markovianity

A Markovian system is one in which the system dynamics are only determined by the instantaneous conditions of the system. In other words, the system is memoryless and does not require knowledge of previous conditions. Markovianity is an important concept when coarse-graining. Fundamentally, Newton's equations are Markovian, as they only require a set of instantaneous positions, velocities and forces that depend on their coordinates. Markovianity is only broken for physical systems when coarse-graining neglects an important feature of the system. Consider the set-up in Figure 3.10. Both particles connected by a spring move under Newton's equations, and their individual motions are Markovian if the positions and momenta of both particles are tracked. However, if the position and momentum of one of the particles is ignored, and only the other particle is tracked, its motion will become non-Newtonian. Its velocity may change suddenly or speed up regardless of its position or velocity. For example, if the particle is found to suddenly change direction, it is likely that it has displaced some distance from the other particle. Therefore, if the particle has suddenly rapidly decelerated, it is likely to keep doing so and change direction. In other words, the velocity may respond to an instantaneous but neglected coordinate which is correlated with a previous velocity of the visible particle. Hence, non-Markovianity is observed in the apparent dynamics.

It is clear that non-Markovianity can arise through coarse-graining if the details neglected are of significance to the system. This is important, not just for the physical coarse-graining observed in Figure 3.10, but for rare event sampling and determining barrier crossing rates. Rare event sampling is a form of coarse-graining, as it reduces the dynamics of the entire system onto the evolution of a collection of reaction coordinates. If an unsuitable reaction coordinate is chosen, the motion along the reaction coordinate becomes non-Markovian. This detail is explored extensively in Chapter 8.

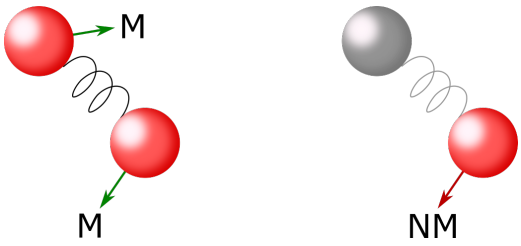

Figure 3.10: Schematic of two moving particles connected by a harmonic spring. The motion of each springs obey Newton's equations, and are therefore Markovian (labelled "M" in the figure). However, when the position and motion of one particle is neglected, represented by the greying of the second particle, the motion of the other particle becomes non-Markovian (labelled "NM" in the figure).

3.4.2 Langevin motion

In its simplest form, Langevin dynamics describe the stochastic motion of a single particle under a frictional force. Designed to describe Brownian motion, the Langevin Equation is a stochastic differential equation for the particle position $x(t)$, defined as follows:

$$\frac{\mathrm{d}^2x}{\mathrm{d}t^2} = -\gamma\frac{\mathrm{d}x}{\mathrm{d}t} + \xi(t) \tag{3.111}$$

where γ represents a damping force and $\xi(t)$ is a stochastic force acting on the the particle. In its simplest form, the force is assumed to be completely uncorrelated with time. This can be represented by considering the autocorrelation function of the force:

$$\langle\xi(t)\xi(t')\rangle = \int\int \xi(t)\xi(t')\rho[\xi(t)]\rho[\xi(t')]\mathrm{d}\xi(t)\mathrm{d}\xi(t') \tag{3.112}$$

where $\rho[\xi(t)]$ is the probability density function of $\xi(t)$. The force can be considered completely uncorrelated if $\langle\xi(t)\xi(t')\rangle$ is zero when $t \neq t'$. In other words, the stochastic force is completely Markovian. This condition can be satisfied if:

$$\langle\xi(t)\xi(t')\rangle = b^2\delta(t-t') \tag{3.113}$$

$\xi(t)$ is therefore a force which is infinite in variance, but changes after an infinitesimal amount of time. While such a force is unintuitive, the particle motion it produces is approximately representative of physical systems. The *fluctuation-dissipation relation* provides a link between the average velocity of the particle and the equipartition theorem, by assuming that the average squared velocity of the particle in equilibrium, $\langle v^2(t\to\infty)\rangle$ is equal to k_BT/m. In this limit, b becomes:

$$b = \left(\frac{2k_B T\gamma}{m}\right)^{\frac{1}{2}} \tag{3.114}$$

The Langevin equation can therefore be recast as:

$$\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} = -\gamma\frac{\mathrm{d}x}{\mathrm{d}t} + \left(\frac{2k_B T\gamma}{m}\right)^{\frac{1}{2}} \eta(t) \tag{3.115}$$

where $\eta(t)$ is a stochastic force satisfying:

$$\langle \eta(t)\eta(t')\rangle = \delta(t-t') \tag{3.116}$$

3.4.3 The Itô process

The Langevin process is a special case of a more general stochastic differential equation, which takes the form:

$$\frac{\mathrm{d}x}{\mathrm{d}t} = a(x,t) + b(x,t)\Xi(t) \tag{3.117}$$

where $\Xi(t)$ is a stochastic variable that satisfies $\langle \Xi(t)\rangle = 0$ and $\langle \Xi(t)\Xi(t')\rangle = \delta(t-t')$. The infinite variance at $t = t'$ is problematic, but can be simplified by integrating over a small but finite time interval δt as follows:

$$x(t+\delta t) - x(t) = \delta x = \int_t^{t+\delta t} a(x(t'),t')\mathrm{d}t' + \int_t^{t+\delta t} b(x(t'),t')\Xi(t')\mathrm{d}t' \tag{3.118}$$

It is further assumed that δt is small enough such that $a(x(t'),t') = a(x(t),t)$ and $b(x(t'),t') = b(x(t),t)$. This allows δx to be given by

$$\delta x = a(x,t)\delta t + b(x,t)\delta W(t) \tag{3.119}$$

where we have defined

$$\delta W(t) = \int_t^{t+\delta t} \Xi(t')\mathrm{d}t' \tag{3.120}$$

Firstly, the condition that $\langle \Xi(t)\rangle = 0$ implies that $\langle \delta W(t)\rangle = 0$. Furthermore, $\langle (\delta W)^2\rangle$ is calculated to be:

$$\begin{aligned}\left\langle (\delta W(t))^2 \right\rangle &= \int_t^{t+\delta t}\int_t^{t+\delta t} \langle \Xi(t')\Xi(t'')\rangle \, \mathrm{d}t''\mathrm{d}t' \\ &= \int_t^{t+\delta t}\int_t^{t+\delta t} \delta(t'-t'')\mathrm{d}t''\mathrm{d}t' \\ &= \int_t^{t+\delta t} \mathrm{d}t' \\ &= \delta t\end{aligned} \tag{3.121}$$

$\delta W(t)$ is therefore a stochastic variable acting in the time interval $t \to t+\delta t$ with a zero mean and a variance equal to δt. An Itô process can therefore be implemented in practice for a specified δt, by drawing a random number $w(t)$ from a probability distribution with zero mean and unit variance, and multiplying by $\sqrt{\delta t}$. The stochastic differential equation can therefore be given by:

$$\delta x = a(x,t)\delta t + b(x,t)w(t)\sqrt{\delta t} \tag{3.122}$$

A Langevin process can therefore be simulated by evolving the particle velocity $v(t)$ as follows:

$$\delta v = -\gamma v \delta t + \left(\frac{2k_B T\gamma}{m}\right)^{\frac{1}{2}} w(t)\sqrt{\delta t} \tag{3.123}$$

The particle position can be integrated using the velocity and the time-step δt. This method is applied in Chapter 8.

3.4.4 The Fokker Planck equation

The Fokker Planck equation determines the time-dependent probability density function, $p(x,t)$, for an evolving Markovian system. $p(x,t)$ may be written in terms of a Markovian transition probability density $\mathcal{T}(\Delta x|x,t)$ such that $\mathcal{T}(\Delta x|x,t)\mathrm{d}\Delta x$ is the probability that a transition is made from x to a region at $x+\Delta x$ with size $\mathrm{d}\Delta x$ in an interval τ starting from time t.

$$p(x,t+\tau) = \int_{-\infty}^{\infty} p(x-\Delta x,t)\mathcal{T}(\Delta x|x-\Delta x,t)\mathrm{d}\Delta x \tag{3.124}$$

By performing a Taylor expansion of $p(x-\Delta x,t)\mathcal{T}(\Delta x|x-\Delta x,t)$ in powers of $-\Delta x$, we arrive at:

$$\frac{1}{\tau}(p(x,t+\tau)-p(x,t)) = \sum_{n=1}^{\infty}\frac{(-1)^n}{n!}\frac{\partial^n(\mathcal{M}_n(x,t)p(x,t))}{\partial x^n} \tag{3.125}$$

where:

$$\mathcal{M}_n(x,t) = \frac{1}{\tau}\int_{-\infty}^{\infty} \mathrm{d}\Delta x (\Delta x)^n \mathcal{T}(\Delta x|x,t) \tag{3.126}$$

Equation 3.125 is known as the Kramers-Moyal equation. The Fokker-Planck equation is a special case of the Kramers-Moyal equation, where the system undergoes an asymmetrical random walk. Under this limit, the terms $\mathcal{M}_n$ become zero for $n > 2$ and the Kramers-Moyal equation reduces to the Fokker-Planck equation:

$$\frac{\partial p(x,t)}{\partial t} = \frac{\partial(\mathcal{M}_1 p(x,t))}{\partial x} + \frac{1}{2}\frac{\partial^2(\mathcal{M}_2 p(x,t))}{\partial x^2} \tag{3.127}$$

The Fokker Planck equation can be shown to be equivalent to a stochastic differential equation such as the Ito process. It can be shown that $\mathcal{M}_1$ and $\mathcal{M}_2$ are related to the coefficients in the stochastic differential equation such that:

$$\delta x = \mathcal{M}_1(x,t)\delta t + \sqrt{\mathcal{M}_2(x,t)}\delta W \tag{3.128}$$

3.4.5 Langevin motion along a potential force

Consider a modified Langevin motion, where an additional potential force $U(x)$ is introduced. The stochastic differential equation becomes:

$$\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} = -\gamma\frac{1}{m}\frac{\mathrm{d}x}{\mathrm{d}t} - \frac{1}{m}\frac{\mathrm{d}U(x)}{\mathrm{d}x} + \left(\frac{2k_B T\gamma}{m}\right)^{\frac{1}{2}}\eta(t) \tag{3.129}$$

Converting this into an Itô process therefore produces the following:

$$\delta v = -\frac{\gamma}{m}v\delta t - \frac{1}{m}\frac{\mathrm{d}U(x)}{\mathrm{d}x} + \left(\frac{2k_B T\gamma}{m}\right)^{\frac{1}{2}}\delta W(t) \tag{3.130}$$

In this instance, we consider the overdamped limit, where $\gamma \to \infty$. Under this limit, both the frictional force and the stochastic element become very large. Assuming the potential force is also sufficiently large, we may set δv to zero, and the above simplifies to:

$$\delta x = -\frac{1}{\gamma m}\frac{\mathrm{d}U(x)}{\mathrm{d}x} + \left(\frac{2k_B T}{\gamma m}\right)^{\frac{1}{2}}\delta W(t) \tag{3.131}$$

where we have replaced $v\delta t$ with δx. Now, using equation 3.128, the corresponding Fokker-Planck equation becomes:

$$\frac{\partial p(x,t)}{\partial t} = \frac{1}{\gamma m}\frac{\partial(\frac{\partial U(x)}{\partial x}p(x,t))}{\partial x} + \frac{1}{2}\frac{2k_BT}{\gamma m}\frac{\partial^2(p(x,t))}{\partial x^2} \tag{3.132}$$

The above equation cannot be solved analytically for a time-dependent $p(x,t)$. However, by considering the $t \to \infty$ limit, the left-hand side of the above equation can be dropped, and $p(x,t)$ can be replaced by $p(x)$. The differential equation therefore becomes:

$$\frac{\partial U(x)}{\partial x}p(x) = -k_BT\frac{\partial p(x)}{\partial x} \tag{3.133}$$

This differential equation is solved by the following normalised $p(x)$:

$$p(s) = \frac{1}{Z}\exp\left(\frac{U(x)}{k_BT}\right) \tag{3.134}$$

where Z is akin to the canonical partition function and is given by:

$$Z = \int_{-\infty}^{\infty}\exp\left(\frac{U(x)}{k_BT}\right)\mathrm{d}x \tag{3.135}$$

It is therefore clear that an Itô process in a potential will asymptotically produce a canonical probability density, and that the potential $U(x)$ is similar to a free energy function $F(x)$. The above derivation may be repeated for a non-overdamped example. However, the derivation is more convoluted, and the above derivation is sufficient for demonstrating that the Langevin process is equivalent to a particle moving according to a canonical ensemble.

3.4.6 Kinetic Monte Carlo

The Kinetic Monte Carlo (KMC) approach makes the assumption that the macroscopic dynamics of a system may be represented by a series of discrete random transitions, each with their own reaction rates. Consider such a system with reaction rates $\{\lambda_i(t)\}$. The probability that the next reaction is reaction i, denoted by $P_i(t)$ is given by:

$$P_i(t) = \frac{\lambda_i(t)}{\Lambda} \tag{3.136}$$

where Λ is the sum of all reaction rates:

$$\Lambda(t) = \sum_i \lambda_i(t) \tag{3.137}$$

The probability of reaction i occurring after time $t+\tau$, $P_i(t,\tau)$ is given by the product of the reaction rate λ_i and the probability that a reaction has not occurred $P_0(t,\tau)$.

$$P_i(t,\tau) = \lambda_i P_0(t,\tau) \tag{3.138}$$

For a small time increment, dt, P_0 can be given by,

$$P_0(t,dt) = (1 - \Lambda dt) \tag{3.139}$$

By defining $\tau = ndt$, where n is a large number, $P_0(t,\tau)$ can be given by:

$$\begin{aligned} P_0(t,\tau) &= (1 - \Lambda(t)dt)^n \\ &= \left(1 - \Lambda(t)\frac{\tau}{n}\right)^n \\ &= \exp\left(\Lambda(t)\tau\right) \quad \text{for } n \to \infty \end{aligned} \tag{3.140}$$

$P_i(\tau)$ can therefore be given by:

$$P_i(t,\tau) = \lambda_i(t) \exp\left(\sum_j \lambda_j(t)\tau\right) \tag{3.141}$$

This is the basis of the KMC algorithm, [116] in which events are picked at random, and the time is updated by an exponentially weighted value, δt, given by:

$$\delta t = \frac{1}{\sum_j \lambda_j(t)} \log\left(\frac{1}{\mu}\right) \tag{3.142}$$

where μ is a uniformly distributed random number in the range (0,1]. The complete KMC is implemented as follows:

1. Determine all possible reactions and evaluate $\lambda_i(t)$ for each reaction. Generate two uniformly distributed random numbers, μ_1, μ_2 in the range (0,1].

2. Determine the next reaction by choosing the lowest i that satisfies:

$$\sum_{j'=1}^{j} p_{j'} > \mu_1 \tag{3.143}$$

   where $p_{j'}$ is the probability of reaction j' occurring, and update the configuration of the system

3. Update the simulation time as follows:

$$t \leftarrow t + \frac{1}{\sum_j \lambda_j(t)} \log\left(\frac{1}{\mu_2}\right) \tag{3.144}$$

METHODOLOGY

4

Chapter 3 provides an overview of the theoretical background and computational techniques required for the work carried out throughout this thesis. This chapter concerns the implementation of such techniques. Throughout this chapter, the techniques used in chapters 5 6 7 and 8 are detailed, as well as the parameters used in the algorithms discussed in Chapter 3.

This chapter begins with the standard parameters of the molecular dynamics simulations used in chapters 5, 6 and 7. The force fields for calcite and water are also discussed. In Section 4.2, the structure of the calcite data files are discussed, as are the means through which calcite terraces, steps and kinks are exposed. This chapter also details computational techniques used in molecular simulation to calculate stress distributions in Section 4.3, and determine average lattice spacings in Section 4.4. Section 4.5 details the types of reaction coordinates used in rare event sampling simulations. This is accompanied by a discussion in Section 4.6 on how to calculate physically meaningful adsorption free energies from the results of such rare event sampling methods. Finally, Section 4.8 details the Kinetic Monte Carlo method used to simulate step growth events in the presence of impurities.

4.1 Simulation details

Throughout this thesis, molecular dynamics simulations are run using the molecular dynamics package Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) [117]. The positions and velocities are updated using the Velocity Verlet algorithm with a time-step of 1 fs. For all simulations, a canonical ensemble is maintained at a temperature of 300 K using the Nosé-Hoover thermostat with a relaxation time of $\tau_{\text{damp}} = 100$ fs. Where explicitly stated, an isobaric, isothermal ensemble is maintained at zero pressure using a Nosé-Hoover barostat with a damping parameter of $\tau_{\text{damp}} = 1$ ps. The long range electrostatics are handled using the PPPM method with an accuracy of 10^{-4}.

Unless explicitly stated, the inter- and intra-molecular interactions of calcium carbonate, as well as their interactions with water, are described using the force fields of Raiteri et al. [45], which are parametrised in order to replicate the experimental solubility of calcite. These force fields are widely regarded as optimal for reproducing thermodynamic quantities for calcite. The inter- and intra-molecular interactions of water are modelled with SPC/Fw, [118] in line with the work of Raiteri et al. Where impurities are modelled, their force fields are discussed individually.

4.2 Simulation cell structures

In chapters 5 and 6, simulations concern exposed terrace, step and kink sites. In order to expose such sites, systems consist of a slab of calcite, periodic in the x- and y-directions, and a $\sim$4 nm water-filled gap dividing the slab from its periodic images in the z-direction. The calcite consisted of six repetitions of the rhombohedral calcite unit in the z-direction. The z-length of the water-filled gap is not given by exactly 4 nm, as the correct zero-pressure volume of the cell as a function of water molecules cannot be calculated without simulation. Instead, the cell is filled with roughly the correct number of water molecules to produce a gap with a length of approximately 4 nm. The simulation cell is then equilibrated and relaxed under zero pressure (NPT). The average z-length of the cell is used for future simulations. This process is repeated for every calcite structure (terrace, step and every unique exposed kink), as well as for any impurity which is introduced into the simulation cell. Every simulation including these NPT simulations consisted of an equilibration process, during which the calcite was frozen and the water was equilibrated for 1 ns. The calcite is then unfrozen and the entire simulation cell is equilibrated for a further 100 ps. In every simulation, the position of the calcite slab is kept in place with a series of harmonic tethers acting on a randomly selected group of Ca-ions roughly in the centre of the slab.

For simulations involving terrace and step sites, the x-length of the simulation cell is equal to 12 repetitions of the calcite unit cell. A monoclinic skew is introduced in the y-direction such that the sum of the y-length and monoclinic skew is equal to 12 repetitions of the unit cell. For simulations involving steps, the step added a seventh layer of calcite to the simulation, and was periodic in the x-direction i.e. had a length of 12 units in the x-direction. In the y-direction, the step spanned six calcite units. The simulation cell therefore consisted of both an acute and obtuse step.

For simulations involving expressed calcite kink sites, the calcite is rotated, and the lengths and monoclinic skew of the simulation box adjusted such that an offset of one Ca unit existed at the boundaries of the simulation box. This is done such that one kink site is exposed at each step. This set-up is detailed in Figure 4.1. The elevated step again spans the entire length of the simulation box along the x-direction, and again spans 6 repetitions of the unit cell in the y-direction. To preserve the continuity of the crystal at the boundaries, the number of repetitions of the calcite unit cell must be adjusted. To expose a- or d-kinks, the box size corresponds to 9 and 13 repetitions of the unit cell in the x- and y-direction respectively, To expose b- and c-kinks, the box size corresponds to 11 and 13 repetitions of the unit cell in the x- and y-direction respectively. Unless explicitly stated, both kink-terminating ions are tethered to their lattice site throughout simulations, using a harmonic tether with a spring constant of 100 kJ/mol/Å^2.

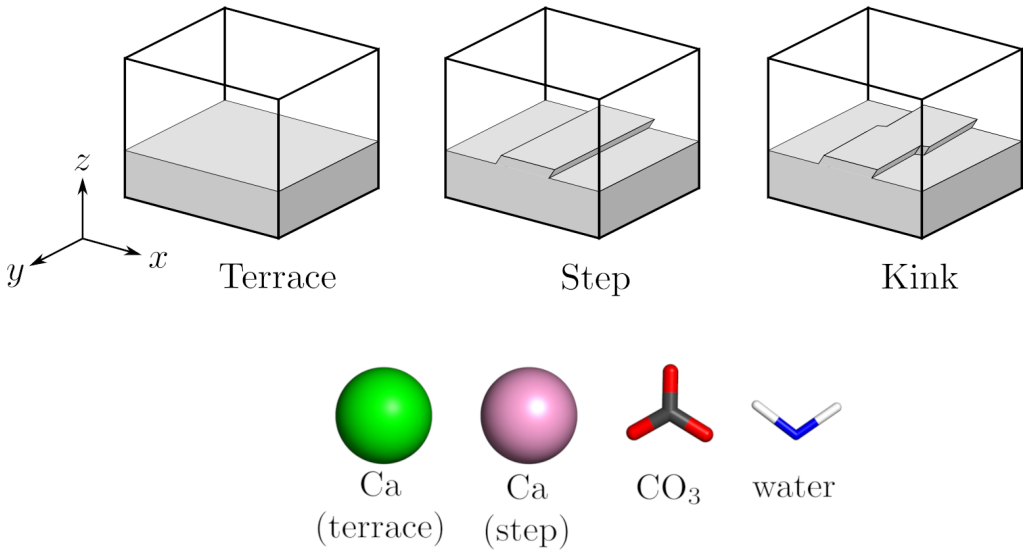

Figure 4.1: (top) Schematic of simulation cells used to expose terrace, step and kink sites. The crystal slab in periodic in the x- and y-directions, and is separated from its periodic image in the z-direction by a ∼4 nm gap filled with water. (bottom) depictions of how Ca, CO_3 and water are typically represented throughout this thesis. The Step Ca ions are coloured in pink for easier visualisation of the step.

4.3 Calculating stress distributions

In Chapter 7, we wish to calculate stress distributions for bulk crystals containing nanoparticles. However, the definition of stress is only applicable in the continuum limit, and is far more ambiguous in the atomistic limit where the space largely consists of the vacuum between atoms. However, since the stress tensor of a supercell is just the volume-average of the atomic virials (stress tensors for individual atoms), it is common to average the atomic virials over an appropriate local volume to produce a local stress field. Branicio and Srolovitz [119] presented a general method for this, later applied to titania nanoparticles [120] and calcite defects, [99] that involves calculating the time averaged atomic virial tensors $\langle W_{\alpha\beta}\rangle_i$, and multiplying their value by a normalised smearing function $P(\vec{x}-\vec{x_i})$ where x_i is the position of the atom. A continuous stress field is obtained by summing the product of the averaged virial and smearing function over all N atoms.

$$\Pi_{\alpha\beta}(r) = \sum_{i \in N} P(r - \langle r_i \rangle)\langle W_{\alpha\beta}\rangle_i \tag{4.1}$$

Here, $P(r - \langle r_i \rangle)$ is chosen to be a radial Gaussian, namely:

$$P(r) = \frac{27}{R^3(2\pi)^{3/2}} \exp\left(-\frac{9r^2}{2R^2}\right) \tag{4.2}$$

where R is equal to three standard deviations of the Gaussian, and parametrises the distance beyond which any contributions to the stress field can be considered negligible. In this thesis, R is set to 9 Å. α and β label the Cartesian components of the virial stress tensor. The hydrostatic stress field, which is equivalent to the negative of the pressure field, is therefore calculated by taking the average of the diagonal elements, $sum_\alpha \left(\Pi_{\alpha\alpha}(\vec{x})\right)/3$. Using the calculated atomic virials from simulations, a grid of points with spacing 1 Å is constructed, and Eq. 4.1 is evaluated at each point.

While the integral of Eq. 4.1 over all space $\vec{x}$ recovers the correct stress tensor of the system, interpreting it at a local level is a little delicate. If the local stress tensor is uniform across a volume that exceeds the smearing volume then it is a physically meaningful measure of stress. Otherwise, its value is sensitive to the means of smearing and so it is, at best, an order-of-magnitude estimate. Crucially though, the sign is a faithful indicator of whether the local stress is compressive (negative) or tensile (positive), and it is this feature of the stress field that we are primarily interested in.

4.4 Radial Distribution Function and lattice spacings

In Chapter 7, we wish to calculate the average x-, y- and z-spacings between atoms, as well as distances between planes. For a perfect crystal, this calculation is trivial, as the average lattice spacings are homogeneous. In this instance, one can simply divide the simulation cell lengths by the number of crystal units in the cell. However, for imperfect crystals, such as those with defects or incorporated impurities, the lattice spacings are not homogeneous, and we may wish to neglect the effect of the defect on the size of the simulation cell. If, for example, a vacancy or defect exists within the lattice, the average lattice spacing may not correspond with the dimensions of a simulation cell. In this limit, the average lattice spacing can be calculated using a modification of the Radial Distribution Function (RDF). The RDF is given as a sum over the entirety of the simulation cell, or over a specific subset of particles as follows:

$$g(r) = \frac{1}{2\pi r N_0 \rho_0} \sum_{i=1}^{N_0} \sum_{j>i}^{N_0} \delta(r - r_{ij}) \tag{4.3}$$

where N_0 is the number of particles in the subset included in the RDF and ρ_0 is their average number density. Consider now the Cartesian components of the RDF. Each component does not need to be normalised by dividing by the area of a shell of radius r. Therefore, a vector equivalent of equation 4.3 is given by:

$$\vec{g}(\vec{r}) = \frac{2}{N\vec{\rho_0}} \sum_{i=1}^{N_0} \sum_{j>i}^{N_0} \delta(\vec{r} - \vec{r_{ij}}) \tag{4.4}$$

where $\vec{\rho_0}$ is a vector of each Cartesian component of the density, i.e. (ρ_x, ρ_y, ρ_z). This RDF can be used to calculate distance between planes by defining a (normalised) vector $\vec{V}$ normal to the plane over which we wish to calculate a distribution of atomic distances. By taking its dot product with each Cartesian component, the distribution of distances between planes, denoted here by $g_V(r)$, is given by:

$$g_V(r) \propto \sum_{i=1}^{N_0} \sum_{j>i}^{N_0} \delta((\vec{r} - \vec{r_{ij}}).\vec{V}) \tag{4.5}$$

where the RDF is no longer normalised, as its magnitude does not affect the calculation of the average lattice spacing. Finally, this can be used to calculate the distribution of distances between neighbouring atoms by introducing a cut-off distance which acts in the direction perpendicular to $\vec{V}$. Denoted by $d_V(r)$, the distribution of distances between neighbouring atoms is given by:

$$d_V(r) \propto \sum_{i=1}^{N_0} \sum_{j>i}^{N_0} \delta((\vec{r} - \vec{r_{ij}}).\vec{V}) \epsilon(\vec{r_{ij}}, d_0) \tag{4.6}$$

where $\epsilon(\vec{r_{ij}}, d_0)$ is given by:

$$\epsilon(\vec{r_{ij}}, d_0) = \begin{cases} 1 & \text{if } r_{ij}^2 - (\vec{r_{ij}}.\vec{V})^2 < d_0^2 \\ 0 & \text{otherwise} \end{cases} \tag{4.7}$$

where d_0 is a cut-off distance perpendicular to $\vec{V}$, beyond which we do not calculate the distances between atoms. The purpose of $\epsilon(\vec{r_{ij}}, d_0)$ is to neglect atoms which are far apart in distance, but whose distance between selected planes is relatively small. Equation 4.6 can be used to calculate the distribution of distances between planes. For a crystalline material, equation 4.6 should produce a series of spikes separated by regions where the values are zero or very low in magnitude. The average lattice spacing, $\langle d_V \rangle$, can therefore be computed using the following integral:

$$\langle d_V \rangle = \frac{\int_a^b r d_V(r) \mathrm{d}r}{\int_a^b d_V(r) \mathrm{d}r} \tag{4.8}$$

where a and b are values encompassing the first spike.

4.5 Rare event sampling details

Where we have explicitly used molecular simulation throughout this thesis, free energies are calculated using metadynamics. In Chapter 8 we apply other rare event sampling techniques to a toy model. In this case, the parameters used are all discussed within the chapter. Unless explicitly stated, metadynamics was implemented using PLUMED 2.0 [121] Well-tempering was used in every metadynamics simulation.

Longer time-scales were accessed through the use of the multiple walkers algorithm [122], in which a set of metadynamics simulations are simultaneously run across different processors, and the cumulative Gaussians for each separate simulation are periodically appended to all other simulations. In this thesis, we applied well-tempering to 12 parallel simulations. Throughout this thesis, metadynamics simulation times varied between approximately 300 ns and 3 μs, depending on the volume of reaction coordinate space required.

In this thesis, we typically consider two sets of reaction coordinate: the first consists of the spatial coordinates of a particle or centre-of-mass position of several particles; the second is the degree of hydration of a Ca ion. For lone Ca ions, we adopt the standard procedure of using the coordination number between Ca and water oxygen (O_w). The coordination number (CN) is given by the continuous, differentiable function:

$$\mathrm{CN} = \sum_{i,j} \frac{1 - \left(\frac{r_{ij} - d_0}{r_0}\right)^n}{1 - \left(\frac{r_{ij} - d_0}{r_0}\right)^m} \tag{4.9}$$

where r_{ij} is the distance between atoms i and j, and d_0, r_0, n and m are parametrisable constants. Throughout this thesis, we use the following values: $d_0 = 2.1$ Å; $r_0 = 1$ Å; $n = 4$; $m = 8$. Equation 4.9 has already been successfully applied to lone Ca ions [45] and step Ca ions [43]. However, equation 4.9 is not applicable for dehydrating Ca-terminated kink sites. This will be discussed in detail in Chapter 5.

When using metadynamics in this thesis, we typically deposit a Gaussian function every 1 ps, or 1000 timesteps. The only exception to this rule is when the position of aspartate (asp) is biased in Chapter 6. In this instance, the rate of deposition is equal to 2 ps, as asp has a greater mass than Ca or CO_3, and diffuses slower, thus requiring a less frequent deposition of Gaussian functions. The Gaussian height is always equal to 2.5 kJ/mol, ($\approx k_B T$). We set the Gaussian width, σ, to 0.2 Å for every reaction coordinate with the only exception being for when equation 4.9 is applied to dehydrate lone Ca ions. In this instance, σ is set to 0.1 because the free energy landscape consists of much sharper features, and therefore requires a smaller Gaussian width.

4.6 Adsorption free energies

Section 3.3 details how free energy functions may be obtained from simulations using rare event sampling. For a free energy function separating an adsorbed state and a dissolved states, the simulation free energy difference, denoted by ΔG_{sim}, can be defined as the free energy difference between the free energy minimum (adsorbed), and the region where the free energy becomes flat (dissolved). Figure 4.2 shows a toy example of such a free energy surface. The adsorbed state can be represented by the region between points a and b. The dissolved state can be represented by the region between points c and d.

Adsorption free energies are distinctly different from simulation free energy differences. Where simulation free energy differences are dependent on the simulation conditions, the adsorption free energy, denoted by ΔG_{ads}, is a standard reference free energy given by

$$\Delta G_{\text{ads}} = -k_B T \log\left(\frac{P_{\text{ads}}}{P_{\text{diss}}}\right) \tag{4.10}$$

where P_{ads} and P_{diss} are the equilibrium probabilities of finding a solute in an adsorbed and dissolved state respectively when the solute is present in solution at a concentration of 1 M. As the adsorption free energy has a well-defined meaning, and does not depend on any experimental or simulation conditions, it makes an ideal quantity to calculate. ΔG_{ads} can be computed with an entropic correction to ΔG_{sim} as we now proceed to discuss.

Here we consider a rare event sampling simulation with the following set-up: metadynamics is used to derive a free energy surface as a function of the position of the adsorbate normal to the {10.4} plane. In this thesis, this position is given by the z-position of the adsorbate (see Figure 4.1). The (x,y)-coordinates of the empty lattice site which facilitates adsorption are set to (0,0). The x- and y-coordinates of the adsorbate are both constrained to a region spanning between -2 Å and 2 Å, using harmonic walls with a spring constant, k, of 100 kJ/mol/Å. An additional offset, o, is added to the harmonic walls such that the harmonic potential is given by $k(x - a + o)^2$ when $x > a$. This is done in order to make the boundaries more rigid in order to minimise the contribution of the regions beyond the walls to the configurational freedom of the dissolved solute. Here, we set o to 0.5 Å. The simulation set-up is depicted in Figure 4.3.

In order to derive ΔG_{ads}, from simulations, first consider the probability of adsorption in the context of the simulation. Consider the example free energy surface in Figure 4.2. The solute can be considered adsorbed between the regions a and b. Between b and c, the solute is not fully adsorbed, but is still somewhat associated with the surface. When the free energy surface becomes flat, i.e. beyond c, the solute is fully dissolved. From Chapter 3 we know the probability of finding the adsorbate between points a and b

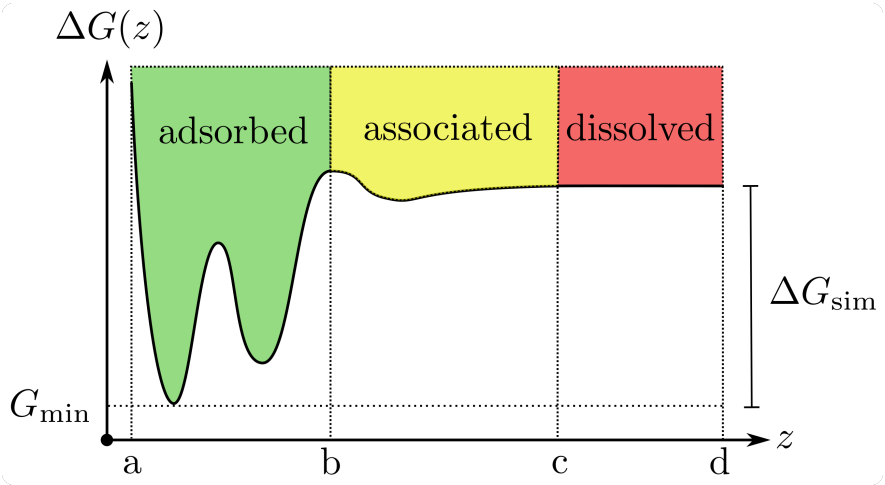

Figure 4.2: Example of a free energy surface in which ΔG_{sim}, G_{min}, a, b, c and d are defined. Note that this example does not represent the exact shape of every free energy surface we study in this thesis, but the theory equally applies in all settings, as long as the free energy surface is flat between c and d. It should also be noted that the exact positions of a and b are not too important, as the integral over the probability density will be dominated by the regions near the thermodynamic minimum.

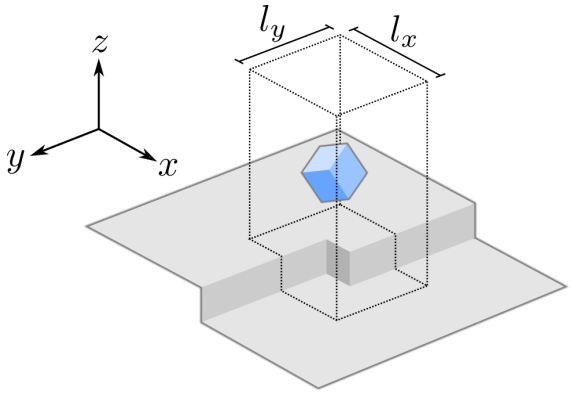

Figure 4.3: Schematic depicting the simulation set-up used to derive adsorption free energies. Here, the adsorbate (blue) is constrained within a pillar using a series of harmonic walls.

(adsorbed), P_{a-b}, divided by the probability of finding the adsorbate between points c and d (dissolved), P_{c-d}, is given by

$$\frac{P_{a-b}}{P_{c-d}} = \frac{\int_a^b \exp\left(-\frac{G(z)}{k_B T}\right) dz}{\int_c^d \exp\left(-\frac{G(z)}{k_B T}\right) dz} \tag{4.11}$$

Since $G(z)$ should be flat between regions c and d, we can define $G(z)$ between c and d to be $G(z) = G_{\text{min}} - \Delta G_{\text{sim}}$, where G_{min} is the minimum value of $G(z)$ and ΔG_{sim} is the simulation free energy difference (see Figure 4.2). Note that ΔG_{sim} is negative, hence the minus sign. The ratio of probabilities therefore

simplifies to

$$\frac{P_{a-b}}{P_{c-d}} = \frac{\exp\left(-\frac{\Delta G_{\text{sim}}}{k_B T}\right)}{(d-c)} \int_a^b \exp\left(-\frac{G(z) - G_{\text{min}}}{k_B T}\right) dz \tag{4.12}$$

We now consider the true probabilities of finding the adsorbate in an adsorbed or dissolved state. The probability of the impurity being adsorbed is entirely determined by the probability of finding the impurity between a and b. Assuming the dissolved impurity behaves as an ideal gas, the probability of finding the impurity in a dissolved state is proportional to the average volume it occupies in solution at a concentration of 1 M, $V_{1\text{M}}$. This probability can be related to P_{c-d} by multiplying by the ratio of accessible volumes i.e. $V_{1\text{M}}/l_x l_y (d-c)$. We therefore have:

$$\begin{aligned} P_{\text{ads}} &\propto P_{a-b} \\ P_{\text{diss}} &\propto P_{c-d} \frac{V_{1\text{M}}}{l_x l_y (d-c)} \end{aligned} \tag{4.13}$$

which simplifies to

$$\frac{P_{\text{ads}}}{P_{\text{diss}}} = \frac{P_{a-b}}{P_{c-d}} \frac{l_x l_y (d-c)}{V_{1\text{M}}} \tag{4.14}$$

It should be noted that the accessible volume in the simulation is not exactly equal to $l_x l_y (d-c)$, as the solute is still able to explore regions beyond the harmonic walls used to constrain the x- and y-positions of the particle within the column. However, given the simulation set-up we apply, the high spring constant used, as well as the use of offsets, minimises this contribution to the accessible volume. We calculate that, ultimately, the regions beyond the harmonic walls reduce the adsorption free energy by 0.05 kJ/mol in every relevant simulation in this thesis. Since this value is much smaller than the likely errors of the simulations in this thesis, we neglect this correction in this thesis.

Equation 4.14 can finally be combined with equations 4.10 and 4.12 to give

$$\begin{aligned} \Delta G_{\text{ads}} &= -k_B T \log\left(\frac{P_{\text{ads}}}{P_{\text{diss}}}\right) \\ &= \Delta G_{\text{sim}} + k_B T \log\left[\frac{V_{1\text{M}}}{l_x l_y \int_a^b \exp\left(-\frac{G(z) - G_{min}}{k_B T}\right) dz}\right] - 0.05 \text{ kJ/mol} \end{aligned} \tag{4.15}$$

Note that the result is no longer dependent on c and d. This can ultimately be represented as an entropy correction as follows:

$$\Delta G_{\text{ads}} = \Delta G_{\text{sim}} + T \Delta S_{\text{ref}} \tag{4.16}$$

where the correction ΔS_{ref} is given by

$$T\Delta S_{\text{ref}} = k_B T \log \left[\frac{V_{1M}}{l_x l_y \int_a^b \exp\left(-\frac{G(z)-G_{min}}{k_B T}\right) dz} \right] - 0.05 \text{ kJ/mol} \quad (4.17)$$

where the correction of -0.05 kJ/mol is discussed above. It should be noted that the exact values of a and b are not especially important, as long as they fully encompass the free energy well associated with adsorption. $V_{1\text{M}}$ is equal to 1660.5 Å^3.

4.7 Minimum Free Energy Pathways

In chapters 5 and 6, we use the nudged elastic band method (see Section 3.3.11) to calculate the Minimum Free Energy Pathway (MFEP). In this thesis, we only calculate the MFEP for two-dimensional free energy surfaces. The elastic band consists of a series of beads separated by a distance of 0.1 (regardless of units) and connected with a series of springs with a spring constant of 200 kJ/mol/ν_i, where ν_i is the unit corresponding to the i-direction (e.g. Å for directions with units of distance). Initial estimates of the MFEP are made by connecting the relevant free energy minima with straight lines. The elastic band is updated using the steepest descent algorithm, where i-components of the k^{th} bead position, $x_i^{(k)}$ are updated as follows:

$$x_i^{(k)} \leftarrow x_i^{(k)} + l \frac{\partial F}{\partial x_i^{(k)}} \quad (4.18)$$

where F is the free energy surface (note that its derivative with respect to $x_i^{(k)}$ is the i-component of the mean force) and l is a parameter determining the magnitude of the displacement with respect to the mean forces. In this thesis, we typically set l to a low value of 0.001 in order to ensure convergence. The number of iterations was typically set to 1000, although this number was increased if more time was required for convergence.

4.8 Kinetic Monte Carlo model

The Kinetic Monte Carlo (KMC) method (see Section 3.4.6) is useful for simulating processes which occur over time- and length-scales inaccessible to molecular simulation. The propagation of kinks and advance of a step is an example of this. In this thesis, we use a KMC scheme to simulate the advancement of a single step along a terrace. The configuration is represented on a two dimensional grid of size 2^{10} by 2^{10}. A value of zero on the grid represents lower terrace and a value of one represents an elavated terrace. The

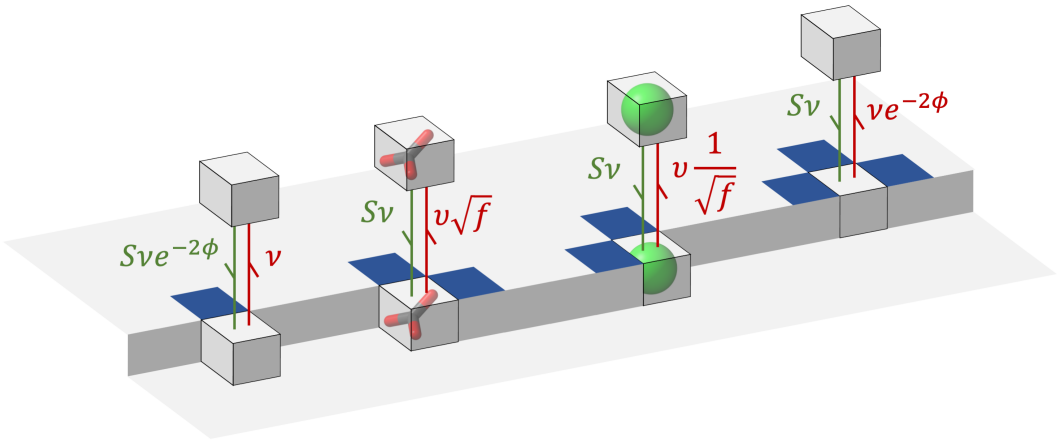

Figure 4.4: Diagram of the KMC model used to describe the growth of a step or island. The attachment and dissolution rates are determined by the number of nearby units, depicted in blue, as discussed in the text.

system is integrated over time using the standard KMC approach discussed in Section 3.4.6.

4.8.1 Attachment and detachment rates

The model used in this thesis is a semi-Kossel model, in which all attachment and detachment rates are assigned a base attachment rate of ν. Figure 4.4 shows a schematic detailing the parameters that define the reaction rates used in simulations. The attachment rates for units to kink sites is given by $S\nu$, where S is the supersaturation given by equation 2.1.

The parameter ν also determines the detachment rate for all units. For a completely Kossel model, the kink dissolution rate for all units would be given by ν. However, we include an additional factor, f, which determines the relative abundance of Ca and CO_3 ions terminating the kink. The kink dissolution rates are given given by $\nu/\sqrt{f}$ for Ca and $\nu\sqrt{f}$ for CO_3. The values of f were individually parametrised, for both acute and obtuse steps, to data on step velocities as a function of solution stoichiometry. [123] Since these functions do not peak at unity, the relative attachment/detachment rates for Ca and CO_3 cannot be equal. Instead, the step velocity must peak when the populations of Ca and CO_3 terminated kink sites are equal. Since f determines the ratio of the populations of Ca- and CO_3-terminated kinks, the position of the peak of the step velocity is therefore given by $1/f$. Using the results in [123], we calculate $f = 1.5$ for acute kinks and $f = 0.2$ for obtuse kinks. Furthermore, since we are only interested in determining velocity inhibition, we are able to use units where ν is set to 1.

For step sites, we further assume that attachment and dissolution rates are identical for all units. We make the assumption that the addition of one neighbouring unit universally lowers the free energy of attachment by

$2\phi k_B T$. We therefore use ϕ to determine the attachment rate, lowering it with respect to kink sites by a factor of $e^{2\phi}$. The dissolution rate for all step sites is kept at ν. For step dissolution, we set the attachment rate to ν and the dissolution rate to $\nu e^{-2\phi}$. The (dimensionless) value of ϕ is set to 3 in line with the model used in ref [62]. We make the further approximation that attachment at terrace sites and dissolution from within the terrace are both too rare to model. If a unit is left with no neighbour sites on the elevated terrace, it is set to dissolve immediately.

4.8.2 Impurities

As well as Ca and CO_3 units, impurities may also be included in simulations for studies of effects such as kink-blocking. In this model, we represent impurities on the grid with a value of 2. Similar to Ca and CO_3 units, the attachment rates for impurities at kink sites are assumed to be independent of the kink type to which they attach. In this model, the rates for attachment, λ^{+}_{imp}, and dissolution, λ^{-}_{imp}, are given by:

$$\begin{aligned} \lambda^{+}_{\text{imp}} &= \frac{\nu}{\sqrt{K_{sp}}} a_{\text{imp}} \\ \lambda^{-}_{\text{imp, k}} &= \frac{\nu}{\sqrt{K_{sp}}} \exp\left(\frac{\Delta G_{\text{ads,k}}}{k_B T}\right) \end{aligned} \tag{4.19}$$

where a_{imp} is the activity of the impurity in the solution and $\Delta G_{\text{ads,k}}$ is the free energy for the impurity adsorbing to kink site k. The adsorption free energies used in this model are calculated through simulations using the methods detailed in sections 4.5 and 4.6. With this set-up the rate of attachment per impurity concentration is equal to that of all other units.

4.8.3 Step velocities

The seemingly most straightforward approach to calculating the growth rate of a crystal would be to simulate a growth island and calculate its growth rate as a function of the impurity concentration. However, we do not opt for this approach, as it makes the calculation of acute and obtuse step velocities difficult. In addition to this, the simulation would have to be terminated before the boundaries of the cell are reached, resulting in poor statistical sampling. Instead, we run individual simulations of an acute and obtuse step by simulating the growth of a single step. From this, we obtain individual step velocities. The boundary conditions are set up such that the units on the boundary will have no interactions beyond the boundaries, akin to growing the step along a confined region. Initially, the entirety of the first five columns are populated with an elevated terrace, and the first column is frozen in place to prevent dissolution. An initial equilibration run is carried out, followed

by a production run, in which the step velocities are calculated. The step is allowed to grow indefinitely by shifting the entire crystal backwards by 600 rows every time the crystal reaches the boundaries of the simulation. The first 600 rows are immediately discounted whenever this occurs.

5 KINK GROWTH OF PURE CALCITE

To recap from Chapter 2, growth of calcite occurs through the propagation of steps. These steps propagate through the adsorption of ions onto the step to form a new step. After this new step sufficiently surpasses a critical length, the new step continues to grow outward via the adsorption of ions into the corner (kink) sites they form at the edges of the step. It is easy for one to visualise this kink growth as a series of single-step processes in which solutes overcome a free energy barrier and adsorb directly into the lattice site. This free energy barrier is often thought to be due to the dehydration of negatively-charged Ca ions, as simulations have previously calculated large residence times for water molecules interacting with calcite Ca ions. Throughout this chapter, we test the validity of this view of calcite growth from the molecular scale using rare-events sampling to calculate the free energy surfaces relevant to the adsorption of ions to calcite kinks.

5.1 Simulation details

The simulation parameters and force fields used for molecular dynamics simulations are given in Chapter 4. In this chapter, we run simulations involving calcite kink sites. The simulation cell dimensions are described in Section 4.2 and can be visualised in Figure A1 in the appendix. In this chapter, we define the location of the empty kink lattice sites (the location roughly at which a solute would adsorb) as having (x,y)-coordinates of (0,0). The x- and y-coordinates of the solute (either Ca or C) are both constrained to a region spanning between -2 Å and 2 Å, using harmonic walls with a spring constant of 100 kJ/mol/Å. An offset was also introduced to the harmonic potential to increase its rigidity (see Section 4.6). In this chapter, we are only interested in calculating adsorption free energies, rather than reaction pathways. Therefore, for the sake of computational efficiency, we only apply a bias potential to the spatial coordinate of the adsorbing ion normal to the {104} terrace. In this chapter, as well as all other chapters, this corresponds to the z-position (see Section 4.2, notably Figure 4.1). For all metadynamics parameters used, see Section 4.5.

5.2 Reaction coordinates for dehydration

As discussed in Chapter 4, lone Ca ions are dehydrated using the coordination number collective variable, given by equation 4.9. However, equation 4.9 is not easily applied to Ca-terminated kink sites. This is because equation 4.9

must be parametrised to the radial distribution function between Ca and O_w, such that its value is one for most of the first peak (which represents a coordinated atom) and decays to zero near the beginning of the second peak (which represents free water). The water structures near acute kink sites are not compatible with these conditions. This can be observed in the radial distribution functions for water near the kink sites studies, shown in Figure 5.1, where the peaks representing bound and free water are not separated by a region of zero density. Instead, we opt for a different approach, where we consider the distance between a kink Ca ion and the nearest water molecule to it. This distance can be approximated as:

$$\mathrm{ND} = \frac{\beta}{\log \sum_j \exp\left(\frac{\beta}{r_j}\right)} \tag{5.1}$$

where r_j is the distance to atom j and β is a parametrisable constant with units of distance, which we set to 10 nm. Equation 5.1 is used to approximate the distance between the Ca atom and the nearest water oxygen atom, while still allowing the function to be differentiable. A similar approach has already been applied effectively to NaCl, [124] where exponential functions are used to weight coordination states by distance. Here, we opt for the nearest distance approximation for its more intuitive output. In all simulations where the binding of CO_3 ions is measured, equation 5.1 is used to dehydrate kink Ca ions. Although the issue with coordination numbers only exists for acute sites, equation 5.1 is also applied to CO_3-terminated obtuse sites for consistency. When equation 5.1 is applied to metadynamics simulations, a Gaussian width of 0.2 Å is used.

It is worth reiterating that equation 5.1 is applied on the condition that we are only attempting to calculate adsorption free energies, rather than reaction pathways. Its applicability for determining reaction pathways is most likely limited; it is unlikely that the adsorption of a solute will involve the displacement of the nearest water molecule from the kink before other water molecules are displaced.

5.3 Water residence times at CO_3-terminated kink sites

As discussed in Chapter 2, water residence times have been calculated for water coordinated with lone Ca ions, as well as Ca ions forming calcite terraces and steps. Little attention has been given to water coordinated with CO_3 ions, as these residence times are relatively small for step sites. It is therefore expected that residence times of water coordinated with CO_3-terminated kink sites should also be small. Nevertheless, it is necessary that we determine whether water molecules coordinated with CO_3-terminated kink sites

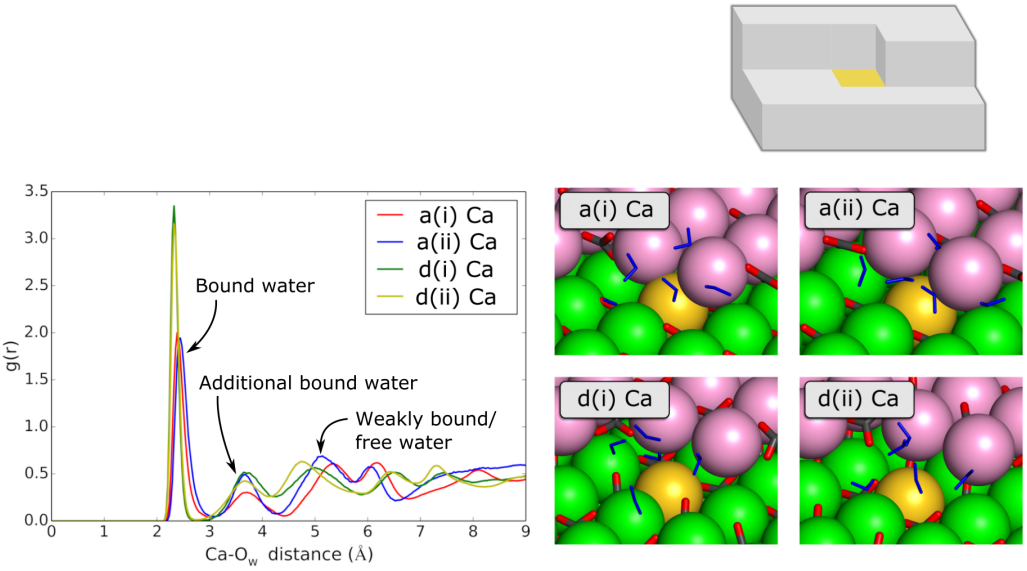

Figure 5.1: Radial Distribution Functions (RDFs) for distance between kink Ca ion (shown in gold and depicted in schematic above) and water oxygen molecules. Snapshots from simulations show the kink Ca (shown in gold) and the local water molecules for each kink type.

presents a barrier to adsorption.

To calculate the residence time of water at the kink site, a series of simulations with a total simulation time summing to about 200 ns were run. From these simulations, the atomic positions within a region of about 5 Å from the kink site were outputted every picosecond. The data from the simulation was then post-processed to build a density map of water oxygen atoms. Every region of high density was identified, and their positions were analysed to identify the water molecule coordinated with the kink site.

Once the water at the kink site (depicted in Figure 5.2) was identified, the identity of the nearest water molecule to the local density peak was identified at every output. The time each water molecule spent at this site was tracked, and a transition was recorded if the identity of the nearest water molecule changed for more than 2 ps as per standard practice [125]. The time the previous atom spent in residence was then added to the list of residence times. The method used to deduce residence times was one initially used by Impey et al. [125] on ionic crystals, and later applied by others to calcite [53, 51]. In this method, the list of residence times is used to calculate a survival function, $P(t)$. Defined as the probability that an atom remains in residence after time t, $P(t)$ is known to obey a multi-exponential dependence on time [125] given by:

$$P(t) = \sum_{i=1}^{n} a_i e^{-\frac{t}{\tau_i}} \tag{5.2}$$

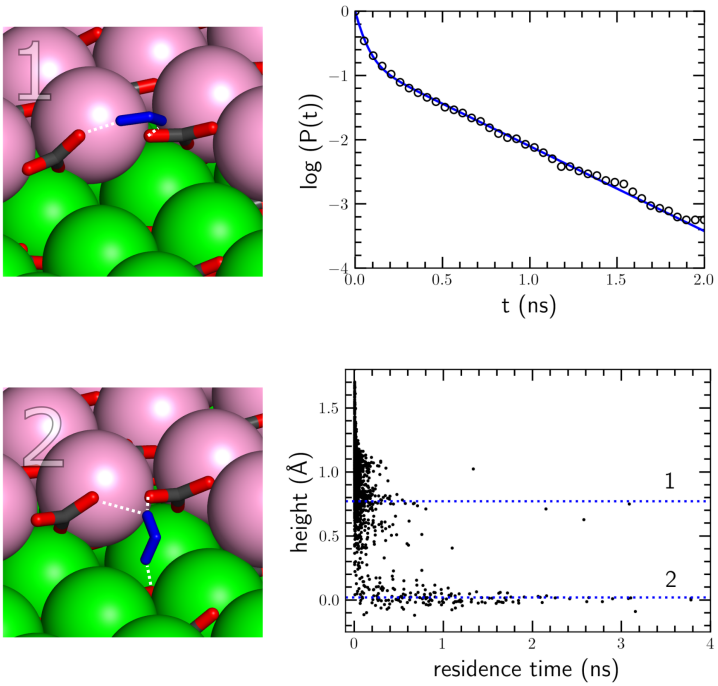

Figure 5.2: Results for simulations of water molecules at a carbonate-terminated kink site. The upper-left and lower-left panels show snapshots of sites 1 and 2. Here, all other water molecules in the simulation are omitted from the image. The top-right panel shows a logarithmic plot of the survival function for the water oxygen atoms at sites. The survival function is fitted with two exponential functions to correspond with sites 1 and 2. The bottom-right panel shows a scatter plot of the average height of the water molecule against its residence time. Two distinct water positions can be deduced from this plot. These states correspond to states 1 and 2.

where n represents the number of unique configurations that correspond to a bound state. τ_i is the mean residence time describing configuration i, and a_i is a normalisation coefficient. In this study, we fit the results by increasing the number of exponentials, starting at one, until a sufficiently accurate fit is observed.

Since these simulations are fairly time consuming, we proceeded to calculate the residence time for only a single kink site. In this instance, we choose the a(i) CO_3 kink site. The data from the simulation was post-processed, and a single water density peak was observed near the kink site. The survival function $P(t)$ was calculated for the nearest water molecule to this density peak. The survival function is shown in Figure 5.2. Fitting the survival function to equation 5.2 revealed that a minimum of two exponential terms are required to fit the data. Analysis of the average water position revealed two distinct configurations, which are also shown in Figure 5.2. In both configurations, hydrogen atoms in water interacts with an oxygen atoms in carbonate. In configuration 1, the hydrogens interact only with the upper terrace. In

configuration 2, the hydrogens also interact with the lower terrace.

Fitting to equation 5.2 produced the following:

$$P(t) = a_1 \exp\left(-\frac{t}{\tau_1}\right) + a_2 \exp\left(-\frac{t}{\tau_2}\right) \tag{5.3}$$

where:

$$\begin{aligned} a_1 &= 0.541 \quad \tau_2 = 0.058 \text{ ns} \\ a_2 &= 0.459 \quad \tau_1 = 0.757 \text{ ns} \end{aligned} \tag{5.4}$$

The values τ_1 and τ_2 correspond to the mean residence times of the water molecules in each configuration. The larger of the two residence times is below 1 ns. Compared to typical simulation times $\sim 100ns$, the water residence time is relatively small, and it can be assumed that water exchanges will happen sufficiently rapidly such that a reaction coordinate used to bias the hydration of the CO_3-terminated kink site will not be necessary. Seeing as the residence times calculated are relatively small, we further reasonably assume the same applies to water molecules coordinated with all CO_3-terminated kink sites.

5.4 Ca kink free energy surfaces with dehydration

Dehydrating Ca ions is important for modelling certain reactions, such as the adsorption of a lone Ca or CO_3 ion onto a step site [43]. In order to determine whether dehydration is necessary here, we adopt the method discussed in Section 4.5 to dehydrate the lone Ca ions, and calculate the free energy as a function of the Ca-kink z-distance and the water coordination number of the lone Ca ion, CN (equation 4.9). Including this second reaction coordinate comes at computational expense, so we wish to avoid repeating this process for every kink site. Instead, we select two CO_3-terminated kink sites, the a(ii) and c(i) kinks, and calculate the free energy of the binding of a Ca ion to the kink. Using metadynamics, we calculate the two-dimensional free energy surface, and use the nudged elastic band method to calculate the Minimum Free Energy Path (MFEP). The results are shown in Figure 5.3.

Perhaps the most noteworthy feature of the free energy surfaces is the existence of two distinct states, separated by their z-distance to the kink site (i.e. z-distance from the upper terrace). Figure 5.3 shows a snapshot of the two configurations for both kinks. In one configuration, the Ca ion resides in the kink lattice site; in the other, it resides above the kink site, forming two interactions with the upper terrace. The two configurations are therefore labelled as 'lattice' (denoted in green) and 'bidentate' (denoted in yellow).

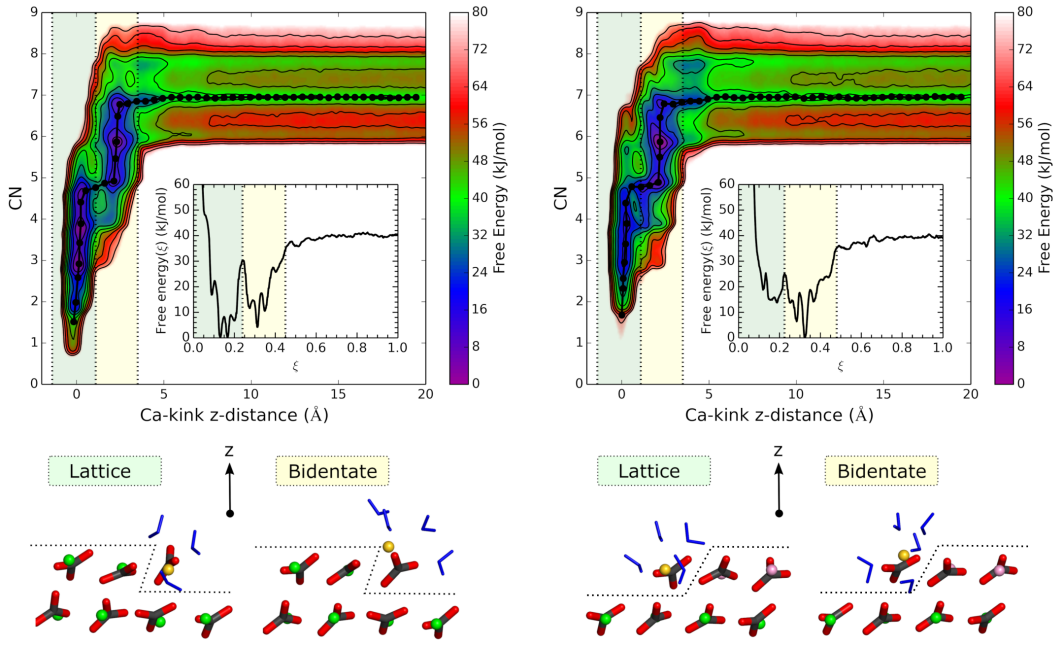

Figure 5.3: Ca-kink interaction free energy as a function of the Ca-kink z-distance and Ca-O_w coordination number for the a(ii) kink (left) and the c(i) kink (right). A value of zero on the x-axis corresponds to the adsorbate residing on the same z-plane as the upper terrace. The dotted beads show the MFEP for Ca attachment to the kink site. The inset graph shows the value of the MFEP as a function of the dimensionless parameter that maps it. The figure also contains side views of cross-sections of calcite along the step, demonstrating an example of lattice and bidentate configurations for both kink types. The outline of the steps on which kinks propagate are traced with dashed lines. Ca ions are shown in green, C in grey and O in red. Water molecules are shown in blue and the inserting Ca ion is shown in gold.

The regions on the free energy surfaces corresponding to each configuration are also highlighted.

Considering the free energy barrier separating the two states, it appears that the largest free energy barrier separating the two states is one presented by the translation of the adsorbate, rather that its hydration. The barriers arising from dehydration are typically small, indicating that including CN as a reaction coordinate may not be required for producing an accurate free energy surface. In order to test this, we repeated the above simulations without the explicit dehydration of the Ca adsorbate, and compared the one-dimensional free energy surface with the two-dimensional free energy surfaces (Figure 5.3) with the hydration component integrated out using equation 3.90. The results are shown in the appendix in Figure A2. It is found that both sets of simulations produce very similar free energy surfaces subject to minor differences which can be attributed to sampling errors.

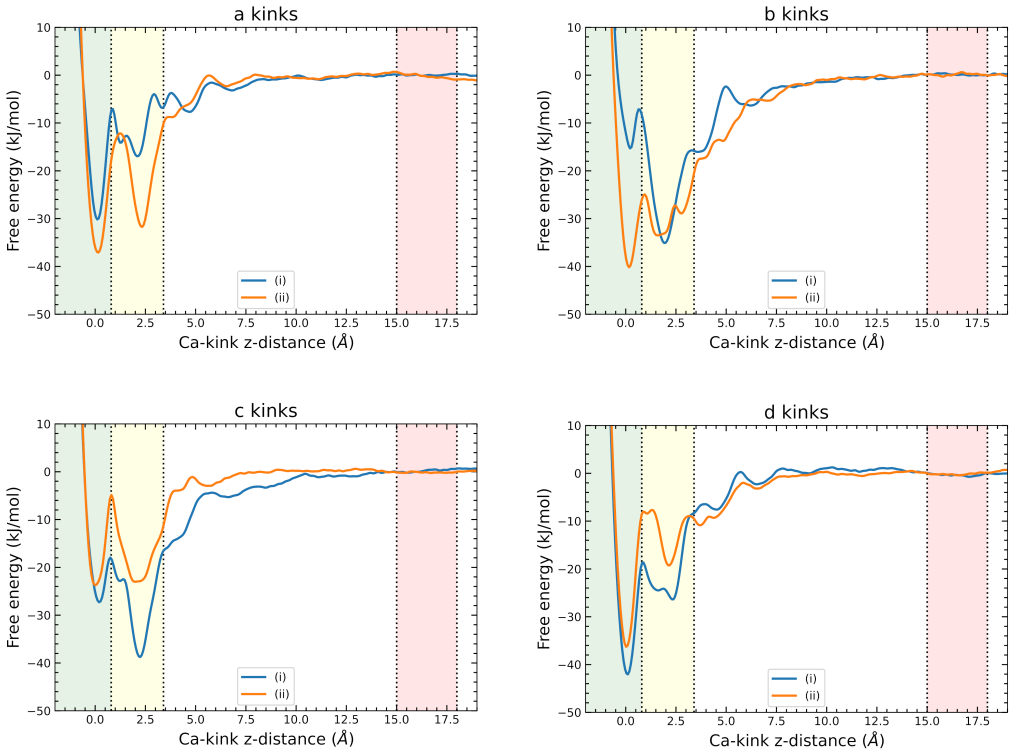

Figure 5.4: Plots of the free energy of all inserting Ca ions as a function of their distance from the kink site. The two plots for each kink type correspond to the two orientations of surface CO_3 ions. The highlighted green and yellow regions in the plots correspond to the lattice and bidentate configurations as per Figure 5.3.

5.5 Ca adsorption

We have demonstrated above that dehydration does not provide a significant barrier for the adsorption of Ca ions to two CO_3-terminated kink sites. Since the removal of a reaction coordinate allows convergence for minimal computational expense, we proceed to calculate the free energy as a function of the Ca-kink z-distance for all corresponding kink sites. These free energies are shown in Figure 5.4. The free energy profiles show that the position of the thermodynamic minimum depends on the kink type; some kinks prefer the lattice configuration while others prefer the bidentate configuration (Figure 5.3). Ca ions prefer to adsorb to the lattice configuration in 6 of the 8 cases. For a and d kink types, all Ca ions have a thermodynamic minimum at the lattice site. For b and c kink types, we find a greater variation in the free energy landscapes, where some prefer the bidentate configuration. For more information on the thermodynamic minima, Figure A3 in the appendix shows the binding configurations corresponding to the thermodynamic minimum for each Ca kink.

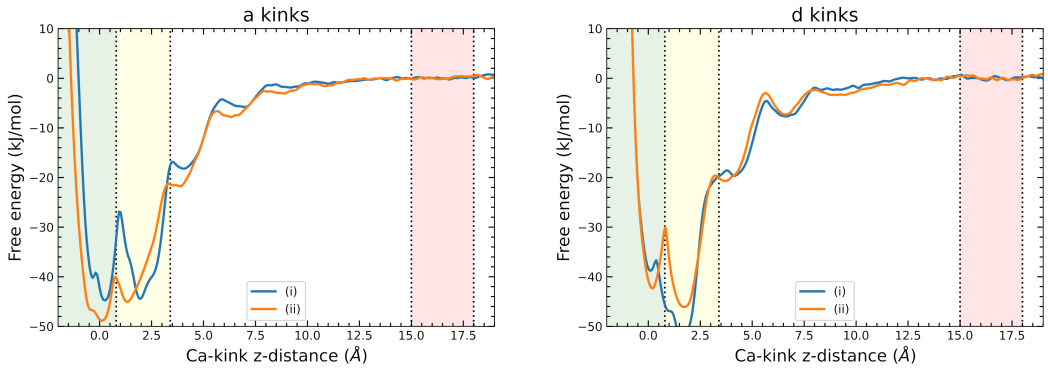

Figure 5.5: Free energy of all inserting CO_3 ions studied as a function of distance from the upper terrace. The two plots for each kink type correspond to the two orientations of surface CO_3 ions. The highlighted regions in the plots correspond to the lattice and bidentate configurations as also shown in figure 5.4.

5.6 CO_3 adsorption

As discussed in Section 5.2, calculating CO_3 adsorption free energies incurs the additional difficulty of dehydrating the kink site. Therefore, two reaction coordinates were used for a CO_3 ion adsorbing to a Ca-terminated kink: the z-coordinate of the CO_3 ion, and, to drive dehydration of the kink, the distance between the kink-terminating Ca ion and its nearest water molecule given by equation 5.1. The free energy surfaces for a CO_3 ion adsorbing to four different Ca-terminated kinks are shown in Figure 5.5, where equation 5.1 is integrated out using equation 3.90. The complete free energy landscapes are shown in Figure A4. Unlike Ca kinks, of which all 8 were studied, we have only shown the results for 4 kinks. This is because any attempts to study b or c CO_3 kinks resulted in a water molecule becoming trapped under the kink-terminating Ca ion during the simulation. In this situation, the Ca ion would otherwise transition to its bidentate configuration. However, due to the harmonic tethering of the Ca ion, it was unable to do so. The result was that the simulation configurations became unstable, and metadynamics simulations ran into convergence issues. This could be solved by applying a dual adsorption method such as the one discussed in Section 5.8. However, these simulations require a far longer convergence time ($\sim$3 μs, see Section 5.8), and running multiple simulations of similar lengths is beyond the scope of this chapter.

Local free energy minima in Figure 5.5 correspond to both lattice and bidentate configurations. The lattice configuration requires the full dehydration of the Ca-terminated kink site to which the CO_3 ion binds, while the bidentate configuration does not. Significantly, half of the CO_3-terminated kink ions prefer to adopt the bidentate configuration. Only the CO_3-terminated a-kinks have a preference for the lattice configuration. By contrast, Ca ions mostly preferred the lattice configuration. This difference is likely explained

by the water molecules at Ca-terminated kinks, the residence times of which are likely to be far larger than those of lone Ca ions [53]. The removal of water at kink sites therefore comes at a larger free energy cost than the removal of water at lone ions. The most stable bound configurations of CO_3 ions are shown in Figure A5, along with the water molecules found near the kink site.

5.7 Adsorption free energies

The results shown in figures 5.4 and 5.5 are used to calculate adsorption free energies using the method outlined in Section 4.6. The Boltzmann exponent of the lattice and bidentate values of the free energy (green and yellow) are integrated over, and the simulation free energy difference is calculated by averaging the free energy over the red regions to account for statistical noise. It can be seen that the free energy surfaces are flat in the red regions, indicating that the solute is fully dissolved. The calculated simulation and adsorption free energies are shown in table 5.1. For the Ca ions adsorbed to CO_3-terminated kinks, we have also included the average number of water molecules coordinated with the Ca ion in its most stable bound configuration, $\langle N_c \rangle$. We again assume that dehydration of the Ca ion is sufficiently rapid in order for us to collect a sufficient sample of N_c, which we define as:

$$N_c = \sum_j \chi_j \qquad \chi_j = \begin{cases} 1, & r_j \leq r_0 \\ 0, & r_j > r_0 \end{cases} \tag{5.5}$$

where r_j is the distance to water oxygen j and $r_0 = 3$ $\mathring{A}$ is a cut-off distance. The value of 3 Å is chosen from Radial Distribution Functions calculated by Raiteri et al. [45]. N_c was calculated by post-processing the trajectories.

The ΔG_{ads} values shown in table 5.1 vary between -14.7 and -30.8 kJ/mol. The variation of these numbers is unsurprising, since similar calculations for step sites show significant variation in binding free energies [43]. Nevertheless, it is clear that different Ca-terminated kink sites have different stabilities. One notable point here is that models such as the Kossel model and solid-on-solid models discussed in Chapter 2 do not seem to apply to calcite, as they assume that the adsorption free energy is independent of kink type.

The values of $\langle N_c \rangle$ are typically dependent on the preferred adopted configuration of the kink site. Where the lattice configuration is preferred, $\langle N_c \rangle$ corresponds to roughly 3, implying that a total of three coordinated water molecules is the most stable configuration (see table 1). Where the bidentate configuration is preferred, $\langle N_c \rangle$ is typically about 4.5, implying that the number of coordinated water molecules fluctuates between 4 and 5. This can be viewed in Figure A3.

This presents an issue when considering the relative abundance of Ca and

Ion	Kink	ΔG_{sim} (kJ/mol)	ΔG_{ads} (kJ/mol)	$\langle N_c \rangle$
Ca	a(i)	-30	-20	3.3
	a(ii)	-36	-26	3.2
	b(i)†	-35	-25	4.6
	b(ii)	-31	-20	3.2
	c(i)†	-39	-29	4.5
	c(ii)	-24	-15	3.1
	d(i)	-42	-31	2.8
	d(ii)	-36	-26	2.7
CO_3	a(i)†	-46	-32	-
	a(ii)	-49	-36	-
	d(i)†	-52	-38	-
	d(ii)†	-46	-33	-

Table 5.1: Simulation free energy differences (ΔG_{sim}) and adsorption free energies (ΔG_{ads}) for Ca and CO_3 ions adsorbing to various kink types. $\langle N_c \rangle$ is the average water coordination number of the Ca adsorbate in its most stable configuration. † The ion adsorbs preferentially to the bidentate configuration.

CO_3-terminated kink sites. Previous studies have determined relationships between the solution stoichiometry ([Ca]/[CO_3]) and step velocities for both acute and obtuse steps [58, 60, 41, 123]. Although studies typically find small offsets, step velocities tend to peak at [Ca]/[CO_3] ≈ 1. This implies that Ca and CO_3 ions have an approximately equal probability of terminating a kink site, a least to an order of magnitude. However, the variation in adsorption free energies we find between Ca and CO_3 indicate that CO_3 has a significantly higher probability of terminating a kink site than Ca, and therefore the step velocities should peak at [Ca]/[CO_3] $\gg 1$. It is clear, therefore, that the adsorption free energies we find are inconsistent with experiment. There are several possible explanations for this. For example, we neglect the existence of bicarbonate ions, something which could limit the number of available ions or kink sites. We do not examine b or c type kinks in as much detail. Neither do we consider the attachment of ion pairs and complexes, the presence of which has been proposed to induce an effect on crystal growth [50]. However, we do not expect any of the above explanations to produce a significant effect on the step velocity/stoichiometry relationship. A more likely explanation is that CO_3 ions primarily facilitate kink nucleation as well as kink propagation, something already concluded from simulations [43]. A lower [Ca]/[CO_3] will therefore increase the kink density, thus producing an increase in step velocity. This effect therefore shifts the peak of the step velocity toward [Ca]/[CO_3] ≈ 1. This explanation does require some degree of fine tuning, although it is worth reiterating that step velocities do not peak exactly at [Ca]/[CO_3] $= 1$ for either acute or obtuse steps.

5.8 Example of a multi-step kink growth mechanism

It is apparent from figures 5.4 and 5.5 that calcite kink propagation cannot be always assumed to involve a step-by-step adsorption of ions to the surface. Moreover, it appears that the the adsorption of ions into kink lattice sites must involve a more complex process where ions are stabilised in their lattice configuration via the adsorption of additional solutes. An example of this is the d(i) CO_3 kink, the lattice configuration of which is unstable. In this instance, the adsorption of another Ca ion (at the very least) is required to stabilise the CO_3 ion at the d(i) lattice site. There remains a possibility that such a process may require the adsorption of multiple ions, and that a chain of bidentate-like ions may exist along the step. To resolve this, we calculate the free energy surface for the position of two ions. We choose the d(i) CO_3 kink for its preference for the bidentate configuration (Figure 5.5), as well as the Ca ion that forms the subsequent d(i) Ca kink. The CO_3 ion is constrained at a maximum of 4 Å above the elevated step. This ensures the ion only transitions between its bidentate and lattice configuration, as well as reducing the computational expense of the simulation. The Ca ion is constrained at a maximum of 15 Å above the terrace. Including the hydration of the adjacent Ca-terminated kink site (d(ii)), there are a total of three reaction coordinates used in this simulation. The simulation required $\sim$3 μs to converge. Figure 5.6 shows the free energy landscape as a function of the Ca and CO_3 positions. The hydration of the d(ii) Ca kink site is integrated out using equation 3.90.

There are four distinct steps to the growth process, labelled A-D in Figure 5.6. First, the CO_3 ion adsorbs to the kink in the bidentate configuration (A). The Ca ion then adsorbs to the bidentate CO_3 ion by sitting approximately 5 Å above the step (B). The CO_3 ion transitions to the lattice site, pulling the Ca ion into a bidentate configuration (C). This comes at a free energy cost. Finally, the Ca ion transitions into its lattice configuration (D). This completes the process of adsorption, and it is found that D is the most stable configuration. This result is significant as it demonstrates that even the CO_3 ion with the least stable lattice configuration is stabilised at the lattice configuration through the insertion of one additional ion. We also note that two of the four kink types studied have a preference for the bidentate configuration, while all a and d Ca kinks prefer to adopt their lattice configuration. We therefore expect a similar multistep process to occur for other a and d CO_3 kink types. It is worth stressing, however, that the mechanism demonstrated here is only an example of a multistep kink growth mechanism, and that we are not assuming that this result will carry over to other kink types which reside in a bidentate configuration. Nevertheless, the results shown in figures 5.4 and 5.5 demonstrate that several terminating ions (4 of the 12 studied) must require one (or more) additional ions to adsorb before a

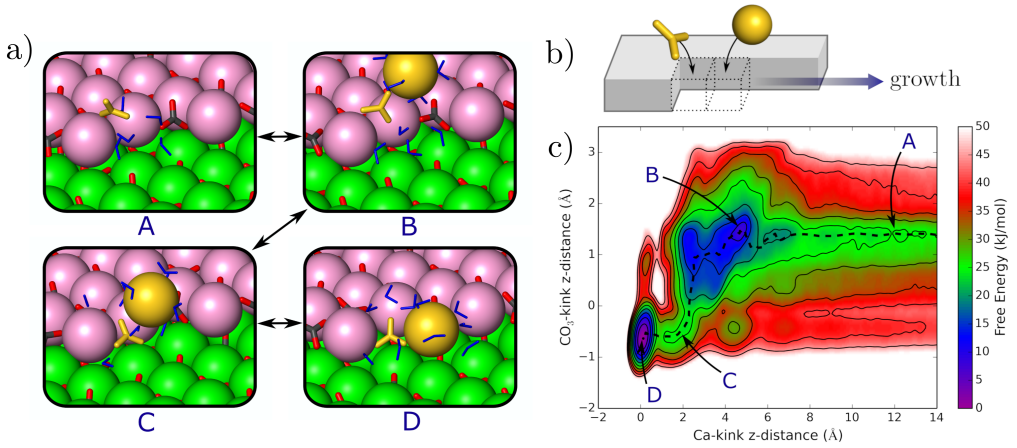

Figure 5.6: a) Four snapshots (A-D) illustrate the multistep growth mechanism. Here, Ca atoms in the upper terrace are shown in pink. The two terminating ions are shown in gold. The perspective of the snapshots is one which directly faces the step, which runs horizontally. The kinks grow from the left side. b) Schematic depicting the perspective of the snapshots and the direction of growth of the kink. c) Free energy as a function of the position of the CO_3 and Ca ions adsorbing to the d(i) kink. A third reaction coordinate that accounts for dehydration has been integrated out. The minimum free energy pathway is traced with a dashed black line.

full transition to the lattice site can take place. Ideally, all kink types which prefer to sit in their bidentate configuration should be studied. However, the free energy plot shown in Figure 5.6 took a total of $\sim$3 μs to convergence. Repeating this process for 5 kink types would require multiple simulations over very large time-scales, and is therefore beyond the scope of this chapter.

5.9 Role of cation dehydration in inhibiting calcite kink growth

Cation dehydration is generally believed to limit the rate of ionic crystal growth [47, 126, 50, 48], although recent evidence suggests this may not be true for the growth of calcium minerals [51]. For adsorption into the bidentate configuration, our simulation results broadly support this new perspective: we find that, for all of the kinks that we have sampled, the ions must overcome only a $\sim 1\, k_B T$ barrier to transition from solution to the the bidentate configuration. The solutes will therefore initially adsorb to kinks at a rate determined by diffusion rather than by a reaction barrier, rather than at a rate determined by the dehydration of any cations. For some kinks, there exists a substantial barrier separating bidentate and lattice configurations. However, since the barrier from bidentate to solution is generally larger than the barrier from bidentate to lattice, the adsorbate is effectively captured by

the kink site as soon as it reaches the bidentate configuration.

5.10 Conclusions

Many of the ions that terminate calcite kinks have a tendency to reside in a bidentate configuration, rather than fit directly into the lattice. They sit above the kink, binding to two ions and causing minimal displacement of water molecules. The integration of these particular ions into the kink lattice site requires the adsorption of an additional ion, and so calcite kinks do not generally grow via a sequence of independent adsorption events as assumed in classical models. This multistep kink propagation process is analogous to what is observed for kink nucleation, in which solutes initially adsorb to the upper terrace before the adsorption of another ion.

The results presented in this chapter provide a framework for the study of the binding of impurities to calcite kinks using molecular simulation, something which has been largely unexplored. Future molecular simulation studies of impurities adsorbing to kinks must therefore take into account whether the kink to which an impurity binds resides in its lattice or bidentate configuration. We will encounter this effect again in the following chapter, which concerns the interaction between amino acids and polyamines with calcite kinks.

6

INTERACTION AND GROWTH-INHIBITING EFFECTS OF IMPURITIES

In the previous chapter, we examined the stable binding configurations and adsorption free energies of Ca and CO_3 units forming kink sites. This allows us to proceed with examining the adsorption of impurities to calcite kink sites. To recap from Chapter 1, computational studies of impurities such as amino acids interacting with calcite lattice sites are infrequently found, especially for kink sites. As a result, little is known about the molecular mechanisms through which amino acids interact with calcite. The relative roles of carboxyl and amine groups are an example of this, as previous computational studies have provided conflicting evidence for which is the primary driver of adsorption. Such simulations have been limited by their poor representation of the kink sites to which the adsorption of impurities is most relevant. The study of impurities interacting with kink sites has been approached with hesitation, due to the large number of kink sites, and a lack of understanding of their equilibrium configurations. Since these issues were discussed in Chapter 5, we now have the tools to address the issue of how amino acids adsorb to calcite. In this chapter, we aim to do so through the calculation of adsorption free energies of amino acids on calcite, as well as determining their binding configurations.

This chapter begins by examining the free energies of binding of amino groups to terrace, step and kink sites. The work presented throughout the first half of this chapter is designed to complement experimental work which found that polyamines had a surprisingly high propensity to occlude in calcite. The second part of this chapter concerns the adsorption free energies and binding configurations for aspartate (asp) interacting with calcite kink sites. The purpose of this half of the chapter is to determine the nature of carboxyl groups' interaction with calcite, and to directly replicate experimental procedures used to calculate the adsorption free energy of asp. [71, 72]

6.1 Binding of amines to calcite lattice sites

This section concerns a case study in which amine-rich molecules were found to have a large propensity to occlude in calcite [87]. It was initially found that polyamines such as putrescine, cadaverine, spermine and spermidine were able to occlude in calcite with a surprisingly high occlusion efficiency. This alone was a surprising result, as it was previously believed that positively charged impurities have a minimal interaction with calcite. An example of this is amino acids, which are generally believed to adsorb to calcite through interactions between negatively charged functional groups and Ca

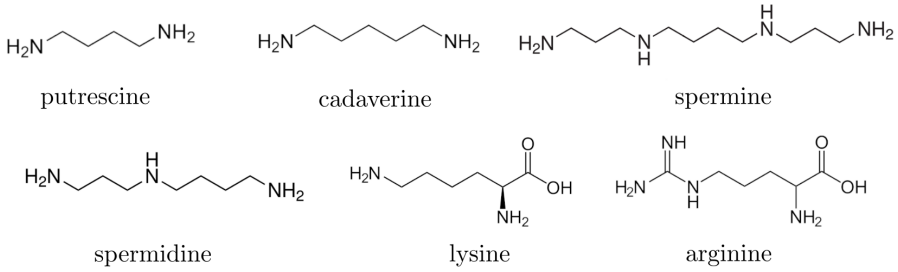

Figure 6.1: Amine-rich molecules studies experimentally and modelled in this section by isolating the primary and secondary amine groups. All molecules studied are observed to occlude in calcite, indicating a significant interaction between amine groups and calcite.

cations. The discovery of positively charged polyamines occluding in calcite prompted the further study of two amino acids: lysine (lys) and arginine (arg). Lys was found to occlude with an efficiency roughly five times that of arg. Additionally, lys was found to produce an elongation of the *c*-axis of calcite, indicating a preferential binding to the acute steps over the obtuse [80, 127]. Conversely, arg was found to induce no morphology change. The purpose of this section is to investigate the binding of polyamines, as well as lys and arg to calcite terrace, step and kink sites in order to gain insight into how amine groups contribute to the adsorption of amino acids to calcite.

Rather than model all amino acids and polyamines studied in their entirety, we opt for the more tractable approach of isolating the relevant side chains, thus requiring fewer reaction coordinates. In this chapter, we study the interaction of three molecules with calcite terrace, step and kink sites: primary amines, secondary amines and the arginine side chain. All molecules are modelled in accordance with the protonation state of lysine (lys) and arginine (arg) at experimental conditions (pKa side chain of Lys = 10.6 and Arg = 12). The primary amine we model is ethylammonium, $C_2H_8N^+$, as it may represent the side chain of ethylammonium or the primary amines in putrescine, cadaverine, spermine and spermidine. Ethylammonium is modelled in its protonated state throughout under the assumption that even molecules with a lower pKa number will still favour binding through protonated functional groups over binding through more prevalent, but non-charged non-protonated groups. The secondary amine modelled is diethylammonium, $C_4H_{12}N^+$, which can represent the secondary amines in spermine and spermidine. The secondary amine is modelled in its protonated state for the same reason as for the primary amine. Finally, the side chain of arg is represented as $C_2H_{10}N_3^+$, again protonated to correspond with experimental pH. In this section, the force fields describing all impurities are obtained from the General Amber Force Fields. [128] The partial charges for each atom were obtained using Antechamber. [129]

In this section, binding free energies are calculated via a two-step process.

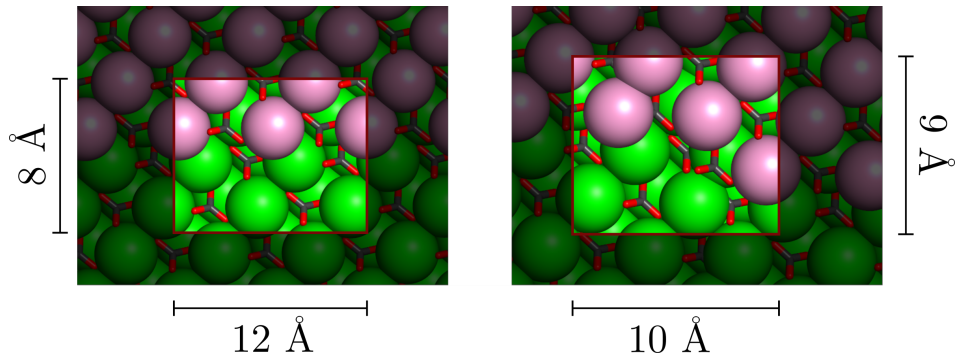

Figure 6.2: View from above simulation cell for the step (left) and kink (right). The shaded regions depict where constraints have been imposed on x- and y-positions of the impurity.

The first step involves calculating a three-dimensional free energy surface for all position coordinates of the amine groups. Constraints are added to this simulation, as per Figure 6.2, in order to reduce the volume of space sampled in the simulation. The constraints exist in the form of a set of harmonic walls with a spring constant of 100 kJ/mol/Å. The three-dimensional free energy function is merely approximated in this process, as deriving the binding free energy would require long simulation times. Instead, the position of the thermodynamic minimum is identified from the simulation, and an additional metadynamics simulation is run. In this simulation, the x- and y-components of the reaction coordinate are constrained at the thermodynamic minimum using a harmonic spring with a spring constant of 100 kJ/mol/Å. In this simulation, the free energy was derived as a function of z-component of the position. From this simulation, the binding free energy is obtained. For ethylammonium and diethylammonium, the position of the nitrogen atom is used as the reaction coordinate, for the arginine side chain, the centre-of-mass of the three nitrogen atoms are used.

It is worth noting that we do not calculate adsorption free energies in this section. This is partly because we are only interested in comparing binding free energies of amine groups with one another, as no experimental derivation of adsorption free energies with which we can compare our results exist. Furthermore, the AMBER force fields have not been explicitly fitted to thermodynamic data, and may produce unrealistic adsorption free energies. Nevertheless, since we are only comparing the interactions of different amine groups with respect to one another, we believe the binding free energies we calculate will be sufficiently realistic at least to a qualitative degree.

6.1.1 Binding energy calculations

In this section, we present every amine binding free energy calculated in this chapter. In Chapter 5 we found that, of all the kink sites studied, the a(ii) CO_3 kink was the only CO_3-terminated kink where the CO_3 has a significant

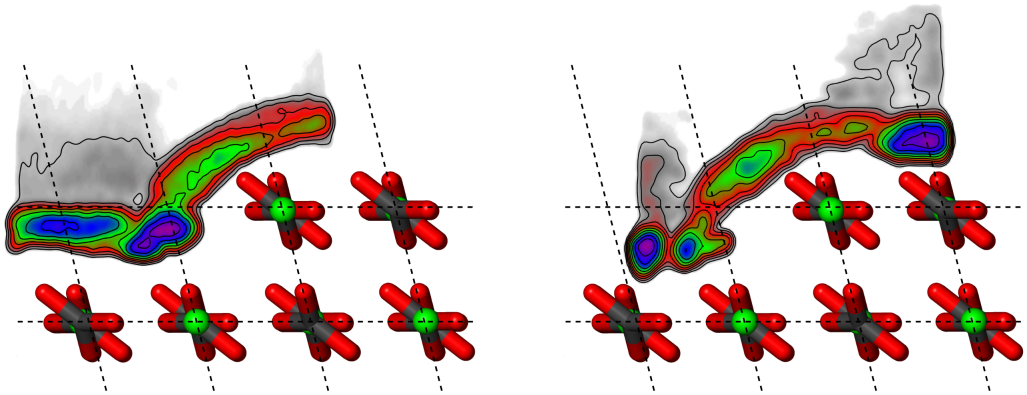

Figure 6.3: Three-dimensional free energy surfaces for the binding of ethylammonium (left) and diethylammonium (right) to the acute step. The free energy surfaces are projected onto two dimensions using equation 3.90. While ethylammonium adsorbs to the step with a stronger interaction than with the terrace, the reverse is true for diethylammonium.

preference to reside in its lattice site. For this reason, this site was chosen for all simulations involving acute kink sites to avoid complication. Similarly, when considering the obtuse kink, we choose the d(ii) CO_3 kink for its relative stability in the lattice site configuration compared with the d(i) CO_3 kink. It is also worth reiterating that, in Chapter 5, we calculated a small residence time for water interacting with the a(ii) CO_3 kink, and concluded that forced dehydration of CO_3-terminated kink sites is not necessary.

The simulation free energy differences calculated are shown in Table 6.1. For ethylammonium, the simulation free energy differences are calculated for the terrace, as well as both acute and obtuse steps and kinks. For the arg side chain, only the acute step and kinks are considered for reasons discussed in Section 6.1.2.

The three-dimensional metadynamics simulations revealed that the secondary amine has only a weak binding to all steps and kinks. Figure 6.3 compares the three-dimensional free energy surface projected onto two-dimensions using equation 3.90. The free energy for diethylammonium shows that binding to terrace sites is stronger than binding to step sites. Any further calculations of free energy surfaces for different step or kink sites revealed the same picture. For this reason, we do not calculate the corresponding binding free energies of the arg side chain to step and kink to the higher level of precision that we do for the terrace, and binding free energies are not reported. Since binding is not observed for the arg side chains to any step or kink sites, it appears that the most likely role of secondary amines is to stabilise the binding of polyamines such as spermine and spermidine, which will bind primarily to kink sites.

Amine	ΔG_{sim} (kJ/mol)				
	terrace	step (a)	kink (a)	step (o)	kink (o)
ethylammonium	-23.0	-41.2	-75.4	-43.4	-89.1
arginine	-15.2	-21.7	-26.3	-	-
diethylammonium	-32.1	-	-	-	-

Table 6.1: Binding free energies for ethylammonium, diethylammonium and the arg side chain to calcite terrace, step and kink sites.

6.1.2 Lysine and arginine

The primary (-NH_3) amine in the lys side chain interacts strongly with the carbonate anions in all adsorption sites; the adsorption free energy increases roughly in proportion to the number of carbonate anions available—one for a terrace, two for a step and three for a kink (Figure 2b). This can be understood using Figure 6.4, which shows lys forming a single bond with a CO_3 ion on the terrace, two on the step and three in the kink. These results are more reminiscent of the solid-on-solid dissolution model described in Chapter 2. In contrast to the lys side chain, the binding free energy of the arg side chain was found to be far weaker at all adsorption sites, especially the step and kink sites. The stereochemistry and charge density of the single amine in the lys side chain therefore likely facilitates stronger binding to calcite than the three amines in the arg side chain.

As discussed above, it was observed experimentally that lys had an occlusion efficiency roughly five times that of arg. Our results show a striking difference in binding free energy. While the binding of the lys side chain is stronger as expected, the difference in binding energy is far too large to produce a factor of five in relative occlusion efficiencies. Instead, it is more likely that arg does not bind through its side chain. Moreover, arg most likely binds through its alpha carboxyl and amino groups. Lys, on the other hand, most likely binds through its side chain amino group. The side chain most likely produces a slightly stronger binding to the surface than the amino groups. This explains why lys has a greater occlusion efficiency than arg. The fact that the mechanisms of adsorption are different may also allude as to why lys produces a morphology change, whereas Arg doesn't. However, no clear mechanism for why this occurs exists.

6.1.3 Acute and obtuse

Since the side chain is found to not be the primary driver of adsorption of arg, we exclude it from calculations of ΔG_{sim} to obtuse sites. We also did not observe any significant binding of diethylammonium to the obtuse steps or kinks. We therefore only present ΔG_{sim} for the obtuse steps for ethylammonium. The binding free energies for ethylammonium show a significantly

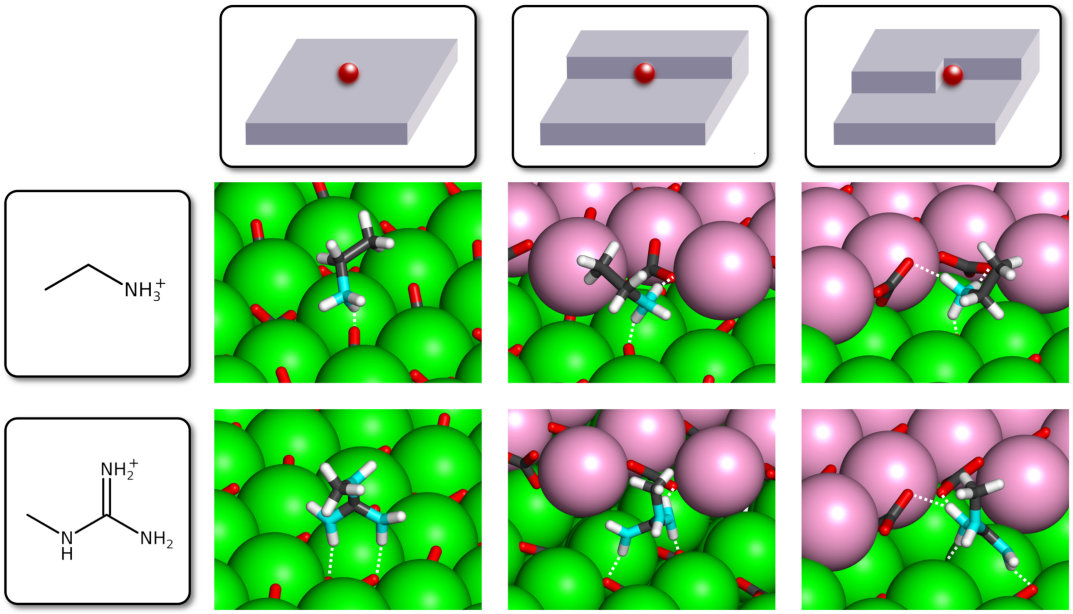

Figure 6.4: Snapshots from trajectories showing binding of the lys (ethylammonium) and arg side chains to a calcite terrace, step, and kink site. The white dashed lines indicate hydrogen bonds. Water is excluded from all images for clarity. Lys forms one bond with the terrace, two with the step and three with the kink. Args adopt a more complex set of configurations in which several hydrogen atoms interact with oxygen in carbonate ions.

stronger binding to the obtuse step and kink. This is not in agreement with experimental findings, which observe a lengthening along the c-axis, indicative of a stronger binding to acute sites. One explanation for this may be that obtuse kink CO_3 ions to not reside in their lattice sites, as found in Chapter 5. The binding of ethylammonium to the kink sites will therefore incur a free energy cost of driving the CO_3 ion into its lattice site. However, this free energy cost of about 5 kJ/mol (d(ii) CO_3, see Figure 5.5) is not sufficient to account for the difference, and ethylammonium will still have a stronger total binding free energy to the obtuse kink. The difference in binding free energy between molecules attaching to acute and obtuse steps is a limitation of the work carried out throughout this chapter. The elongation along the c-axis has often been thought to be due to a low binding to obtuse sites, although no evidence for this has been provided throughout this section or, indeed, anywhere else so far. This will be revisited in the following section.

6.2 Aspartate adsorption free energies and binding configurations

To recap from Chapter 2, asp and its effect on calcite growth has been the focus of extensive experimental study, although the molecular details behind

such effects remains in contention due to a lack of understanding of the kink sites to which asp primarily adsorbs. The objective of this section is to reproduce these adsorption free energies using computational methods, and use them to examine the configurations of asp adsorbing to calcite. Particularly, we are interested in interactions with calcite kink sites. Interactions with steps or the terrace are likely to be much weaker, and therefore have a negligible effect on calcite growth. We therefore only consider interactions with kink sites in this section. This is a complex task which must be split into several sections. We begin by exploring potential binding configurations by forcing different configurations and equilibrating to determine which configurations are stable. We subsequently proceed to calculate the binding free energy of aspartate for a single kink site using metadynamics. The result from this section allows us to simplify the calculation such that binding free energies can be calculated for all necessary kink sites. These binding free energies are used to calculate the corresponding adsorption free energies. Finally, the calculated adsorption free energies are used to parametrise a Kinetic Monte Carlo (KMC) scheme, with which the inhibition of step velocities are calculated. This is used to derive a macroscopic adsorption free energy for both step velocities and for the entire crystal, which are compared with experimental values.

In the previous section, we modelled the intermolecular and intra-molecular interaction of amine groups using the AMBER force fields. However, we do not opt for these force fields in this section. This is because preliminary simulations with asp described using the AMBER force fields found that carboxyl groups had an extremely weak interaction with Ca atoms including kink sites. We suspect this to be an inaccurate representation of reality, since the carboxyl-rich asp has a higher occlusion rate into calcite than other amino acids, including the amine-rich Lys and Arg. Instead, the asp force fields of Raiteri et al. [130] are explicitly fitted to solvation free energies, similar to their calcite force fields [45], thus making them an ideal candidate for calculating thermodynamic quantities such as adsorption free energies. We therefore use the force fields of Raiteri et al. throughout this section. The simulation cell set-ups are identical to those used in Chapter 5 with a terminating CO_3 ion swapped for an asp molecule. The set-up is described in Section 4.2 and depicted in Figure A1.

6.2.1 Exploring stable configurations

Asp consists of two carboxyl groups and one amino group. The existence of these three functional groups creates a large array of potential combinations of interactions between functional groups and calcite lattice sites. Since water residence times for water molecules coordinated with calcium ions near the step are typically large, [53] it is possible that other calcium ions may require dehydration in order to calculate binding free energies. It is therefore important that all possible configurations are tested for their stability before

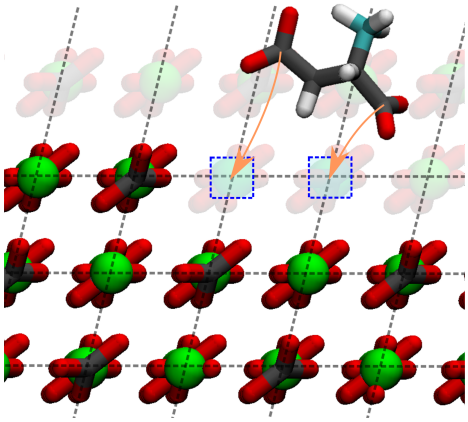

Figure 6.5: Diagram of the method used to force the aspartate into different initial configurations. The two carboxyl carbon atoms are forced into cubic boxes with a volume of 1 $Å^3$ (depicted in blue). The positions of the blue regions are given by the positions an imaginary Ca or CO_3 ion would occupy in a bulk crystal. This exact example of aspartate bound to a step site is not considered in this study, it is shown purely because it makes for an easier visualisation than a kink site.

proceeding to calculate binding free energies.

Here, we consider every possible binding configuration in which asp interacts with a kink site, i.e. every permutation of functional group and lattice site local to the kink. We determine whether configurations are stable using a process outlined as follows. Each relevant functional group was forced into either the kink site or a nearby lattice site to the kink sites. This was done by starting a simulation with the calcite frozen in its lowest-energy bulk configuration and placed in a vacuum. The additive was dragged into the chosen configurations via its functional groups (carbon for carboxyl groups and nitrogen for amine groups) using a constant force with magnitude 50 kJ/mol acting on chosen functional groups. The negatively charged carboxyl groups will bind to Ca sites. Therefore, the starting positions of the carbonate ions were chosen to be the exact position of the lattice sites in which a CO_3 ion would occupy in a bulk crystal. Since the positions of the carboxyl groups will not correspond exactly to these lattice sites, a window of one cubic Å was allowed for the functional groups to move freely. Likewise, the positively charged amine group will interact with CO_3 ions, so the same method was applied in this instance. Figure 6.5 shows an example of the method used to drag the aspartate functional groups into various configurations. It is worth noting that Figure 6.5 shows a configuration (i.e. aspartate bound to a step and terrace site) not explored in this study. This configuration is shown purely because it makes for an easier visualisation than a kink site.

Once the initial configuration was set up, water was added to the simulation and equilibrated for 100 ps with the calcite still frozen. The calcite was unfrozen and a further equilibration of 50 ps was carried out. Finally,

the dragging forces holding the functional groups in place were removed, and a simulation lasting 10 ns was carried out. If a functional group drifted away from its initial configuration within a few 100s of ps and did not return to its initial configuration throughout the simulation, the configuration was considered unstable. Should the functional group remain in its initial configuration throughout the entire simulation, or for a suitably long time, it was considered stable (or metastable) and the use of additional reaction coordinates may be required.

We began by considering binding to the a(ii) CO_3 kink site, the same as in Section 6.1. However, before even exploring configurations, it was found that the interaction between the amino group and the CO_3-terminated kink site was so weak that the amino group diffused out of the kink site within a few ps of removing the tethers. This result alone is a striking contrast to what was observed in Section 6.1. Rather than exhibiting a strong binding to the kink site, there is no binding at all. This will be discussed further in Section 6.2.9. For the time being, we proceed under the assumption that the binding of amino groups is negligible for asp.

Assuming that binding of the amino group can be discounted, only configurations involving the binding of carboxyl groups to Ca ions required testing. We assume that interactions with CO_3-terminated sites will be relatively small, and therefore only consider Ca-terminated kink sites from here on. In Chapter 5, the a(ii) CO_3 kink was found to be the most stable in its lattice configuration. We assume that a carboxyl group will behave similarly to a CO_3 ion, and begin by considering interactions involving the a(i) Ca kink (to which a CO_3 ion will adsorb to form an a(ii) CO_3 kink). For the a(ii) kink, a total of six unique configurations were explored. Each configuration included a single carboxyl group occupying the kink lattice site. The other carboxyl group occupied either a site above the step, a lattice site along the step or a site on the terrace adjacent to the kink. The positions of the carboxyl groups were additionally switched, thus doubling the total number of configurations explored to six. Figure 6.6 shows snapshots of every configuration tested for its stability.

It was found that the only configurations to remain stable were those where either carboxyl group resided in the kink site. These bonds were found to remain throughout the course of the simulations. Other bonds between carboxyl groups and lattice sites were broken between 100 ps and 2 ns, and were not re-formed throughout the simulation. It appears, therefore, that only interactions between both carboxyl groups and the kink site are stable (or at least metastable). We are therefore able to proceed to calculate binding free energies under the assumption that multi-site adsorption is not an issue.

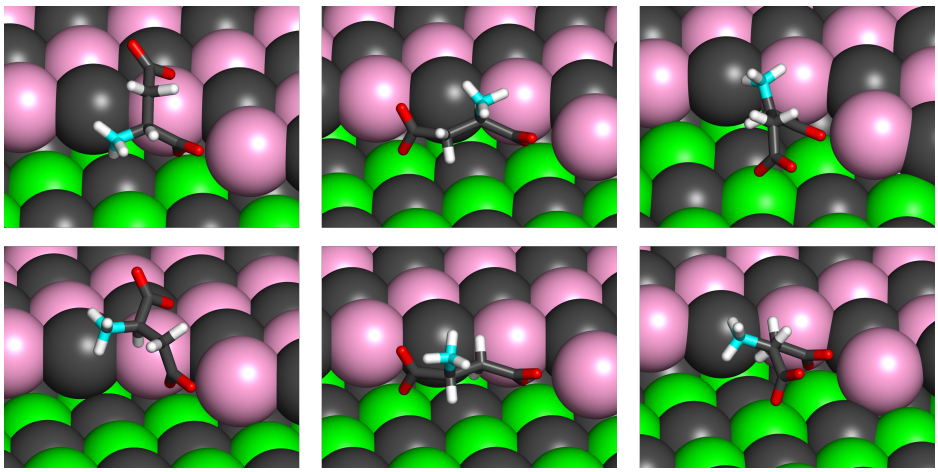

Figure 6.6: Configurations explicitly tested for their stability. Every configuration includes one carboxyl group occupying the kink site, an the other occupying a different lattice site.

6.2.2 Binding free energy of asp to the a(i) Ca kink

With the knowledge that only bonds between carboxyl groups and Ca-terminated kink sites are stable for the a(i) Ca kink, binding energies can be calculated using metadynamics involving the dehydration of the kink site using equation 5.1. We proceeded by calculating the binding free energy of asp interacting with the a(i) Ca kink site. Here, we assumed the kink must be dehydrated as is done in Chapter 5. We again adopted the method presented in Chapter 5 to dehydrate the kink. Here, we did not opt for the method outlined in Section 6.1, in which the thermodynamic minimum is located and a separate metadynamics simulation is run. Instead, we opted for the method akin to the one described in used in Chapter 5, in which the x- and y-positions of the centre of mass of the asp molecule were constrained to within 3 Å of the kink site. We do this because using this method makes it far easier to calculate adsorption free energies. The average z-position of the two carboxyl group C atoms were used as the reaction coordinate, and a metadynamics simulation was run with Gaussians deposited every 2 ps. All other parameters are detailed in Section 4.5. The resulting free energy function is a two-dimensional surface with a dependence on the average z-position of the two carboxyl groups (x-axis), as well as equation 5.1 (y-axis). This free energy surface is shown in Figure 6.7, as well as a series of snapshots depicting the asp configuration corresponding with the thermodynamic minimum.

The results of Figure 6.7 are rather striking. It can be observed that the thermodynamic minimum corresponds to a configuration in which the kink site remains hydrated, and the asp molecule adopts a bidentate configuration akin to those observed for CO_3 in Chapter 5, in which there is no interaction with the lower terrace. It appears that asp does not displace water when binding to this kink site. This finding is significant, as the oppo-

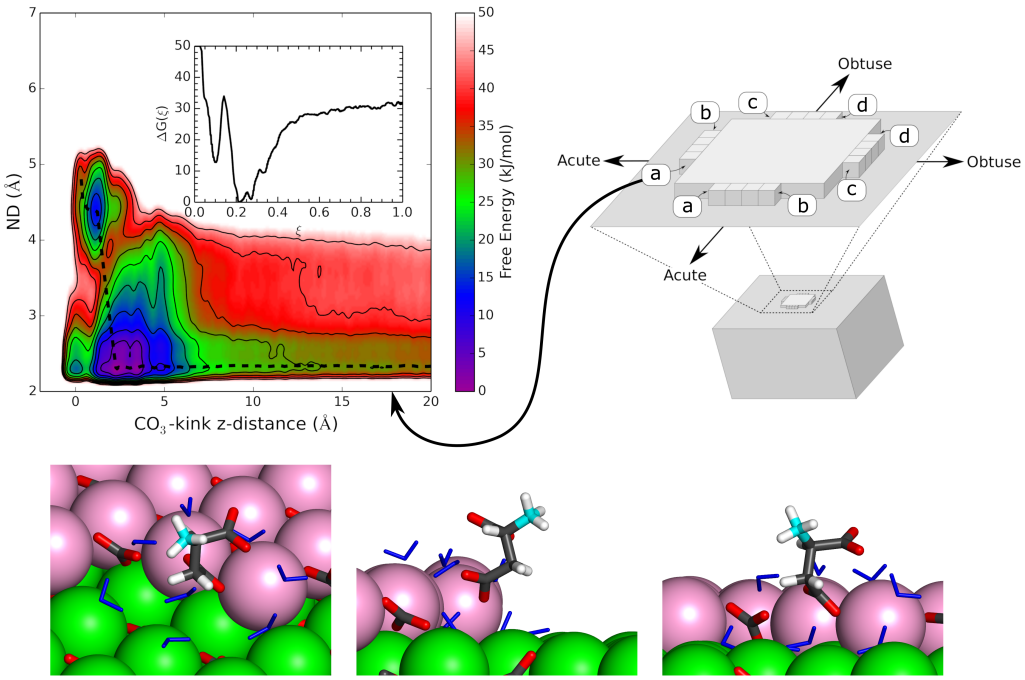

Figure 6.7: a) Two-dimensional free energy as a function of the average position of the two carboxyl group carbons with respect to the kink site (x-axis), and the nearest water molecule to the kink site (y-axis). b) One-dimensional free energy surface obtained by projecting the free energy surface in (a) onto one dimension. c) Snapshots from three perspectives show a bound state of asp with an a(ii) Ca kink site. The amino acid sits above a hydrated kink site and interacts directly with the terminating Ca. Here, a low value along the y-axis ($y \approx 2.3$) corresponds to a fully hydrated kink site. The well seen at the top-left of the graph corresponds to a fully dehydrated kink site occupied by a functional group.

site has previously been assumed true. Our findings are in some contention with previous hypotheses that impurities such as asp are able to improve mass transport onto the surface through dehydrating positively-charged lattice sites. It is likely that this may be extended to other molecules that bind to calcite through carboxyl groups.

It can be seen in the snapshots of Figure 6.7 that binding primarily occurs through the side-chain carboxyl group, although trajectories did also reveal equivalent binding configurations for the alpha carboxyl group. This offers an explanation as to why aspartate has a high occlusion efficiency with respect to other primary amino acids. The alpha-carboxyl group, which all amino acids share, has a similar but slightly lower propensity to adsorb onto the surface compared with the asp side-chain.

6.2.3 Generalisation and adsorption free energies

The results of Section 6.2.2 have demonstrated that asp does not prefer to adsorb directly onto the kink lattice site to which CO_3 has the greatest propensity to adsorb, of all a- and d-kinks. This is important, as it indicates that the thermodynamic minimum can most likely be obtained without fully dehydrating the kink. Since we observe this for the kink to which CO_3 has the largest propensity to adsorb to the lattice site, we can proceed under the assumption that asp will not adsorb directly onto any kink lattice site. This allows us to calculate free energies for asp binding to every Ca-terminated kink sites without the need for explicit dehydration of the kink sites, thus saving a great deal of computational expense. It should be noted that, while this approximation is very likely to be valid for all a- and d-kinks, there does exist some degree of ambiguity for the b- and c-kinks, for which CO_3 adsorption free energies are not calculated.

To calculate adsorption energies, we ran metadynamics simulations using the exact method discussed previously, except we remove equation 5.1 as a reaction coordinate. This is done for every Ca-terminated kink. For most kinks, we tether the position of the terminating Ca ion to its lattice site as outlined in Section 4.2. The only exceptions to this were the b(i), c(i) and c(ii) kinks, which also have a tendency to adopt their bidentate configurations. For these ions, we remove the tether and include their z-position as a reaction coordinate. The adjacent CO_3 ion on the step was then constrained as per Section 4.2.

The eight calculated free energy functions are shown in Figure 6.8. For the b(i), c(i) and c(ii) kinks, the two-dimensional free energy surfaces are projected onto one dimension using equation 3.90. The original two-dimensional free energy surfaces are shown in Figure A6. The binding free energies were calculated from the free energy functions by averaging the free energy for values of the asp-kink z-distance between 15 and 20 Å. The adsorption free energies were subsequently calculated by applying the entropy corrections outlined in Section 4.6. The region corresponding to the bound state is given by the green-shaded region, which is integrated over to produce the adsorption free energies. The results are shown in Table 6.2.

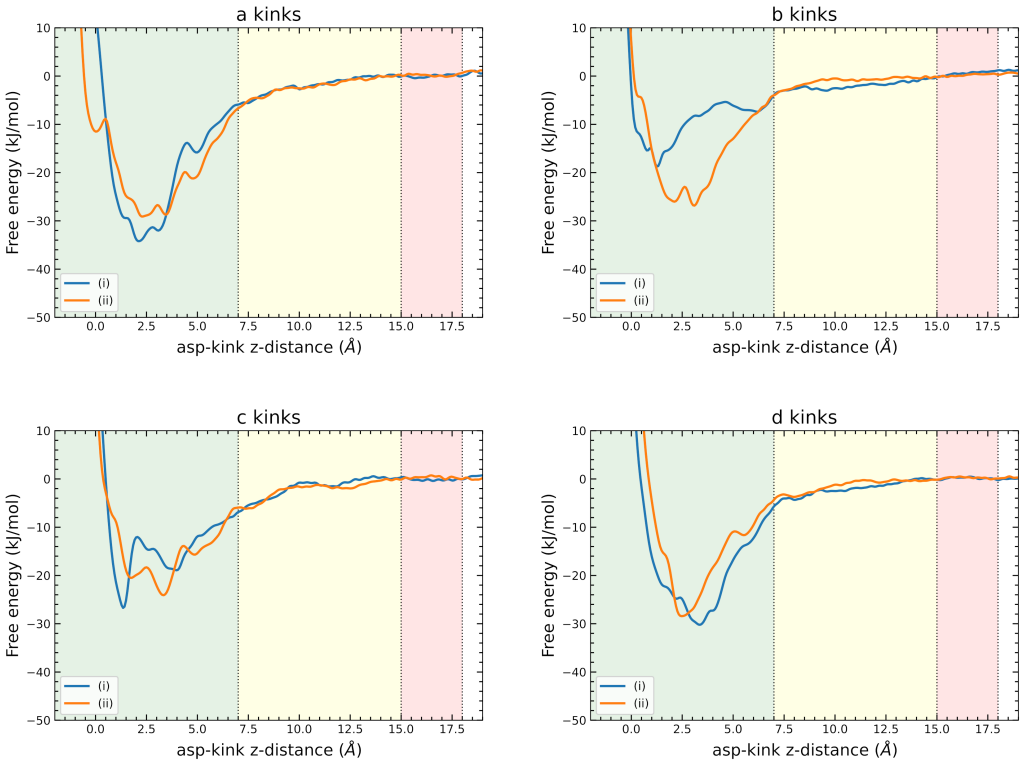

Figure 6.8: Free energy of asp as a function of its z-distance (average of two carboxyl group C atoms) from the kink site for every Ca-terminated kink site. The three highlighted regions represent bound (green), associated (yellow) and dissolved (red) states, corresponding to Figure 4.2. The simulation free energies, ΔG_{sim}, are calculated by averaging over the dissolved region.

Kink type		ΔG_{sim} (kJ/mol)	ΔG_{ads} (kJ/mol)
a	(i)	-31	-21
	(ii)	-34	-25
b	(i)	-18	-6
	(ii)	-26	-18
c	(i)	-26	-15
	(ii)	-24	-14
d	(i)	-30	-21
	(ii)	-28	-19

Table 6.2: Simulation and adsorption free energies for each kink type. ΔG_{ads} is calculated from ΔG_{sim} using the method discussed in Section 4.6. The free energy values are expressed to the nearest kJ/mol, as we expect the error of the free energy surface to be on the order of magnitude of 1 kJ/mol.

Most of the adsorption free energies are relatively similar, with the most notable exception being for the b(i) kink, to which asp adsorbs with a free energy of -6 kJ/mol. This outlier is a result of the b(i) kink naturally adopt-

ing its bidentate configuration (Chapter 5). The plots shown in Figure A6 demonstrate that the most stable binding configuration corresponds to one where the terminating Ca ion does not reside in its lattice site. This explains the anomalously high adsorption free energy for asp at this kink.

Experimental measurements of adsorption free energies have produced results of -17.1 kJ/mol [72] and -21.0 kJ/mol [71] by measuring the fractional inhibition of the crystal growth rate. Although this growth inhibition ultimately arises from a combination of several kink velocity inhibitions, the values in Table 6.2 can be compared with these experimental values to some degree. The lower values for each kink type in Table 6.2 appear in agreement with these experimental findings, although it is so far unclear exactly how the individual kink adsorption free energies will combine to produce a macroscopic adsorption free energy for an entire growth island or crystal. It is also worth noting that the adsorption free energies of asp are smaller in magnitude compared with those calculated for Ca (which vary between -15 and -31 kJ/mol) and CO_3 (which vary between -33 and -38 kJ/mol). This is consistent with the experimental observation that asp produces only a kink blocking effect, rather than a step pinning effect which would be seen for impurities which have a stronger binding to growth sites than Ca or CO_3 units.

6.2.4 Step velocity inhibition

It is important, at this point, to draw a distinction between adsorption free energies for individual kinks and experimentally measured adsorption free energies. Adsorption free energies are fundamentally related to the fractional growth inhibition and, in this study, are given by the Langmuir Adsorption Isotherm (equation 2.4). Our calculated kink adsorption free energies in Table 6.2 will therefore determine the fractional inhibition of the growth of that particular kink site. Experimentally found adsorption free energies relate to the fractional inhibition of the growth of the entire crystal. It is so far unclear as to how the multiple kink adsorption free energies will translate to step adsorption free energies (i.e. the adsorption energy as determined by the step velocity inhibition), and the crystal adsorption free energy observed in experiments.

To resolve this issue, we employed a KMC scheme to directly measure the fractional inhibition of step velocities. We ran individual calculations for acute and obtuse steps. The KMC model used is that introduced Chapter 4. It is worth noting here that we employ Ca and CO_3 dissolution rates parametrised to experimental step velocity dependencies on solution stoichiometry. In chapter 5, we derived adsorption free energies for most kinks, which could potentially be used to parametrise dissolution rates as is done for impurities (see section 4.8.2). However, we do not do this for several reasons: first, the free energy profiles in chapter 5 are more complex than those

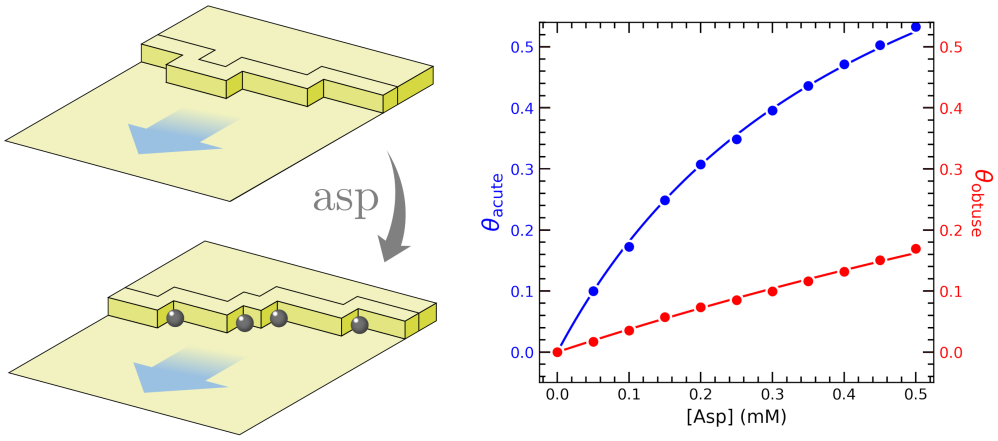

Figure 6.9: Step velocity fractional inhibition (θ) plots for acute and obtuse steps calculated using KMC. Both plots are fitted to equation 2.4, from which ΔG_{ads} is calculated. The schematic shows the build-up of kinks when higher concentrations of asp is added to the simulation, due to the stronger binding in one kink direction than the other. The inhibition of the step growth is therefore determined by the stronger binding.

we calculate for asp, consisting of two bound states, and we cannot assume a universal attachment rate as we do for asp; secondly, we have only run these calculations for deprotonated CO_3 where, in reality, the solution will also consist of bicarbonates. We instead opt for fitting all dissolution rates to experiment, and assume a semi-Kossel system, in which all Ca-terminated kinks have the same dissolution rates as each other, as do all CO_3 ions.

In order to determine the growth-inhibiting effects on acute and obtuse steps, we ran a total of 11 calculations with increasing asp concentrations for each step, ranging between 0 and 0.5 mM, and record the step velocity in each instance. Figure 6.9 shows the fractional inhibition of the step velocity for both acute and obtuse sites. Both plots were fitted to equation 2.4, which is used to calculate ΔG_{ads} for both step sites, denoted by $\Delta G_{\text{ads,a}}$ and $\Delta G_{\text{ads,o}}$ for the acute and obtuse steps respectively. The result is:

$$\begin{aligned} \Delta G_{\text{ads,a}} &= -19 \text{ kJ/mol} \\ \Delta G_{\text{ads,o}} &= -15 \text{ kJ/mol} \end{aligned} \tag{6.1}$$

Our results have demonstrated that the step velocity inhibition of the acute step is significantly greater than that of the obtuse step, which translates to a lower adsorption free energy for the acute step. This result is in some agreement with expectations, as it has been long assumed that binding to obtuse sites is far weaker than binding to acute sites. The evidence for this claim is in the morphology of calcite growth in the presence of asp, which has found minimal morphological effects on obtuse steps of growth islands and on obtuse corners of crystals, as discussed in Chapter 2. Nevertheless, studies have demonstrated a significant decrease in step velocity for

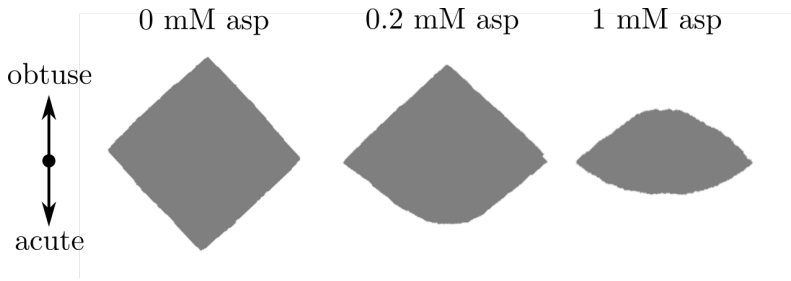

Figure 6.10: Morphologies of growth islands in the presence of asp at different concentrations calculated using KMC simulations. The growth island is depicted by the grey region, and the acute and obtuse directions are labelled.

the acute step when asp is included in solution [78]. Furthermore, studies of asp polymers found a greater morphological effect on obtuse steps for larger asp concentration [81]. Our results are therefore largely in agreement with previous studies, which ultimately suggest that binding to obtuse steps is weaker, but still significant. One limitation of the work carried out through this chapter, however, is a lack of explanation as to why asp does not induce morphology change for obtuse steps.

A further point of interest is that the step velocities obey the Langmuir adsorption isotherm despite the complex set of adsorption free energies for different kink sites. For example, propagation of the a-kink is inhibited more than that of the b-kink which propagates in the opposite direction (see Table 6.2). This leads to some ambiguity as to how the growth of the step is overall inhibited. However, KMC simulations revealed that the inclusion of asp results in a build-up of a-kinks due to the greater inhibition of propagation of a-kinks as illustrated in Figure 6.9. The density of a-kinks therefore becomes significantly larger, and the density of b-kinks drops to nearly zero. The growth inhibition of the step is therefore entirely determined by the inhibition of the a-kink.

6.2.5 A brief note on morphology

The KMC model applied to Section 6.2.4 may be extended to observing the morphology change that results from introducing asp to the solution. It can be tested whether our calculated adsorption free energies may reproduce the rounding of the acute steps without producing a rounding of the obtuse steps. By starting KMC simulations with a small growth island with a length and width of 100 units, and running until the simulation boundaries were reached, we were able to examine the morphologies of the growth islands. Figure 6.10 shows an example of the morphologies produced at different asp concentrations. It can be seen that, while the introduction of asp at small concentrations produces only a change in the acute morphology, the obtuse steps also experience a rounding effect at higher concentrations.

It appears that either our free calculations are incorrect, or that the in-

fluence of asp on obtuse step morphology is a more complex set of processes than what our model is able to accommodate. It is quite possible, for example, that the adsorption free energies at obtuse sites, notably the d-kinks are overestimated, since it is likely that the morphologies are extremely sensitive to the adsorption free energies, and a slight overestimate could have a significant impact. Another possibility is that the lack of influence of asp on obtuse morphologies is due to differences in the mechanisms of growth and dissolution found between the acute and obtuse steps. For example the morphology of etch pits formed in dissolving crystals indicates that the mechanisms of obtuse step dissolution significantly differ from those of acute step dissolution [131]. This alone, however, is not enough to explain why we observe a morphology change in the obtuse step. The lack of consistency with experimental morphologies is a significant limitation of our models used, and is something which requires addressing in future studies. It is clear that asp does have a significant interaction with obtuse growth sites, but the molecular mechanisms underpinning the lack of obtuse step morphology change remain unclear.

Although questions remain regarding the morphological impact, we do not consider this a major hindrance to our calculations. As stated above, the morphological impact is most likely highly sensitive to the calculated adsorption free energies, and any major morphological discrepancy can be explained by minor errors in the calculations of adsorption free energies. Furthermore, a previous study by Elhadj et al. [78] using AFM indicated that asp has a significant effect on the obtuse step velocities. By fitting the step velocities of Elhadj et al. to the Langmuir equation, we derive an asp adsorption free energy of -12.3 kJ/mol for the obtuse step, a value similar to our calculated value of -15 kJ/mol. We therefore proceed under the assumption that our results are reasonably accurate despite the issue with morphology.

6.2.6 Normal growth inhibition

Having determined the step velocities as a function of asp concentration, we wish to calculate the overall growth inhibition of the crystal. ΔG_{ads} is calculated experimentally by measuring the total rate of attachment of ions to the crystal. This value is proportional to the normal growth rate, i.e. the rate of growth of an island normal to the $[10\bar{4}]$ terrace. Therefore, the fractional inhibition of the normal growth rate will determine the fractional inhibition measured in experiments. Throughout this section, we consider growth centred around a screw dislocation to match the growth experimentally found for low supersaturations. The normal growth rate can be written in terms of the acute and obtuse step velocities as:

$$R(S) \propto \frac{v^{\pm}}{\lambda^{\pm}} \tag{6.2}$$

where R is the normal growth rate, S is the supersaturation, $v^{\pm}$ are the acute (-) or obtuse (+) step velocities and $\lambda^{\pm}$ are the corresponding terrace widths. The step velocities and terrace widths can be calculated using the method of Darkins et al. [62]. This method involves re-weighting the step velocities found in Section 6.2.4 such that they match experimentally observed step velocities for [asp]=0, using the following relation: [62, 132, 19]

$$\begin{aligned} v^{+} &= 64.910(S-1) \text{ nm/s} \\ v^{-} &= 16.488(S-1) \text{ nm/s} \end{aligned} \tag{6.3}$$

Assuming that the average critical length $\langle L_c \rangle$ is independent of [asp], we can calculate the average critical length [132, 65], and therefore the terrace widths. We can therefore re-write equation 6.2 as:

$$R \propto \frac{v^{\pm} \sin \phi}{1 + \frac{v^{\pm}}{v^{\mp}}} \tag{6.4}$$

where ϕ is the angle between adjacent spiral turns as described in [132]. Using equation 6.4, we calculated the fractional inhibition of R as a function of [asp], and fitted the results to equation 2.4. The plot of fractional inhibition is shown in Figure A7. We finally calculated ΔG_{ads} to be:

$$\Delta G_{\text{ads}} = -19 \text{ kJ/mol} \tag{6.5}$$

This value can be directly compared to experimental estimates of ΔG_{ads}. which vary between -21 kJ/mol and -17.1 kJ/mol. Our final adsorption free energy is clearly within the range of experimental error. This is a major success of the techniques used throughout this chapter, as well as the force fields used to describe asp. There are limitations to the results presented in this section, such as the weak binding of the amino groups and the inability to reproduce experimental morphologies. Nevertheless, the reproduction of quantitative experimental results related to growth rate is a success.

6.2.7 Binding configurations

As well as calculating adsorption free energies, the simulations carried out in the previous section allow us to elucidate the stable binding configurations. From the simulation trajectories, outputted every 10^5 time-steps, we observed that multiple binding configurations exist which are dependent on the value of the reaction coordinate. In order to determine the most frequently observed binding configurations, we searched through the trajectory and discounted any frames in which the reaction coordinate (the centre-of-mass of the two carboxyl group C atoms) is not within 0.2 Å of the value corresponding to the thermodynamic minimum, as determined by Figure 6.8. When the

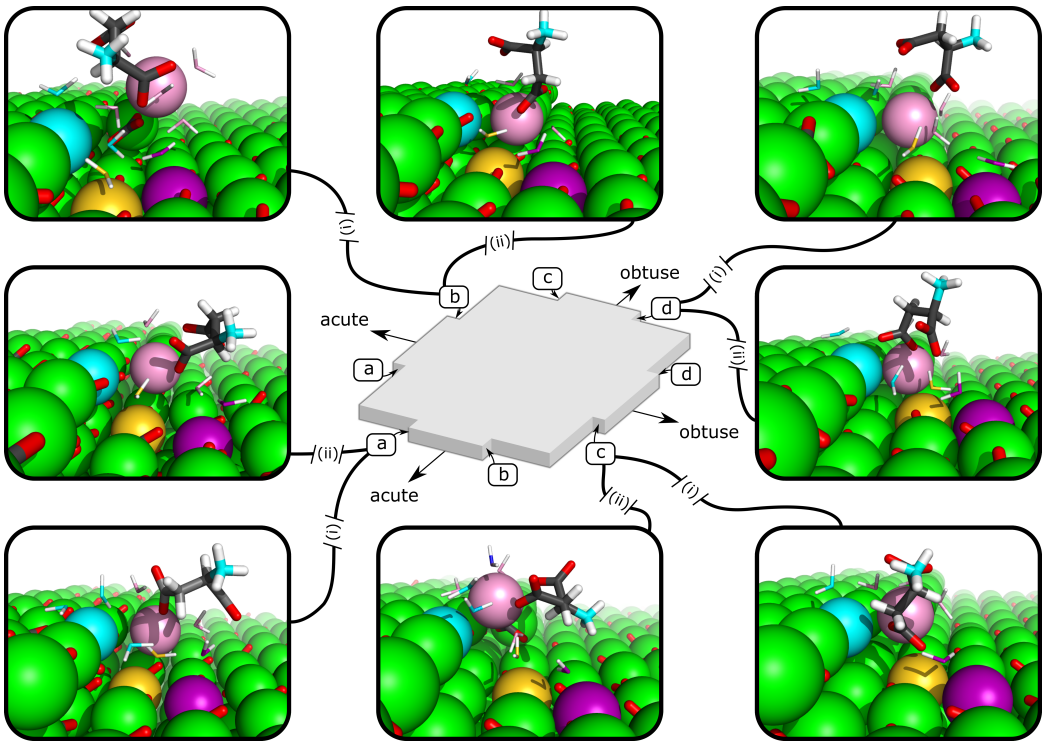

Figure 6.11: Stable binding configurations of asp adsorbing to all calcite Ca-terminated kink sites. The colour has been modified for the following ions: kink Ca ion (gold); terminating ion (pink); step ion next to the kink (cyan); nearby terrace ion (purple). The water molecules coordinated with the listed Ca ions are assigned the colour of the Ca ion to which they bind. Here, the step ions are shown in green, rather than the typical pink depiction.

position of the terminating Ca was also biased, the constraint was imposed on the second reaction coordinate (z-component of the terminating ion). The resulting trajectories in all cases were dominated by a single configuration. These configurations are shown in Figure 6.11. In Figure 6.11, the colours of the Ca ions near the kink site have been modified, and the water molecules coordinated with each Ca ion are also modified to be the same colour as the Ca to which they bind. The binding configurations show a general theme: one in which the carboxyl groups reside above the kink in a configuration akin to the bidentate configuration observed for Ca and CO_3 in Chapter 5. The one notable exception is the binding to the b(i) kink, in which the Ca ion remains in its bidentate configuration. For both c-kinks, the adsorption of asp results in the terminating Ca ion transitioning into its lattice configuration.

6.2.8 Water displacement

In Section 6.2.2 it was determined that the water molecule nearest to the kink site was not displaced upon the adsorption of asp to the kink site. Building upon this surprising result, we proceeded to examine the local water structure

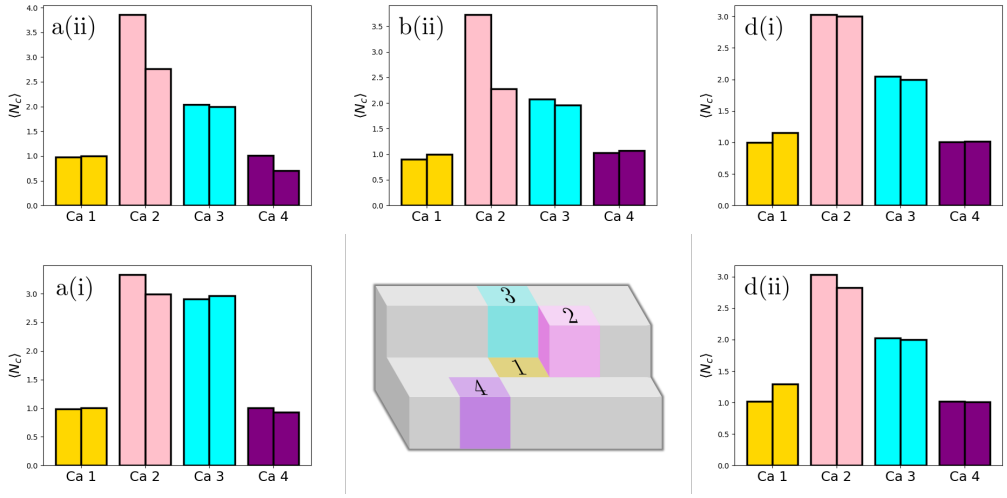

Figure 6.12: Calculations of the average number of coordinated water molecules $\langle N_c \rangle$ for four different Ca ions, labelled 1, 2, 3 and 4 as shown in the schematic, in the vicinity of several kink sites. The bars on the left-hand side represent $\langle N_c \rangle$ when the asp is fully dissolved, and the bars on the right-hand side represent $\langle N_c \rangle$ when the asp is residing in their most stable binding configurations. The results for the remaining kink sites are not shown as the terminating Ca ion prefers to adopt the bidentate configuration (Chapter 5), and their water structures are less well defined.

near the kink site both in the absence, and the presence of asp. Here, we examined the frames of the trajectory which correspond to a bound state, as was done in the previous section. For these frames, we calculated the average coordination number of water molecules with the four Ca ions highlighted in Figure 6.11, defined as $\langle N_c \rangle$.

$$N_c = \sum_j \chi_j \qquad \chi_j = \begin{cases} 1, & r_j \leq r_0 \\ 0, & r_j > r_0 \end{cases} \tag{6.6}$$

where r_0 is a cut-off distance which we set to 3.0 Å. Note that this is the same quantity calculated previously in equation 5.5. The results are shown in Figure 6.12.

The most striking result shown in Figure 6.12 is the lack in difference of water coordination between when asp is adsorbed to the kink and when it is dissolved in solution. It appears that the binding of asp causes no displacement of water molecules in the majority of cases. It should be noted that there are minor differences observed in most cases, but these can be attributed to water molecules which briefly enter the coordination sphere, but do not become fully adsorbed. The only significant differences seen are for the a(ii) and b(ii) kinks, for which the terminating ion sheds a single water molecule upon the adsorption of asp. It is particularly interesting that asp causes such a minimal displacement of water molecules, as it disagrees

with existing hypotheses that impurities such as asp are able to dehydrate the surface.

Another important caveat to these results is that, while asp mostly does not displace water molecules, there is some displacement observed which is not explicitly facilitated by a bias potential, as we have only biased the position of the asp. Although we established that dehydration is not an issue for lone Ca ions, we cannot assume the same for surface ions. To test this issue, we ran a calculation in which we tracked the identity of the water molecules coordinated to the terminating Ca ion of the a(ii) kink. We do not perform the same level of analysis as we did in section 5.3, as the residence times calculated were generally too large to obtain sufficient statistics. We typically observed residence times of about 10 ns. This is problematic for our simulations, as they typically spanned about 300 ns, thus only allowing a handful of dehydration events to take place. It is therefore possible that our free energy profiles calculated for the a(ii) and b(ii) kink sites are inaccurate. This casts some doubt onto the validity of our results, and should be taken into consideration in future, despite the success of this study in reproducing experimental adsorption free energies.

6.2.9 A note on force fields

It is important to note that the simulation free energies calculated in this section are vastly different from those calculated in Section 6.1. Using the force fields of Raiteri et al., we find that binding to calcite occurs exclusively through the carboxyl group, and the amino group does not bind. In contrast, the Amber force fields produce a strong interaction between amino groups and CO_3 ions, but a very weak interaction between carboxyl groups and Ca ions. This stark contrast highlights the sensitivity of quantitative and qualitative observables on the choice of force fields, and the challenges presented for fitting force fields.

The free energy for ethylammonium binding to the kink site calculated in Section 6.1 is significantly larger than those calculated for asp in this section. Since the results in this section produce an accurate adsorption free energy, it must be the case that the Amber force fields drastically overestimate the binding free energy. This does not necessarily invalidate the conclusions of Section 6.1, as the binding configurations are unlikely to change if the binding is weaker, and we are only interested in comparing the results of Section 6.1. Conversely, it is likely that the binding through the amino group is underestimated by Raiteri et al. The experimentally observed occlusion of polyamines and lysine, and their morphological impact, are proof that amino groups must be able to bind to calcite kinks. We do not expect the underestimated binding of amino groups to have much impact on the results of this section, as we believe that the interaction of asp with calcite is mostly driven by the carboxyl groups, and that the amino groups will have little

impact. However, the work carried out in this chapter indicates that future challenges exist for the development of thermodynamically consistent force fields.

6.3 Conclusions

This chapter began with a series of calculations of binding free energies of amine groups to calcite terrace, step and kink sites. The primary amine, ethylammonium, was found to bind strongly to all sites. The arginine side chain was found to bind far more weakly to all sites compared with ethylammonium. This led to the conclusion that arginine does not bind through its side chain, and instead through its primary amine and carboxyl groups. The secondary amine, diethylammonium, was found to bind strongly to terraces, but did not bind to steps or kinks, indicating that its main purpose is to stabilise larger polyamines which bind to steps and kinks through its primary amine groups. Our results, in conjunction with experimental results, have demonstrated that positively-charged functional groups are, contrary to the general consensus, able to bind strongly to calcite, and cannot be neglected when considering the binding of e.g. amino acids to calcite.

The second part of this chapter concerned determining the free energy of aspartate adsorbing to calcite, adopting more sophisticated force fields than those used previously. Contrary to the results of the first part of this chapter, it was found that aspartate adsorbs to calcite primarily through its two carboxyl groups, with minimal contribution from the amine group, and multi-site adsorption is not observed for asp in simulations. These results agree with previous expectations that aspartate binds primarily through its carboxyl group, and indicates why the carboxyl-rich aspartate has a strong degree of occlusion, as observed in experiments. However, the discrepancy between these findings and those above have highlighted an ever-present issue in molecular simulation i.e. the inconsistency across different force fields. The role of amine groups in driving the adsorption of amino acids remains largely ambiguous.

The most stable binding configurations of asp were found to be where the amino acid adopts a configuration akin to the bidentate configuration observed in Chapter 5. The carboxyl groups either interact directly with ions in the upper terrace, or reside above the upper terrace, and the water molecule nearest to the kink site is not removed from the kink site. The minimal displacement of water throws into question the validity of the proposed mechanism in which mass transport to the surface is improved by the dehydrating of the surface by impurities such as asp [78, 88].

Aspartate adsorption free energies to most kink sites were found to be relatively similar, presumably due to aspartate binding through the same mechanism. The only notable exceptions are for where the terminating Ca

ion does not sit in its lattice site. The calculated adsorption free energies were used to parametrise a Kinetic Monte Carlo scheme which measured the fractional inhibition of step velocities as a function of aspartate concentration. The step velocities obey a Langmuir adsorption isotherm for both steps due to the reduction in the density of the kink sites to which asp binds less strongly to. This causes the most expressed kink site to be the kink sites to which asp binds the strongest. The growth of acute steps is inhibited more by asp, which is consistent with experimental observations that obtuse step morphologies are less affected by asp, although the directly observed morphologies from simulations highlight that more work is required in understanding the nature of the morphological impact of impurities such as asp on calcite. Finally, the adsorption free energy of aspartate as observed in experiments was calculated to be -19 kJ/mol, which is within the range of experimentally derived values. This result is a major success for the force fields and techniques used in this chapter.

7 MECHANICAL PROPERTIES OF IMPURE CALCITE

Chapter 6 concerned the adsorption of impurities to the calcite surface during the growth process. Of the impurities that adsorb to calcite, some become kinetically trapped into the lattice as the crystal grows around the impurity. When this occurs, the impurity is permanently occluded into the host crystal. This is how impure crystals typically form, although it should be noted that there are many other mechanisms through which impurities may be incorporated into the lattice during the nucleation and growth processes. Calcite provides a good example of the property-enhancing effects of impurities. While pure calcite is a brittle material, i.e. fractures under a relatively low stress, studies have found that the hardness of the crystal can be increased by introducing impurities such as asp [92]. Where an understanding of impurity directed crystal growth is vital for our understanding of biomineralisation, it is also crucial to gain an understanding of how impurities may alter the structural and mechanical properties of crystals at the molecular scale.

This chapter concerns two case studies where impurities have produced novel effects on the crystal lattice. The first section concerns a case study where polymer-coated gold nanoparticles were found to be able to occlude in calcite at high concentrations while preserving the single crystallinity of the material. This came as a surprising result, particularly as the polymers in question contained only weakly-charged functional groups, and are therefore expected to have little interaction with calcite. Nevertheless, the nanoparticles were able to occlude at concentrations of up to 37 wt% while preserving the crystal structure of calcite. As well as occluding at high concentrations, the nanoparticles produced a small change in the lattice spacing of the host crystal. Most surprising was the result that the nanoparticles produced a decrease in the lattice spacing along the a-direction, but an increase in lattice spacing in the c-direction. The first section of this chapter concerns the attempt to understand this using molecular simulation.

The second section of this chapter concerns a unique case study of biogenic calcite within the brittlestar Ophiocoma wendtii. Calcitic lenses within the skeleton of the brittlestar contain arrays of Mg-rich regions which are believed to improve the fracture toughness of the material by inducing a compressive stress on the host matrix. Throughout this section, we model such a system: we calculate hydrostatic stress tensors, as well as directly simulate a crack propagation event in order to identify the toughening mechanisms only observable at the molecular scale.

7.1 Hydroxyl-rich Au nanoparticles

The conventional view on biomineralisation is that highly charged, acidic molecules are the key to controlling crystallisation of minerals such as calcium carbonate [133] and, as a result, studies of bio-inspired crystallisation methods have primarily focussed on acidic molecules, which have been noted to exert control over morphology [134], orientations [135] and polymorphism [136]. Recently, a study by Kim et al. [86] has challenged this view by demonstrating that low-charge macromolecules can outperform their acidic counterparts in terms occlusion in calcite. They found that polymer-coated gold nanoparticles were able to occlude in calcite at concentrations of up to 37 wt%. Despite such high levels of occlusion, these functionalised nanoparticles resulted in no signature change in morphology or crystallinity, which is contrary to expectations, although unsurprising considering the low charge of the functional groups (and therefore the expected weakness of their interactions with calcite. Most noteworthy are the results found for the occlusion of 4 nm poly (glycerol monomethacrylate)-functionalised gold (PGMA-Au) nanoparticles, which were found to occlude at the highest concentrations.

Aside from the fact that such low-charge nanoparticles can counter-intuitively occlude at such concentrations, a surprising result emerged from powder x-ray diffraction (powder XRD) experiments. Naturally, the inclusion of the nanoparticles induced a shift in lattice spacing in the crystal. However, the sign of the change in lattice spacing was dependent on both the concentration of the PGMA-Au nanoparticles, and the plane normal to which lattice spacings were measured. Most notably, the nanoparticles induced a compressive strain in the a-direction, and a tensile strain in the c-direction. However, when the concentration of the nanoparticles was increased, a compressive strain was universally observed. These results are surprising and unintuitive. Assuming the nanoparticles induce a universal tension on the crystal, one would expect either a universal expansion or contraction. However, molecular simulation may be able to uncover whether or not this is, in fact, due to the properties of calcite rather than the properties of the nanoparticles and any direct effects they induce.

Since the polymers that coat the gold are weakly interacting with the calcite, it can be assumed that they should exert minimal forces on the lattice in which they occlude. It can also be assumed that they have minimal effects on the interface of the hole of the lattice that they leave behind after occluding. For this reason, we assume that the nanoparticle itself can be approximated as a vacuum occluded into the crystal. This greatly simplifies the problem, as it only involves removing a hole from the lattice. The problem remains, however, that we don't know the shape the hole that is left behind will be. For this reason, two approximations were chosen: either a spherical hole or a {10.4} rhombohedral hole. The spherical hole represents a case where the nanoparticle is less fluid, and the morphology of the hole is kept

fixed in the form it takes as the nanoparticle occludes (which we assume to be spherical). The {10.4} rhombohedral hole represents an approximation where the spherical hole is fluid enough for the hole to relax to its lowest energy morphology.

In order to approximate a high concentration of PGMA-Au nanoparticles in calcite, a periodic calcite crystal was simulated. For simplicity, the simulation box was set up such that the x and z directions of the simulation box represented the a and c directions of calcite respectively. The box lengths corresponded to 18, 18 and 6 repetitions of calcite's a,b and c unit cells respectively. A monoclinic simulation cell was used so as to accommodate the calcite's unit cell. To create a spherical hole, ions within a defined region were removed from the lattice. In order to ensure charge neutrality, excess ions at the boundary of the hole were selected at random and removed until the charge was neutral. The size of rhombohedral hole was calibrated such that the number of atoms remaining in the simulation was as similar as possible to the spherical case.

All simulations consisted of an equilibration process where the calcite was relaxed under NVT for 100 ps, then under NPT for 200 ps. For the NPT simulations, all simulation box lengths and triclinic parameters were allowed to relax. Conditions were then kept under NPT for data harvesting, which continued for a further 700 ps. After preliminary simulations were run, it was noted that the same process could not be used to simulate low nanoparticle occlusion, as a much larger box is needed, and simulations took too long in this limit.

7.1.1 Local stress field

Figure 7.1 shows the results for the hydrostatic stress field (defined by $\sum_\alpha \Pi_{\alpha\alpha}(r)$, where $\Pi_{\alpha\beta}(r)$ is the stress tensor) obtained from simulations of calcite with the two holes included in the lattice. Studying Figure 7.1, the results are largely as expected. A tensile stress is present, on average, at the interface of the hole in each case. The bulk crystal is under some competing pressure, although it is worth noting that, in each case, the sign of the stress is not constant throughout the bulk crystal. It appears that, due to the small magnitude of the stresses at the interface, the compensating pressure decays relatively quickly. This implies that most of the effects that result in lattice distortion are likely to occur mostly near the interface of the hole. This change in sign also implies that there is either very little or no interaction with the hole and its periodic images. However, in the rhombohedral case, the tensile stress does continue across the boundary along the a direction where it bisects the hole. Therefore, some interaction with periodic images may be present.

The magnitudes of the stresses are higher in the spherical case, although this is to be expected, since the lowest energy surface is the {10.4} surface.

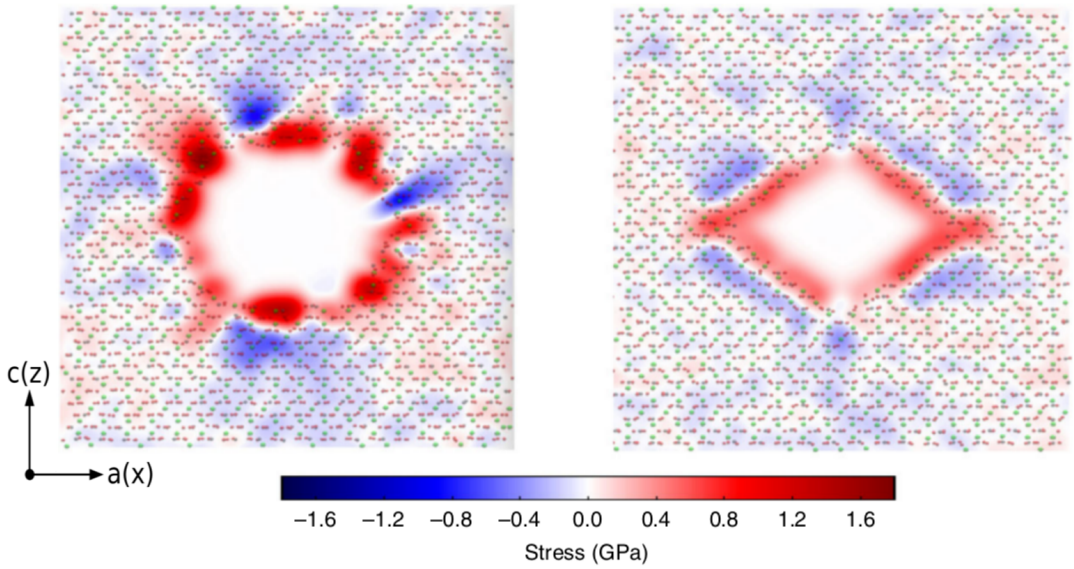

Figure 7.1: Hydrostatic stress fields obtained from simulations of calcite with a spherical (left) and rhombohedral (right) hole.

The stress is more uniform at the interface in the case of the rhombohedral hole, but this is again to be expected since the rhombohedral hole presents low energy surfaces of calcite.

7.1.2 Lattice spacing

For PGMA-Au nanoparticles in calcite, there are likely to be local tensile stresses present at the interface of the hole in the crystal. Since the changes in lattice parameters found experimentally are very small, and the PGMA-Au nanoparticles occlude at such high density, it is likely that the stresses at the interface will have an effect on the overall average lattice spacings. For this reason, measuring the change in simulation box dimensions is insufficient in deducing the average lattice spacing. For this reason, the method introduced in Section 4.4 was used. This allows not only a more accurate measurement of average lattice spacing, but an insight into how the distribution of lattice spacing changes. This allows us to draw conclusions as to whether peak broadening would be observed in an X-ray diffraction experiment.

In order to implement the method described in Section 4.4, LAMMPS was customised in order to project the Radial Distribution Function onto arbitrary planes. The x-component and z-components of the RDF were calculated separately during a single simulation. Only the calcium and carbon atoms were included in the calculation, as including the oxygen atoms in the calculation would interfere with the calculation of the lattice spacings. The perpendicular cut-off, d_c, was set to 1.5 Å in order to account for thermal fluctuations while preventing adjacent atoms from contributing to the RDF. It is worth noting, though, that choosing a cut-off anywhere between 1 Å and

	Experimental	Spherical hole	Rhombohedral hole
a lattice strain (%)	−0.0299	−0.0278	−0.0334
c lattice strain (%)	0.0137	−0.0145	0.0083

Table 7.1: Comparison of experimental data of strain in the a and c direction obtained from powder-XRD experiments with simulations of a spherical and rhombohedral hole obtained from the projected RDF.

2 Å had a negligible effect on the results.

Table 7.1 shows the results of the average lattice spacing found from the projected RDF obtained from simulations. The experimental column shows the novel combination of contraction in one dimension and expansion in the other. It can be seen that spherical hole is unable to replicate this effect. A rhombohedral hole, on the other hand, is seen to reproduce the effect. Not only does it do this, but the strains are of the same order of magnitude as seen in experiment in both dimensions. Considering the number of approximations made in this simulation (most notably, the replacing of the PGMA-Au nanoparticle with a vacuum), this is a surprisingly accurate result.

It is quite intuitive that an interface under tension should result in a universal contraction of the lattice, so it is unsurprising that the spherical hole results in exactly that. It is less intuitive that the rhombohedral hole should replicate the unusual result of contraction in one dimension and expansion in the other. However, analysis of the average box dimension also found contraction in the a-direction and expansion in the c-direction. Seeing as the stress maps reveal little or no interaction with periodic images along the c-direction, it is likely that the hole itself has undergone contraction in the a-direction and expansion in the c-direction. Since the elastic coefficient for calcite is higher in the a direction [137] than the c-direction, it is likely that the tensile stress at the interface is higher in the x-direction. This causes a stronger contraction in the a-direction, which forces the hole apart in the c-direction. This causes the corresponding expansion/contraction in the bulk crystal.

7.2 Mg-rich coherent nanoparticles

One unique case study of impurities producing crystals with enhanced mechanical properties concerns the brittlestar Ophiocoma wendtii. The arm plates of these brittlestars are covered by roughly 10-micrometre sized lenses composed of calcite [138]. The function of these lenses is to focus light onto photoreceptor nerve bundles positioned beneath the lenses. Each lens is aligned with its optical axis parallel to the c axis of calcite in order to minimise the effects of birefringence. Calcite, being transparent and highly abundant in nature, provides an ideal material for these lenses. However,

as discussed, calcite is a famously brittle mineral with a low fracture toughness. The brittlestar is able to overcome this issue through the introduction of coherent, Mg-rich nanoparticles within the host crystal. As discussed in Chapter 2, fracture toughness may be improved through formation of complex hierarchical structures of hard and soft materials. However, this method would be inappropriate, as it would be detrimental to the optical properties of the material. However, in 2017, Polishchuk et al. [3] identified the mechanism of coherent nanoparticles in the brittlestar when examining the nanostructure of the lenses. It was discovered that each calcitic lens was found to contain a dense array of magnesium calcite nanoprecipitates, about 4 nm in diameter. The coherency of the nanoprecipitates allows the preservation of the optical properties of calcite. These nanoprecipitates are believed to form during crystallisation of calcite from amorphous calcium carbonate (ACC) [139]. The solubility of Mg in ACC is much higher than in calcite, therefore during crystallisation Mg rich regions would be expected to form. Coherent nanoparticles, often referred to as Guinier-Preston (GP) zones [140, 141], are already well known in metallurgy, and have been the subject of previous molecular dynamics studies [142, 143]. Prior to the findings of Polishchuk et al., no examples of such a phenomenon had been observed in any living organism. Whereas the formation of GP zones in metallurgy involves extensive heating [144], the brittlestar is able to achieve this process in calcite under ambient conditions. In metals, GP zones increase the tensile strength of the host material [145], generally at the expense of increased brittleness. However, Polishchuk et al. propose that the nanoprecipitates have a different effect on the mechanical properties of calcite: rather than increasing the tensile strength by inhibiting dislocation motion, as in metal alloys, the nanoprecipitates in calcite increase the fracture toughness by inducing a compressive stress in the host matrix. Such a prestressing mechanism is employed in other brittle materials, such as tempered glass and prestressed concrete. In calcite, the compressive stress in the host matrix is induced by the coherent Mg-rich nanoprecipitates. Due to the small size of the Mg ion compared to the Ca ion, there is a local tensile stress within the nanoprecipitate, causing a compression of the lattice spacing. The continuity of the lattice planes ensure that, where the nanoprecipitates are under a tensile stress, the surrounding matrix is under a compressive stress. It is proposed that this compressive stress inhibits crack propagation and, therefore, increases fracture toughness.

While a proposed mechanism exists for the purpose for the Mg-rich nanoprecipitates in the skeleton of Ophiocoma wendtii, it has yet to be confirmed unambiguously using molecular simulation. Generally speaking, impurities and their mechanical effect on calcite at the molecular scale is a field which has been examined with molecular simulation in only a select few studies [92, 99] (see Chapter 2). The case study of the brittlestar, however, provides an ideal context for the study of impure calcite crystals, as we have the tools at our disposal to simulate a calcite crystal containing a Mg-rich coherent nanoparticle.

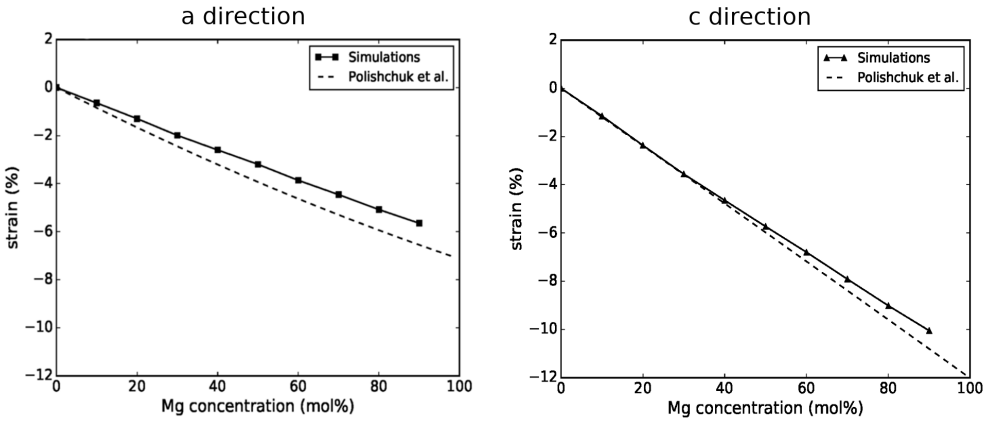

Figure 7.2: Comparison of lattice strain in the a direction and c direction between molecular dynamics simulations and the empirically derived relation found by Polishchuk et al [3].

Throughout this section, we use molecular dynamics to investigate the effects of magnesium incorporation in calcite. We examine the effects of different concentrations of magnesium on calcite lattice parameters. We investigate the hydrostatic stress field and the magnitude of the compensating stress field in the host matrix. Finally, we use crack propagation simulations to examine the effect of magnesium nanoprecipitate incorporation on calcite fracture toughness.

In this section, the calcite interactions were modelled with the rigid ion version of the Pavese et al. [94, 146] force field. These force fields are preferred over the standard force fields of Raiteri et al. as they were fitted explicitly to the elasticity tensor of calcite. The Mg-ion interactions were obtained from Raiteri et al, due to the lack of any equivalent force fields similar to Pavese et al. for Mg. Periodic simulation cells with edge lengths roughly equal to 8 nm in all dimensions were populated with calcite. Magnesium ions were introduced into the calcite lattice by randomly substituting for calcium ions. A probability of assignment for Ca and Mg was used to set the percentage of magnesium ions. When modelling spherical nanoprecipitates, the Mg-substitutions were confined to a spherical volume with a diameter of 4 nm, roughly the size of the observed nanoprecipitate. In all simulations, the nanoprecipitates were populated with 40 mol% Mg in order to replicate the empirically found concentration observed by Polishchuk et al. [3].

7.2.1 Force field validation: strain dependence on Mg content

To test the suitability of the combination of force fields employed in this section, we began with a uniform random distribution of Mg ions in a bulk calcite lattice, and relaxed the cell vectors to eliminate the stress. The re-

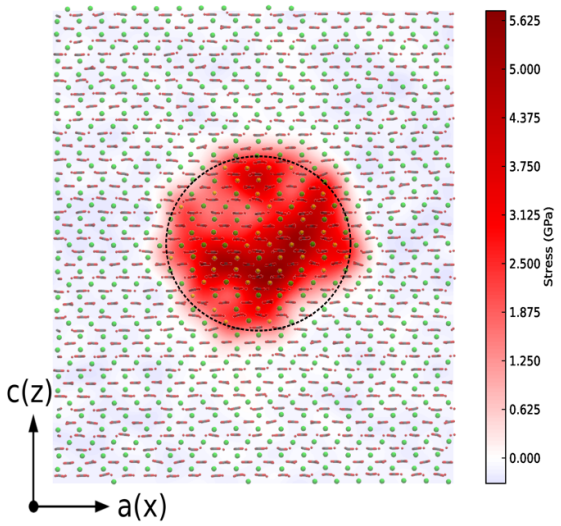

Figure 7.3: Hydrostatic stress tensor for a periodic (bulk) calcite crystal with an incorporated 40 mol% Mg nanoprecipitate of diameter 4 nm. Atoms from a cross-section of the cell are also shown. Ca is shown in green, C in grey, O in red and Mg in yellow. Here, red represents regions under tensile stress and blue represents regions under compressive stress. The cell dimensions are 86 Å in the x-direction and 105 Å in the z-direction. The dashed circle indicates the Mg-rich domain. The hydrostatic stress field is roughly homogeneous throughout the host crystal.

sulting cell strain as a function of Mg concentration is shown in Figure 7.2. The results are compared with the empirically derived relations for the lattice parameters in the a and c-directions as a function of Mg concentration as found by Polishchuck et al. It is clear from the results that the force fields predict a reasonably accurate relationship between strain and concentration, especially in the more elastic c-direction.

7.2.2 Stress distribution for calcite with 40 mol% Mg nanoprecipitates

For simplicity, the simulation box was set up such that the x-direction corresponded to the a-axis and the z-direction corresponded to the c-axis of the crystal. The box dimensions were chosen such that the lengths corresponded to 18, 18 and 6 repetitions of the primitive hexagonal cell along the a, b and c axes respectively. All cell dimensions were relaxed under NPT at atmospheric pressure to allow for stress relaxation. The crystal was equilibrated for 100 ps, and a 500 ps simulation was carried out for the calculation of the local stress field.

The local hydrostatic stress field was computed for a cross-section of the simulation cell that bisected the centre of the nanoprecipitate, shown in Figure 7.3. The radial and angular components of the stress tensor field are

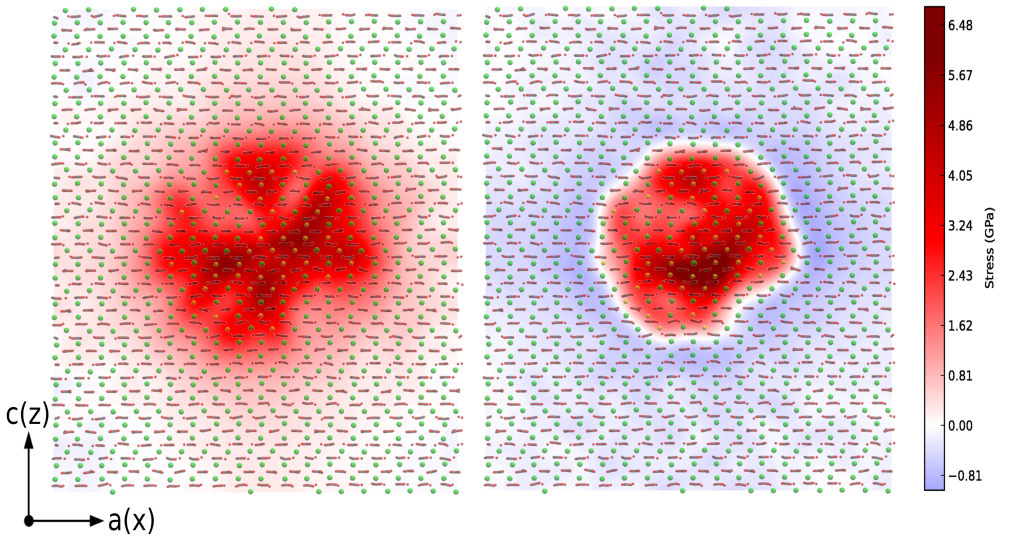

Figure 7.4: Radial (left) and tangential (right) components of the stress tensor for the same system as Figure 7.3 in cylindrical coordinates. The smaller lattice spacing within the nanoprecipitate causes a tensile radial stress (red). This causes an angular compression (blue).

shown in Fig 7.4. Figure 7.3 shows a high tensile stress within the nanoprecipitate, caused by the smaller size of the Mg ions, as can be seen in the red region encompassing the nanoprecipitate in Figure 7.3. The surrounding host matrix experiences a compressive stress in order to compensate for the tensile stress, as can be seen in the blue regions outside the nanoprecipitate. The hydrostatic pressure appears roughly homogeneous throughout the bulk crystal, although Figure 7.4 sheds more light on the complexity of the stress distribution: the radial stress distribution (left) shows a host matrix almost entirely under tensile stress. Furthermore, the radial stress is larger in magnitude nearer the nanoprecipitate. The tangential component, however, shows an entirely compressive stress. It is apparent, therefore, that the stress encountered by a propagating crack will be dependent on where it is relative to the nanoprecipitate. These results provide insight into how crack propagation may be inhibited in some regions, and deflected in others, as elaborated in the following section.

The calculation of the atomic virials also allows the average stress induced in the host matrix to be calculated. By excluding all atoms within the defined spherical region, and accounting for the resulting volume change, the stress tensor components for regions outside the nanoprecipitate ($\sigma_{\alpha\beta}$) can be calculated by averaging the remaining atomic virials over the remaining volume. The results of this process for the xx and zz components, along with the hydrostatic stress, σ_H, are given in table 7.2, where they are compared with the empirical estimates of Polishchuck et al. [3] derived from a continuum elasticity model. While the values are consistently lower, the agreement

is generally quite impressive considering the various approximations in both data sets.

	Simulations	Polishchuck et al.
σ_{xx} (MPa)	−145	−180
σ_{zz} (MPa)	−127	−140
σ_H (MPa)	−139	−170

Table 7.2: Comparison of stress tensor components for the host matrix between simulations and estimates of Polishchuck et al.[3]

7.2.3 Crack propagation

In order to model a calcite crystal with a pre-existing crack, a thin slab of calcite was simulated. A monoclinic, periodic simulation cell was used to allow periodicity along two of the rhombohedral crystal's three axes. The crystal was oriented such that the free (10.4) surface of the slab was normal to the z-direction of the simulation cell. The length of the cell in the x and y-directions was 75.22 Å. Along the z-direction, the crystal accounted for 81.46 Å. The simulation cell was also periodic in the z-direction, but with a 40 Å vacuum region to emulate free surfaces. At the upper surface of the crystal, an initial crack was introduced to the slab by defining two adjacent regions of atoms. Each region was effectively infinite along the y-direction, had a thickness of 15.7 Å along the x-direction, and penetrated 15.1 Å into the crystal in the z-direction. All interactions between these two regions were disabled so as to create an effective interface that would nucleate a crack upon being strained. During simulations, the Ca ions at the lower face of the slab were constrained to their initial z-coordinates using a harmonic potential with a spring constant of 50 kJ/(mol)/Å^2 in order to simulate the effect of bulk crystal at the base of the slab.

Simulations consisted of an equilibration process of 0.2 nanoseconds, where all cell vectors except the z-length were relaxed under NPT at atmospheric pressure. The barostat was then changed such that only the cell length in the x-direction was able to fluctuate at atmospheric pressure. The cell was deformed along the y-direction using a constant engineering strain rate of 0.01 ps^{-1} over a period of 10 ps. During this 10 ps period, and following a 2 ps equilibration period, the yy-component of the stress tensor was evaluated every 0.1 ps and averaged over the remaining 8 ps. Simulations were repeated 10 times using different velocity seeds and nanoprecipitate configurations (i.e. different seeds for determining the Mg distribution within the nanoprecipitates).

Figure 7.5 shows visualisations of crack propagation in calcite under several different conditions. In the control situation, where the calcite contains no Mg (Figure 7.5(a)), the crack propagates cleanly down the low-energy

(10.4) plane and, with the exception of the constrained region at the lower surface, the two surfaces are completely separated. The inclusion of the nanoprecipitate, however, brings about a visibly significant change in the nature of the crack propagation. When the crack bisects the nanoprecipitate (Figure 7.5(b)), propagation still occurs along the (10.4) plane, but the fracture is far more disordered, with ions bridging the two resulting surfaces. The Mg-rich domain itself is also observed to fail before the crack has reached the precipitate, presumably because of the tensile stress in the nanoprecipitate. When the crack reaches the nanoprecipitate slightly off-centre, as shown in Figure 7.5(c), the crack changes direction when propagating through the Mg-rich domain, producing a rough, high-energy interface. This again may be attributed to the high tensile stress within the nanoprecipitate. Where, as shown in Figure 7.5(d), the crack propagates around the nanoprecipitate, the crack is deflected from the low-energy (10.4) plane towards the nanoprecipitate. In light of the stress distribution in Figure 7.4, it is unsurprising that the crack would be deflected towards the nanoprecipitate, since the radial stresses are highest near the Mg-rich domain, and a crack propagating through this region is able to relieve said radial stress. This reveals an interesting consequence of the presence of a Mg-rich domain: the complex nature of the stress distribution may result in a deflection of crack propagation away from the lowest energy (10.4) surface, which will contribute to the increased toughness.

The stress-strain curves calculated during crack propagation are shown in Figure 7.6. Note that the stress values are averaged over 10 simulations for each point. The points of maximum stress on these curves correspond to the initiation of the crack propagation. The curves corresponding to Figure 7.5(a) and (d) are indistinguishable before the point of fracture, suggesting that the precipitates have little influence on cracks that do not approach closely. The curves corresponding to Figure 7.5(b) and (c) display two important features. Firstly a higher tensile stress (an additional 223 and 180 MPa respectively) is reached before crack propagation, and secondly some stress is retained at the end of the simulation. The increase in maximum tensile stress is the result of the compressive stress in the matrix, as suggested by Polishchuk et al. The residual stress is a result of the incomplete fracture, as the fracture surfaces are bridged by a disordered region of ions at the end of the calculation. Both of these effects would increase the fracture toughness. The crack that propagates off centre (Figure 7.5(d)) does not display a significant increase in the tensile stress at the initiation of crack propagation but it does show significant residual stress, partially due to bridging of the surfaces.

As seen in Figure 7.5: rather than the crack propagating cleanly down a (10.4) surface, the breaking of the crystal is more complex; this is clearly an effect of the crack being deflected towards, and even through, the Mg-rich domain. This result demonstrates a possible inhibitory mechanism for crack propagation not previously reported to our knowledge, and may suggest a novel mechanism for improving fracture toughness. Polishchuk et al. pro-

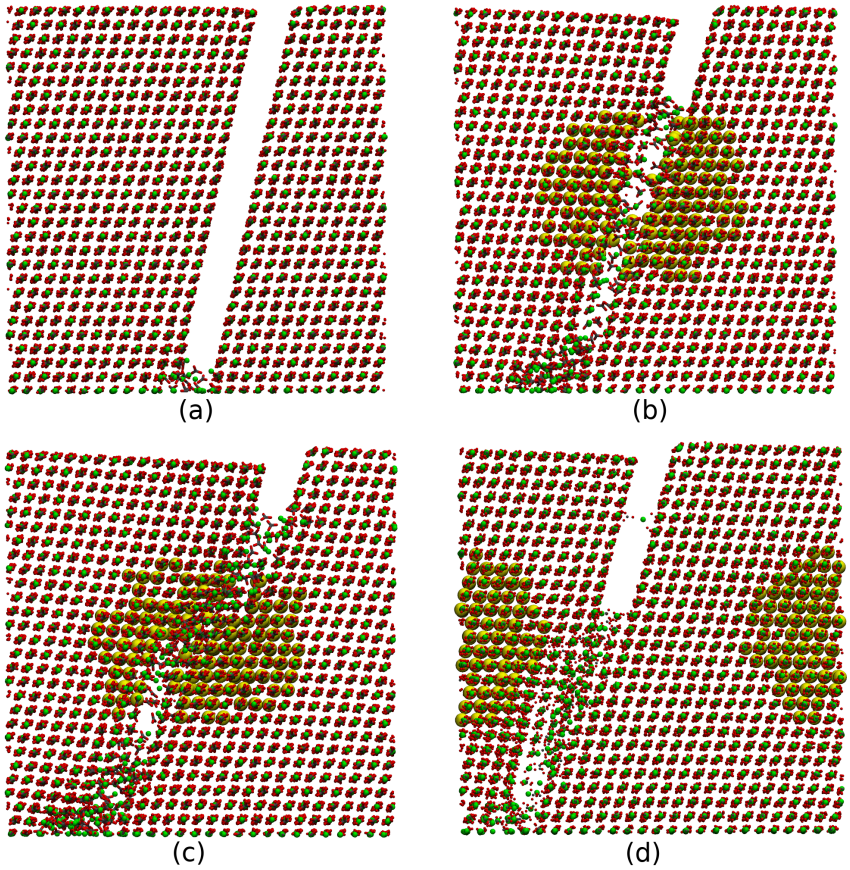

Figure 7.5: Snapshots of a crack propagating under different conditions including: crack propagating through pure calcite (a); crack bisecting the nanoprecipitate (b); crack propagating slightly off-centre of the nanoprecipitate (c); crack propagating around the nanoprecipitate (d). Here, the Mg ions have been magnified for easier visualisation.

posed a toughening mechanism whereby cracks are deflected on a macroscale due to variations in the material density [3]. Our observations suggest that a nanoscale deflection may also be an important part of the story.

It is worth emphasising that the stress and strain at the point of crack propagation is only increased when the plane of the crack directly bisects the nanoprecipitate, as can be observed in Figure 7.6. This can be explained by the results in Figure 7.4, which demonstrate that different regions of the crystal will be under different stresses. As a crack propagates, the stress opposing propagation will be the stress perpendicular to the direction of propagation. If the crack propagates towards the precipitate, then it will encounter the compressive tangential stress observed in Figure 7.4, whereas if the crack propagates around the precipitate, the relevant component of the stress tensor field becomes the radial component. This is why the stress at the point of propagation of the crystal is only increased when the crack bisects the nanoprecipitate. The difference of 223 MPa observed in Figure

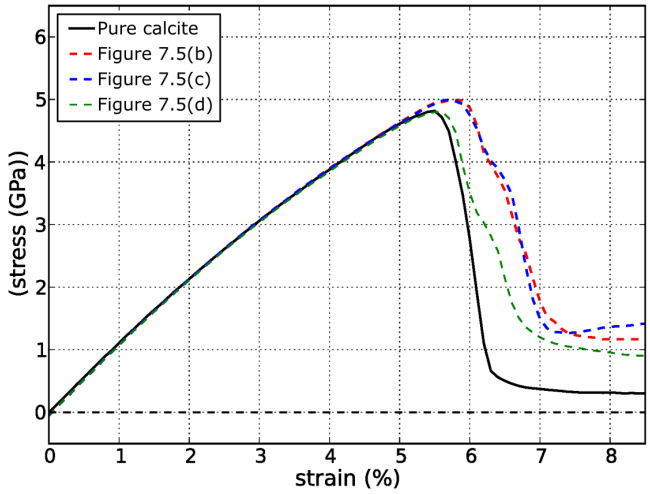

Figure 7.6: Stress-strain plots averaged over 10 simulations. An insignificant difference with respect to pure calcite is seen when the crack is far from the nanoprecipitate (d), but the difference becomes more apparent when the plane of the crack directly bisects the nanoprecipitate (b) and (c). In each case, the residual stress is higher when a nanoprecipitate is present.

7.6 is of the same order of magnitude as the tangential stress observed in the host matrix in Figure 7.4, indicating that the prestressing of the host matrix does contribute to the increased fracture toughness. This relies on the assumption that the crack approaches the nanoprecipitate, although it can be assumed that, given the high precipitate density, a propagating crack through a real crystal would, at some point, approach a nanoprecipitate.

7.3 Conclusions

The results of lattice distortion found from simulations have replicated the qualitative effects seen experimentally, and even reproduced the correct order of magnitude, despite the assumptions made in approximating the PGMA-Au nanoparticle as a vacuum. This leads to three conclusions: firstly, the novel combination of expansion and contraction observed experimentally in the case of PGMA-Au nanoparticles can be understood in terms of the mechanics of the interface of the hole left behind by the nanoparticles; secondly, the nanoparticles are sufficiently weakly interacting and fluid that their presence can be considered irrelevant and approximated as a vacuum; thirdly, the hole left behind by the nanoparticle is likely to take a rhombohedral morphology.

In the second section, we used molecular dynamics to investigate the effects of magnesium-rich nanoprecipitates on the stress distribution and crack propagation in calcite domains. Such nanoprecipitates have been identified in calcitic microlenses within the skeleton of the brittlestar Ophiocoma wendtii, and are observed to increase the fracture toughness in comparison to non-

biogenic calcite. Our simulations provide unique insight into the atomistic mechanisms responsible for the increased toughness. We have identified three possible sources. Firstly, the small size of the magnesium ions compared to calcium, together with the coherent nature of the nanoprecipitates, induces a compressive stress in the surrounding matrix, which counteracts the applied tensile stress during crack propagation, similar to the toughening of tempered glass and prestressed concrete. This mechanism was suggested by Polishchuk et al. who also estimated the magnitude of this effect using continuum mechanics. Here, we were able to calculate the components of the stress tensor at atomic resolution and show that, although the hydrostatic stress is relatively homogeneous throughout the matrix, the radial and tangential components show separated regions of tensile and compressive stress. As well as providing an explanation for why crack propagation is inhibited only when the crack approaches the nanoprecipitate, the inhomogeneous stress distribution is significant because it induces deflections of the cracks away from the lowest energy surfaces. This second source contributes to increased energy adsorption during fracture and, consequently, increased fracture toughness. The third contribution to increased fracture toughness identified by the simulations is the mechanism by which the cracks propagate through the nanoprecipitates. The highly inhomogeneous stress distribution within the precipitates results in strongly disordered fracture surfaces which form bridges between the crack surfaces that support residual stress at the termination of the simulation. Such crack bridges would absorb energy during fracture and may be a significant contribution of the nanoprecipitates to increased toughness. In summary, our atomistic simulations of stress and fracture in calcite with embedded Mg-rich nanoprecipitates have identified novel mechanisms by which such nanoprecipitates may increase the toughness of calcite.

8 NON-MARKOVIANITY: ORIGINS AND IMPACT ON RARE EVENT SAMPLING

As discussed in Chapter 3, a Markovian system is one where the future configurations of the system depend only on the conditions of the system at the present time. A non-Markovian system will also be dependent on not just the system's current configuration, but the configurations at previous times. It is well established that non-Markovianity can have a significant effect on diffusion [147, 148]. However, it is unclear exactly how these non-Markovian effects on kinetics may affect the crossing rates of free energy barriers. Central to this is the concept of the reaction coordinate, which reduces the $3N$-dimensional kinetics of a system as a few-dimensional process over a potential of mean force (PMF) as discussed in Chapter 3. A physical reaction, for example the dissolution of an adsorbate from a calcite kink site we encountered in chapters 5 and 6, can be reduced to the transition along one or several reaction coordinates. In the context of rare event sampling in simulations, the reaction coordinate, or coordinates, must be chosen such that the ergodicity of the simulation is allowed under the influence of the bias potential. In other words, there must be no rare events other than those which can be described using the set of reaction coordinates. A classic example of this is the rotation angles ϕ and ψ of alanine dipeptide [149, 150]. When one angle is neglected, rare event sampling techniques such as metadynamics typically run into convergence issues.

As well as being an important consideration of rare event sampling, the choice of the reaction coordinate may affect calculations of barrier crossing rates. Analytical solutions to barrier crossing rates are typically dependent on a free energy barrier height. For example, in Transition State Theory (TST) [151], the free energy barrier crossing rate is only dependent on the height of the free energy barrier separating reactant and product states, as well as a diffusion constant. Another example is Classical Nucleation Theory (CNT), which gives a nucleation rate proportional to $\exp(-\Delta F^*/k_B T)$, where ΔF is the free energy cost at the top of the barrier separating the vapour and condensed phases. CNT is based on the original Becker-Döring equations [152], which assume a constant rate of passage along the N-coordinate determined by a set reaction rates dependent on the cluster size. While subtle, a crucial assumption is made in the Becker-Döring equations that the reaction rates are only dependent on the immediate cluster size, and that the previous history of the cluster size does not impact the reaction rates. In other words, the system is Markovian. It is important to note non-Markovianity can only exist in a coarse-grained system. While Newton's equations of motion are dependent only on the instantaneous positions, momenta and forces, the act of coarse graining the $3N$-dimensional Newtonian kinetics of the system onto a set of

reaction coordinates can invoke non-Markovianity. An intuitive example of non-Markovianity can arise from the nucleation example: a cluster growing sufficiently quickly may deplete the local environment of solutes, leading to a lower attachment rate to the cluster. Here, the reaction rates have become dependent on the history of the system, and the nucleation rate calculated from CNT will ultimately be incorrect. The crucial assumption here is that the kinetics along the reaction coordinate, in this case the cluster size N, are Markovian. This is also the case for TST, where the barrier crossing rate is only dependent on the height of the free energy barrier and a diffusion constant. The free energy landscape, including the barrier height, may be dependent on the choice of reaction coordinates. The calculated nucleation rate may therefore depend on the reaction coordinate choice, even though the system kinetics are fundamentally identical. The only assumption made in TST is that the kinetics along the reaction coordinate are Markovian. The free energy barrier crossing rate must also therefore be dependent on any non-Markovian kinetics which arise through coarse graining.

It is clear that coarse-graining yields non-Markovian kinetics to a degree dependent on the reaction coordinate choice. So far, however, the fundamental origins of non-Markovianity in coarse-graining have not been characterised. In this chapter, we aim to answer the following questions: first, can introducing non-Markovian kinetics to a system correct free energy barrier crossing rates; second, how can we identify a poor reaction coordinate choice from rare event sampling simulations? We begin by calculating transition rates for a series of discrete processes, and introduce corrections by deriving non-Markovian transition rates. We then proceed to apply rare event sampling techniques to a two dimensional Langevin process and develop the means to identify poor reaction coordinate choices from the output of simulations.

8.1 Two-dimensional toy free energy surface

In this chapter, we are interested in working with a toy free energy surface in which we can use a single parameter to tune the appropriateness of a single reaction coordinate. For simplicity, we considered a two-dimensional toy free energy surface, and took a single component of the free energy surface as the reaction coordinate. An example two-dimensional toy free energy surface is the Müller-Brown potential [2], which contains two saddle points and three free energy wells. However, for the sake of simplicity, we would prefer a free energy surface which is easily modified to change the suitability of the reaction coordinate. Instead, we used the two-dimensional potential, $U(x, y)$, which is treated as a Potential of Mean Force (PMF) throughout this chapter, and is given by:

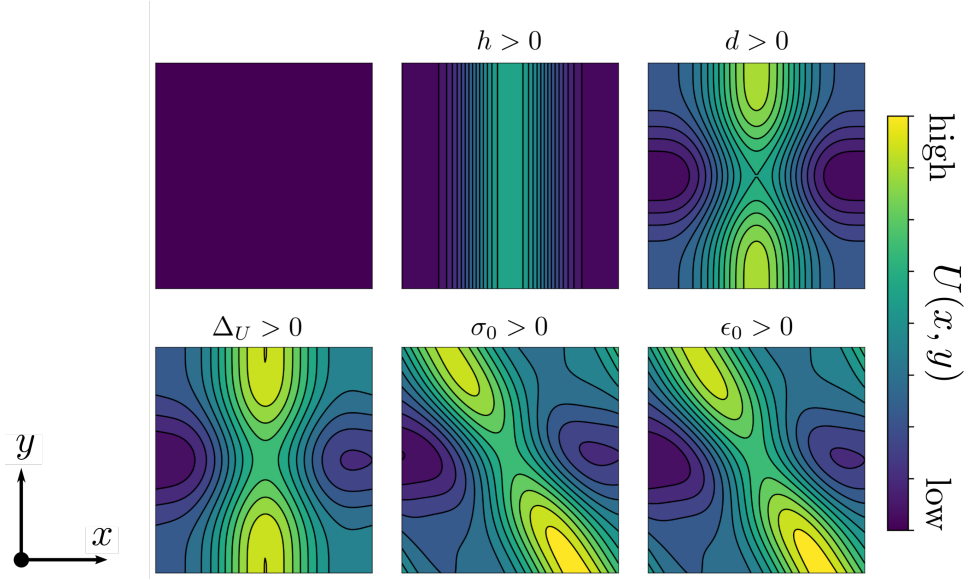

Figure 8.1: Demonstration of the effect of the parameters which control $U(x, y)$.

$$U(x,y) = \Delta_U \frac{x - x_{lo}}{L_x} + h \sin^n \left[\pi \left(\frac{x - x_{lo}}{L_x} + \sigma(y) \right) \right] + d \left(1 - \epsilon(y) \cos^m \left[\pi \left(\frac{y - y_{lo}}{L_y} - \frac{1}{2} \right) \right] \right) \tag{8.1}$$

where h represents the saddle point height and d represents the depth of the saddle point. The parameters m and n determine the sharpness of the free energy peak. x_{lo} and x_{hi} determine the boundaries of the potential in the x-direction. y_{lo} and y_{hi} determine the boundaries in the y-direction. L_x is defined as $L_x = x_{hi} - x_{lo}$. The same applies for L_y. The parameter Δ_U determines the relative height of the potential at x_{xlo} and x_{hi}. For example, a value of $\Delta_U = 3$ will set the potential to be 3 k_BT higher at x_{hi} than at x_{lo}. The functions $\sigma(y)$ and $\epsilon(y)$ are given below, where σ_0 and ϵ_0 are adjustable parameters.

$$\epsilon(y) = 1 + \epsilon_0 \left(\frac{y - y_{lo}}{L} - \frac{1}{2} \right) \qquad \sigma(y) = \sigma_0 \left(\frac{y - y_{lo}}{L} - \frac{1}{2} \right) \tag{8.2}$$

The function $\epsilon(y)$ makes the saddle point asymmetric along the y-direction. The function $\sigma(y)$ adds a skew to the barrier height, the extent of which is controlled with the parameter σ_0. Figure 8.1 shows a demonstration of how all the parameters in equation 8.1 affect the shape of the potential.

Throughout this chapter, we use the x-coordinate as a reaction coordinate. The reduced free energy as a function of the reaction coordinate x, $U(x)$, is defined by $U(x) = \int_y \exp[-U(x,y)/k_BT]dy$ (equation 3.90). Figure

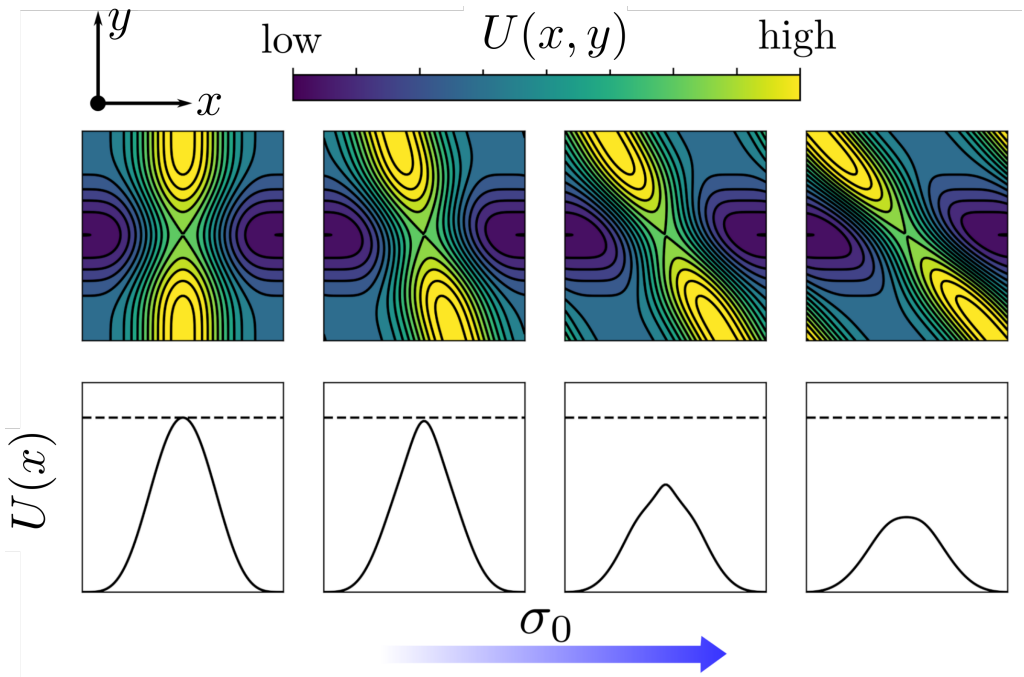

Figure 8.2: PMF as a function of x and y (top) and as a function of only x (bottom) for $\sigma_0 = 0.0, 0.3, 0.6, 0.9$ respectively (left to right). The dashed line on the right-hand side indicates perfect agreement between the saddle point height and the height of the coarse-grained barrier. The angle of the saddle point with respect to the x-direction is observed to have a large effect on the height of the coarse-grained free energy barrier.

8.2 shows a series of examples of $U(x, y)$ for different σ_0 parameters. The free energy as a function of the reaction coordinate x is also calculated using equation 3.90 and shown in f Figure 8.2. It can be seen that, as σ_0 is increased, the one-dimensional barrier height decreases, even though the true barrier height remains the same. The parameter σ_0 therefore tests the appropriateness of using the x-coordinate as the reaction coordinate. It is also clear from Figure 8.2 how methods such as transition state theory cease to apply as σ_0 is increased. A higher value of σ_0 corresponds to a poor reaction coordinate choice, whereas if $\sigma_0 \approx 0$, the use of the x-direction as a reaction coordinate is suitable.

Throughout this chapter, we set $k_B T$, m and the frictional force (γ in Chapter 3) to 1 (see equation 3.123). For the PMF given by equation 8.1, we used the following parameters: $h = 15$; $d = 7$; $\epsilon_0 = 0.2$; $n = m = 4$; $x_{lo} = 0$; $x_{hi} = 10$. Where we examine non-Markovian kinetics, we set Δ_U to zero for simplicity. Where we examine rare event sampling, we set Δ_U to -5. This was done to ensure no details were overlooked by taking the specific case where both wells are the same depth.

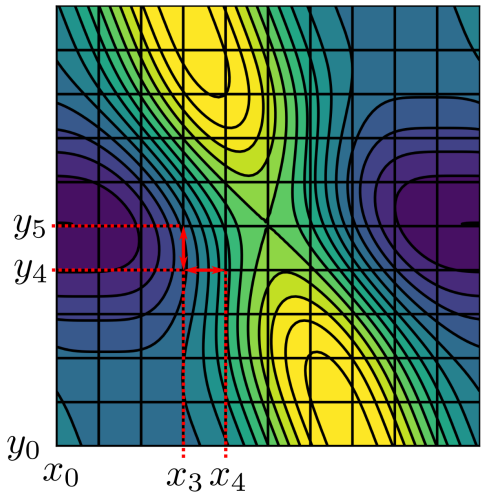

Figure 8.3: visualisation of the discretisation of the system as a series of transitions to nearest neighbours over the two-dimensional PMF. An example of the discrete x- and y-values is shown.

8.2 Emergence of non-Markovianity in discrete transitions

We began by considering discrete transitions over the two-dimensional PMF given by equation 8.1. Non-Markovianity is highly non-trivial to define in a continuous setting, and attempts to do so typically involve solving the generalised Langevin equation [153] for straightforward examples of memory kernels. For coarse-graining a two-dimensional Markovian PMF into a one-dimensional non-Markovian PMF, this approach becomes infeasible. However, constructing a series of discrete processes becomes far more palatable, as the number of past configurations associated with an initial configuration becomes a finite function of time. For example, for transitions only to nearest neighbours allowed, the number of possible previous configurations is equal to $(2^d)^m$, where d is the dimensionality of the system and m is the number of transitions we consider.

8.2.1 Model for discrete transitions along the two-dimensional PMF

In order to model discrete transitions over the PMF in equation 8.1, we define a grid of 11 by 11 points and consider transitions between these points. The x- and y-coordinates are therefore represented by

$$\begin{aligned} x &= x_i \qquad 0 \leq i \leq 10 \\ y &= y_j \qquad 0 \leq j \leq 10 \end{aligned} \tag{8.3}$$

where i and j are integers. Only transitions to nearest neighbours in the x- or y-directions were allowed, e.g. $x_i \rightarrow x_{i\pm1}$. Here, we apply rigid boundary conditions (i.e. do not allow transitions beyond the simulation boundaries), and only allow a single transition to take place at any given point. Since only transitions to nearest neighbours were allowed, a maximum of four reactions may occur at any given point (except at the boundaries where the number of possible reactions is smaller due to the rigid boundary conditions). We can therefore define a set of reaction rates, which we denote with $\Gamma_{x_i \rightarrow x_{i\pm1},\ y_j}$ representing a transition along the x-direction, and $\Gamma_{x_i,\ y_j \rightarrow y_{j\pm1}}$ representing a transition along the y-direction. In this chapter, we wish to treat $U(x,y)$ as a PMF. We therefore set the forward reaction rates ($x_i \rightarrow x_{i+1}$ and $y_j \rightarrow y_{j+1}$) to $1t_0^{-1}$, where t_0 is the time unit of the system (and is set to 1 throughout this chapter) and calculated the backward rates ($x_i \rightarrow x_{i-1}$ and $y_j \rightarrow y_{j-1}$) using detailed balance. We can therefore define a set of master equations as follows:

$$\begin{aligned} \Gamma_{x_i \rightarrow x_{i+1}, y_j} &= 1, \quad \Gamma_{x_i \rightarrow x_{i-1}, y_j} = \exp\left(-\frac{U(x_{i-1}, y_j) - U(x_i, y_j)}{k_B T}\right) \\ \Gamma_{x_i, y_j \rightarrow, y_{j+1}} &= 1, \quad \Gamma_{x_i, y_j \rightarrow, y_{j-1}} = \exp\left(-\frac{U(x_i, y_{j-1}) - U(x_i, y_j)}{k_B T}\right) \end{aligned} \tag{8.4}$$

where the backward rates are determined from the forward rates using detailed balance. Since we are interested in integrating this model of a finite number of discrete transitions over time, given a set of reactions, the algorithm we use to update the i and j indices is the Gillespie algorithm, which is the same algorithm as the one used to integrate KMC models over time (see Section 3.4.6).

8.2.2 Coarse-graining and non-Markovian transitions

Our intention here is to uncover barrier crossing rates for a coarse-grained PMF in both Marvkovian and non-Markovian settings. In other words, we want to calculate the corresponding average reaction rates in the x-direction and use them to run a one-dimensional simulation. Calculating the average reaction rate in a Markovian setting is straightforward. We can define $\Gamma_{x_i \rightarrow x_{i\pm1}}$ as the average reaction rate representing a transition in the x-direction only as a function of the x-coordinate, The values of $\Gamma_{x_i \rightarrow x_{i\pm1}}$ in a Markovian setting are obtained by integrating over the y-direction. In our discrete setting, the integral is replaced by a sum over all possible j values:

$$\Gamma_{x_i \rightarrow x_{i\pm1}} = \langle \Gamma_{x_i \rightarrow x_{i\pm1}, y_j} \rangle = \sum_j \mathcal{P}_{y_j|x_i} \Gamma_{x_i \rightarrow x_{i\pm1}, y_j} \tag{8.5}$$

where $\mathcal{P}_{y_j|x_i}$ is the probability of finding a y-position of y_j given that its x-value is equal to x_i. It should be noted here that $\sum_j \mathcal{P}_{y_j|x_i} \equiv 1$. While

the values of $\Gamma_{x_i \to x_{i\pm1}}$ are calculable analytically, we instead derive the value numerically by running a two-dimensional simulation, and calculating all $\Gamma_{x_i \to x_{i\pm1}}$ values by averaging the time taken for all $x_i \to x_{i\pm1}$ reactions to occur:

$$\Gamma_{x_i \to x_{i\pm1}} = \left\langle \mathcal{T}_{x_i \to x_{i\pm1}} \right\rangle^{-1} \tag{8.6}$$

where $\mathcal{T}_{x_i \to x_{i\pm1}}$ is the average time taken for the transition of $x_i \to x_{i+\pm1}$ to occur, which can be directly calculated during simulations. While equations 8.5 and 8.6 are both fairly intuitive, their equivalence is not immediately obvious. A proof of their equivalence can be found in Section A.1.1.

While the above method is applicable for a Markovian setting, the problem becomes more difficult when we wish to derive a set of non-Markovian master equations. If the system is completely Markovian, then $\Gamma_{x_i \to x_{i\pm1}}$ will only be dependent on the value of x_i. However, we can see in figure 8.2 that this cannot be true when σ_0 is non-zero. Instead, we assume that $\Gamma_{x_i \to x_{i\pm1}}$ is non-Markovian, and instead attempt to derive values for $\Gamma_{x_i \to x_{i\pm1}}$ based also on previous x_i-values. Methods do exist which generalise the Gillespie algorithm to allow for non-Markovian processes [154, 155]. However, these algorithms are difficult to implement when directly deriving memory from prior simulations. Instead, we define our own method in which which we calculate a series of reaction rates dependent on previous conditions. Consider Figure 8.4. When the transitions are limited to $x_i \to x_{i\pm1}$, there are a number of possible pathways leading to point x_3 at step n. The total number of pathways leading to point n (ignoring boundary conditions) is given by 2^m, where m is the number of previous transitions that we wish to remember. This number is the same for all x_i-values if we ignore boundary conditions.

With this in mind, a set of non-Markovian reaction rates can be derived by calculating separate transition rates *given* each possible set of previous configurations. Instead of deriving two transition rates at each n, a total of 2^m rates *given* previous pathways of length m can be derived. In a completely Markovian setting, the reaction rates will be identical given every previous pathway. In a non-Markovian setting, all 2^m rates are independent. For example, if we wish to take into account one previous configuration, we can define $\Gamma_{x_i \to x_{i\pm1}|x_{i-1}\to x_i}$ as the reaction rates for the $x_i \to x_{i\pm1}$ reactions, given that the previous reaction leading to x_i was an $x_{i-1} \to x_i$. We can also extend this to account for an arbitrary number of previous configurations e.g. the green path in Figure 8.4 represents a transition from x_3 at step n to x_4 at step $n+1$. The reaction rate is given by $\Gamma_{x_3 \to x_4|x_3\to x_2\to x_3}$. For the sake of efficiency, we define a series of coefficients, $\mathcal{M}$, representing each reaction rate. For example, if we wish to store all reaction rates given knowledge of two previous configurations, we can define $\mathcal{M}_{\alpha,\beta,i,\pm1} = \Gamma_{x_i \to x_{i\pm1}|x_{i-\alpha-\beta}\to x_{i-\alpha}\to x_i}$ for $\alpha, \beta, \gamma = \pm1$. Taking the green path in Figure 8.4 as an example, we have $\mathcal{M}_{-1,1,3,1} = \Gamma_{x_3 \to x_4|x_3\to x_2\to x_3}$. The number of indices subtracted by 2 gives the number of previous configurations remembered. For example, if we wished to

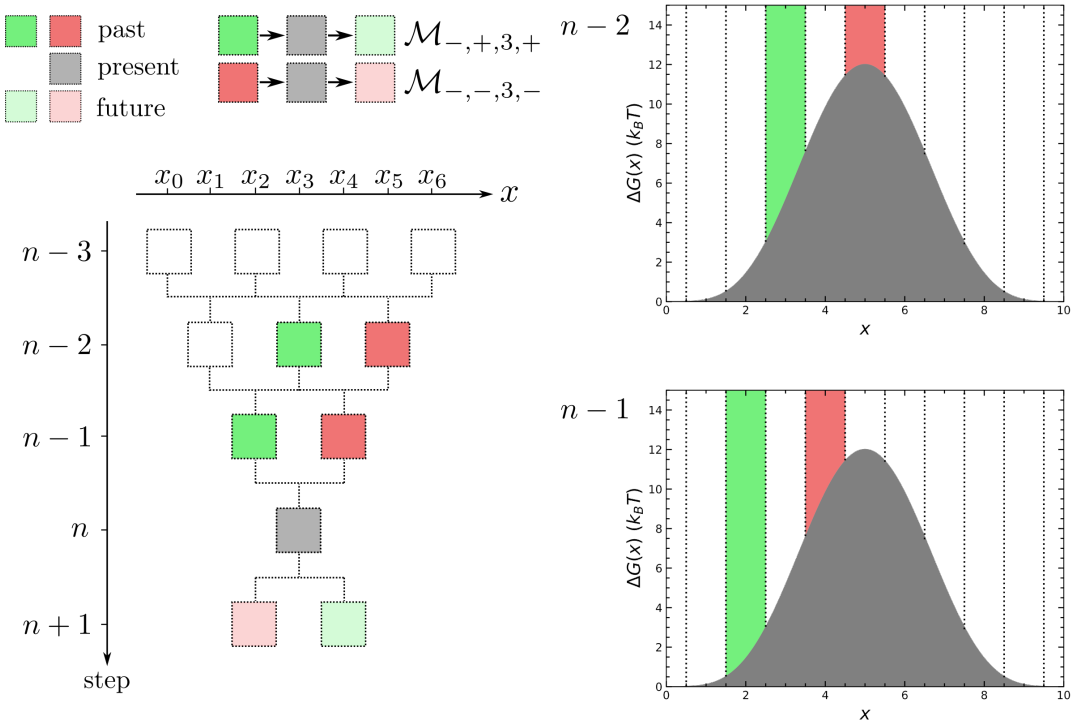

Figure 8.4: Diagram of all possible pathways leading to point x_3 at step n given only one-dimensional transitions to neighbouring points. Two such pathways are highlighted in red and green.

forget the configuration at step $n-2$, we could write $\mathcal{M}_{1,3,1} = \Gamma_{x_3 \to x_4 | x_2 \to x_3}$.

For any number of coefficients remembered, the coefficients of $\mathcal{M}$ can be calculated numerically from a stationary simulation of diffusion over the two-dimensional PMF by averaging the time taken for each transition to occur given each possible pathway. In this section, we derive reaction rates with a memory of up to four previous transitions as follows:

$$\begin{aligned}
\mathcal{M}_{\alpha,i,\pm 1} &= \left\langle \mathcal{T}_{x_i \to x_{i\pm 1} | x_{i-\alpha \to x_i}} \right\rangle^{-1} \\
\mathcal{M}_{\beta,\alpha,i,\pm 1} &= \left\langle \mathcal{T}_{x_i \to x_{i\pm 1} | x_{i-\alpha-\beta \to x_{i-\alpha \to x_i}}} \right\rangle^{-1} \\
\mathcal{M}_{\gamma,\beta,\alpha,i,\pm 1} &= \left\langle \mathcal{T}_{x_i \to x_{i\pm 1} | x_{i-\alpha-\beta-\gamma \to x_{i-\alpha-\beta \to x_{i-\alpha \to x_i}}}} \right\rangle^{-1} \\
\mathcal{M}_{\delta,\gamma,\beta,\alpha,i,\pm 1} &= \left\langle \mathcal{T}_{x_i \to x_{i\pm 1} | x_{i-\delta-\alpha-\beta-\gamma \to x_{i-\alpha-\beta-\gamma \to x_{i-\alpha-\beta \to x_{i-\alpha \to x_i}}}}} \right\rangle^{-1}
\end{aligned} \tag{8.7}$$

where $\alpha, \beta, \gamma, \delta = \pm 1$ and $\mathcal{T}$ are the transition times analogous to equation 8.6 but with separate transition times given a previous pathway. Note that equation 8.7 is effectively a non-Markovian analogue of equation 8.6.

We examined the free energy barrier crossing rate in two settings: first, we examined the crossing rate for a simulation over the two-dimensional PMF with transitions allowed in both x- and y-directions; second, we examine the

crossing rate in a one-dimensional coarse-grained space in which the transition rates are given by equation 8.6 in the Markovian case, and equation 8.7 in the non-Markovian case. If the calculated barrier crossing rates are similar, the coarse-graining can be considered accurate. If they are not similar, however, the coarse graining has neglected crucial detail.

We began by carrying out two-dimensional simulations, where we calculate the correct barrier crossing rate. The imaginary particle, which represents the current configuration, begins at $(x_{lo}, y_{hi} + y_{lo}/2)$ which corresponds to one of the free energy minima. Every time the particle's x-position reaches x_{hi}, it is returned to $(x_{lo}, y_{hi} + y_{lo}/2)$ and a crossing is recorded. We then ran another simulation in two dimensions where we use equations 8.6 and 8.7 to derive Markovian and non-Markovian transition rates along the x-direction. These reaction rates are then used to simulate barrier crossings in one dimension. In this instance, we calculated crossing rates by returning the particle to x_{lo} as soon as it reaches x_{hi}. The processes, along with the number of iterations used in each case, is summarised below:

- Calculate the crossing rate from simulations in two dimensions by allowing the particle to return to x_0 when it reaches x_{10}: 2×10^8 iterations.
- Run a simulation in two-dimensions without enabling crossings (i.e. not returning the particle to x_0 when it reaches x_{10}). As the simulation runs, store the previous four configurations, and use the reaction times to calculate $\Gamma_{x_i \to x_{i\pm1}}$ (Markovian) and all $\mathcal{M}_{\alpha,i,\pm1}$, $\mathcal{M}_{\beta,\alpha,i,\pm1}$, $\mathcal{M}_{\gamma,\beta,\alpha,i,\pm1}$ and $\mathcal{M}_{\delta,\gamma,\beta,\alpha,i,\pm1}$ coefficients (non-Markovian): 6×10^8 iterations.
- Run a simulation in one-dimension using the derived rates and coefficients as reaction rates. Calculate the crossing rate by allowing the particle to return to x_0 when it reaches x_{10}, and compare with the first, two-dimensional simulation: 1×10^8 iterations.

8.2.3 Results

The process detailed above was carried out for a series of random walks over the potential given by equation 8.1, using variations of σ_0 between 0 and 0.8 in increments of 0.1. The non-Markovianity of the system varied between a memory of one one-dimensional transition to a memory of four one-dimensional transitions in increments of one.

Figure 8.5 shows the results of the ratio between coarse-grained crossing rates and true crossing rates. The dashed line at 1.0 indicates perfect agreement. Turning our attention first to the Markovian limit, it is apparent that an increase in σ_0 causes the ratio of crossing rates to diverge, i.e. the one-dimensional random walk produces a far larger barrier crossing rate than the correct crossing rate obtained from a random walk over the two-dimensional PMF. This can be easily understood with reference to Figure 8.2. The skewing of the saddle point results in a random walk over a one-dimensional

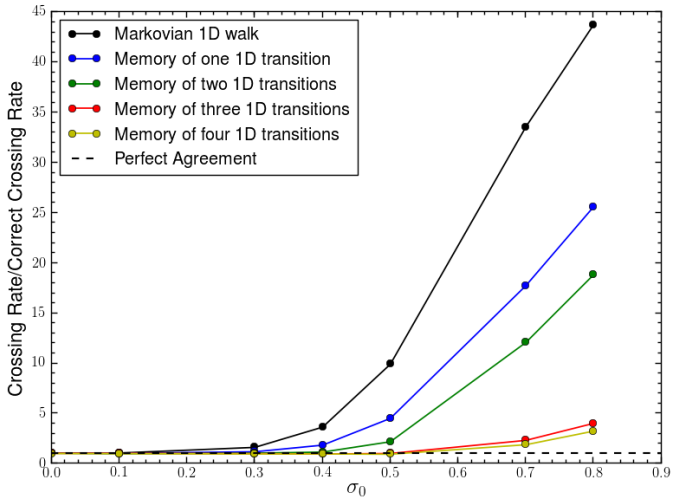

Figure 8.5: Ratio of coarse-grained crossing rates using reaction rate derived using the method presented in Section 8.2.2. Here, the process is repeated for a memory of zero, one, two, three or four previous transitions along the x-coordinate. The dashed line indicated perfect agreement between the (correct) two-dimensional simulations and the (coarse-grained) one-dimensional simulations.

PMF with a smaller peak. By deriving the average reaction rates over the x-direction, we effectively coarse-grain the potential over which the particle travels to produce a walk over a PMF with a smaller peak.

It is clear that coarse graining under certain conditions does not preserve the kinetics in the Markovian limit. However, Figure 8.5 also shows the ratio of crossing rates for non-Markovian walks with a memory ranging from one previous transition to four previous transitions. Under these conditions, the crossing rate ratio is less sensitive to σ_0. The sensitivity to σ_0 is reduced for every increased number of transitions memorised. With a memory of three or four transitions, the crossing rate becomes over an order of magnitude smaller than the Markovian limit. It is clear that imposing non-Markovian kinetics on the coarse-grained system produces barrier crossing rates more similar to those of the original system. To understand why this is, consider Figure 8.6. If the particle resides at x_4, having been initiated at x_0, it is most likely to reside below the saddle point (i.e. have a lower y-coordinate value than the saddle point). The most likely transition from x_4 to x_5 will therefore occur below the saddle point. However, the most likely transition from x_5 to x_6 will occur above the saddle point. However, a large barrier in the y-direction exists separating these two reactions. Without memory, the one-dimensional transition rates do not take into account this separation, and the crossing rate is significantly increased, due to the favouring of transitions occurring on opposite sides of the saddle point. However, consider the case where reaction rates are also dependent on the previous configuration. A transition from x_4 to x_5 is again most likely to occur below the saddle point. However, the transition from x_5 to x_6 *given* a previous transition from x_4 to x_5 has

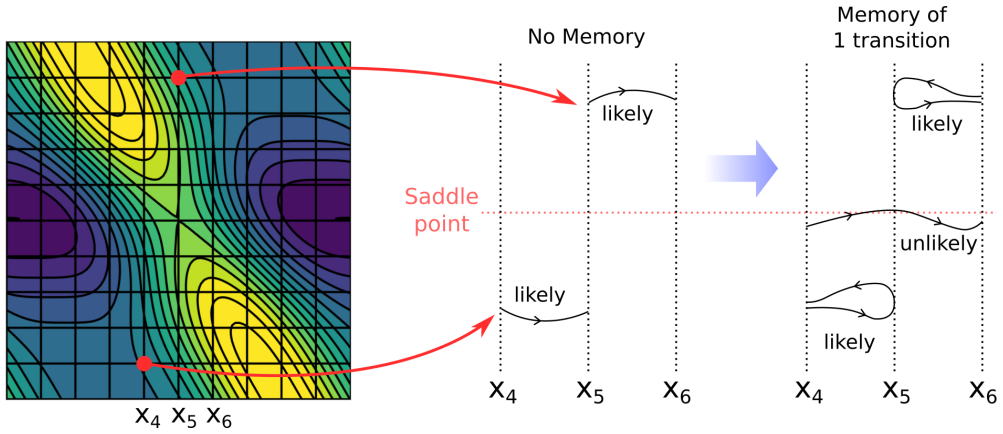

Figure 8.6: Schematic demonstrating why invoking memory gives rise to improved kinetics. If the coarse-grained system is treated as Markovian, the main contribution to the reaction rates for neighbouring sites may be over regions either side of the saddle point. However, if memory is added, the reaction rates take into account the increased likelihood that a transition from x_4 to x_5 will result in a transition back to x_4.

already indicated that the particle is most likely residing below the saddle point. Based on the information supplied by the previous configuration, the system now knows that it is far more likely that the particle will return to x_4, rather than progress to x_6. This is reflected in the transition rates. If the skew of the saddle point (σ_0) is increased further, the transition from x_5 to x_6 may eventually also be most likely to occur below the saddle point. At this point, a memory of two or more transitions will recover the kinetics to a more accurate degree.

Figure 8.5 has demonstrated how non-Markovian kinetics can arise, even in straightforward cases such as our toy model, from the neglect of detail of configuration along dimensions of phase space normal to the reaction coordinate. A Markovian description of the kinetics is only valid when as little detail as possible is neglected, such as when $\sigma_0 \approx 0$ in our example. In the absence of non-Markovian corrections to the dynamics, barrier crossing rates may be highly sensitive to our choice of reaction coordinates.

8.3 Characterising the failures of rare event sampling methods owing to unsuitable reaction coordinates

As discussed above, it is well established that rare event sampling techniques require a careful selection of reaction coordinates in order to accurately calculate free energy functions. What is not established is how exactly a poor reaction coordinate manifests in the convergence of the free energy function,

and how to use the data from rare event sampling to determine whether the reaction coordinates used are appropriate. Fortunately, our two-dimensional model offers a framework through which these details may be established. By knowing when a reaction coordinate is a poor choice, the model can be used to compare the output of simulations between high and low values of σ_0. Throughout this section, we use the parameters listed previously, and differentiate between an appropriate and inappropriate reaction choice by setting σ_0 to either 0 or 0.5. We proceed to use various rare event sampling techniques with the x-axis as the reaction coordinate in an attempt to replicate the coarse-grained one-dimensional PMF.

8.3.1 Kinetics over the two-dimensional PMF in the continuous limit

In order to run rare event sampling techniques on our two-dimensional PMF, we must design a continuous analogue of the model described in Section 8.2.1. In order to do so, a two-dimensional Langevin process, with equation 8.1 supplying the potential, was used. In order to allow the process to be evaluated numerically, we solved for an Itô process, by specifying a finite time interval δt and using the techniques discussed in Section 3.4.3. The noise on both dimensions is therefore given by a Gaussian-weighted random number with a mean of zero and a variance of δt as required (see Chapter 3). The application of an Itô process under the potential $U(x, y)$ produces a canonical distribution as demonstrated in Section 3.4.5. The potential $U(x, y)$ can therefore be treated as a PMF as is done in the previous section. We therefore refer to $U(x)$ as $\Delta G(x)$ throughout this Section to clarify its purpose as a PMF.

8.3.2 Rare event sampling techniques

In order to characterise the breakdown of numerical rare event sampling techniques in the absence of suitable reaction coordinates, we use several techniques in an attempt to reconstruct the coarse-grained one-dimensional free energy surface for different values of σ_0. We apply these techniques to our two-dimensional Langevin process. In every instance, we choose the x-coordinate as our reaction coordinate. When $\sigma_0 \neq 0$, the skew inhibits the suitability of the x-coordinate as a reaction coordinate.

The first, rare event sampling technique we consider is metadynamics. As metadynamics is used extensively in chapters 5 and 6, it makes a good starting point for this section. With reference to equation 3.95, we set $W_0 = 1\ k_BT$, σ to 0.2 and τ to 200 time-steps in this chapter. We also apply well-tempering, with the bias-factor set to 12 (the height of the saddle point). The simulation consisted of a brief equilibration of 5×10^4 timesteps and a production run totalling 3×10^7 timesteps was carried out.

Our second choice of rare event sampling technique is umbrella sampling. Here, we follow the method discussed in Section 3.3.4. We use a series of harmonic springs, each with a spring constant of 50 k_BT to tether the particle, and the spacing between windows (see chapter 3) was set at 0.12. For every new window, a random number in the y-direction was chosen, and a brief equilibration run of 4×10^3 timesteps was carried out, and a production run totalling 3×10^5 timesteps was then carried out.

Our third choice of technique is Mean Force Integration (MFI). Here, we use the same simulation setup as for the well-tempered metadynamics simulations, and post-process the results of the simulations using the MFI algorithm (see Section 3.3.7). We also use a kernel width (h) of 0.2 (see Section 3.3.8), and output the collective variable (x) every 20 timesteps (so 10 kernels are added for every Gaussian added to the bias).

Our final choice of rare event sampling method is the Jarzynski approach. As discussed in Section 3.3.8, we tether the particle to a moving harmonic spring with a spring constant of 50 k_BT. The spring moved along the x-direction (the reaction coordinate) between $x = -2$ and $x = 12$. The particle started at either $x = -2$ or $x = 12$ and, after an equilibration process lasting 2×10^4 time-steps, travelled to the other x-coordinate over a timespan of 2×10^5 time-steps. This process was repeated 100 times.

8.3.3 Breakdown of rare event sampling methods

Before looking into each method, it is useful to determine exactly what happens when rare event sampling methods are applied to simulations when the reaction coordinate is inappropriate. To resolve this, we run 50 separate simulations using metadynamics, umbrella sampling and mean force integration (we do not consider the Jarzynski method here for reasons discussed later). We use the parameters discussed in Section 8.3.2 for each simulation. We output the final derived free energy profiles in each case and average them as shown in Figure 8.7. The standard deviations are also shown, as are the results from each separate simulation. We do this only for the poor reaction coordinate case ($\sigma_0 = 0.5$). In the limit of $\sigma_0 \approx 0$, all free energy methods are unsurprisingly able to reproduce the correct free energy surface. We do not show the results of this though, as they are of no interest.

From Figure 8.7, it is clear that all methods fail to replicate the correct one-dimensional free energy surface. The reasons for this are as follows: a statistically accurate method will reproduce a probability density as given by the correct one-dimensional free energy surface. However, when a poor reaction coordinate is used, the height of the one-dimensional barrier is lower than the true height of the saddle point (see Figure 8.2). This means that a reproduction of the correct free energy surface will not yield a bias with enough force to drive the particle through the saddle point. This results in a final free energy surface which is inaccurate, and varies roughly around

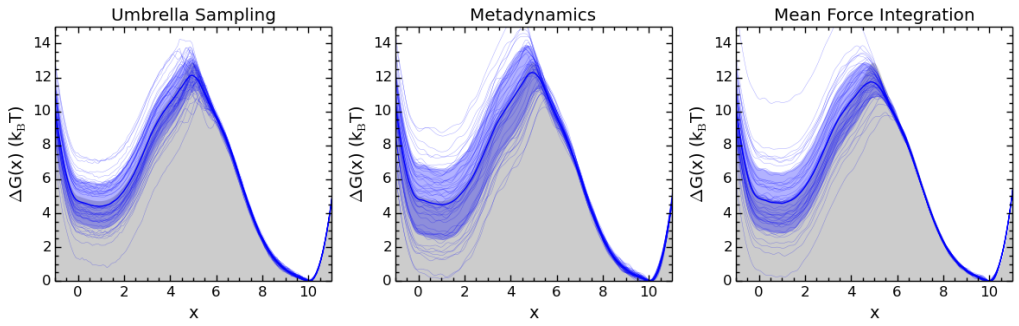

Figure 8.7: Results of 50 umbrella sampling, metadynamics, and mean force integration simulations for our two-dimensional PMF with $\sigma_0 = 0.5$, with the x-coordinate used as the reaction coordinate. The grey shaded region represents the correct PMF, obtained by integrating equation 8.1 over y using equation 3.90. The narrow, blue lines show the individual PMFs derived from every separate simulation, and the thicker blue line indicates the PMF averaged across all simulations. These results of each simulation, along with the mean, are used to calculate a standard deviation. This standard deviation is shown in the blue shaded region. The standard deviation calculated here is not to be confused with the sampling error of each separate simulation. Rather, it gives an indication of the degree of convergence reached for each rare event sampling technique.

the correct free energy surface with a standard deviation dependent on the rare event sampling method of choice. This is due to either an asymmetric accumulation of Gaussians on one side of the well (metadynamics), or the rapid transition from high-energy to low-energy states that does not appear to affect the probability density (umbrella sampling). We further elaborate on these details in the following sections.

It is important to note here that only the output of a single simulation will be available in practice. Determining the validity of the reaction coordinate choice is therefore a decision which must be made from the data of a single simulation. We therefore proceeded to examine the output and errors of single simulations in order to determine whether it is possible to assess the validity of the reaction coordinate. We also wish to be able to quantify the error that arises from using a poor reaction coordinate and determine whether it represents the true error of the simulation arising from a poor reaction coordinate choice.

A closer look at umbrella sampling

Of the three methods used, the standard deviations in Figure 8.7 show that umbrella sampling is the least susceptible to the errors that arise from using poor reaction coordinates. This may appear to indicate that umbrella sampling is a better choice than metadynamics or mean force integration. However, there are several caveats to this. Figure 8.8 shows a single simulation where the derived PMF produced differs significantly from the correct

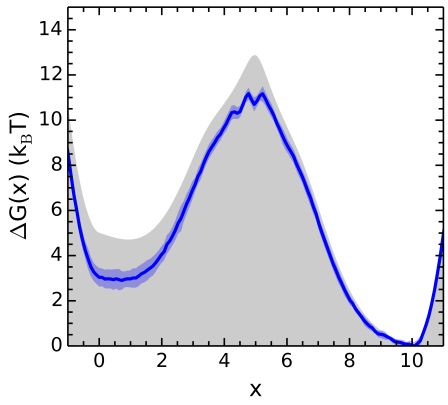

Figure 8.8: Result of a single umbrella sampling simulation with the sampling error highlighted. The simulation has produced an inaccurate free energy surface, which is not accounted for by the sampling error.

PMF. Here, the sampling error is calculated using Monte Carlo (MC) bootstrapping, in which subsets of each histogram are sampled at random. For a true estimate of the error, an agreement with the standard deviation shown in Figure 8.7 should be observed. This is not what is observed here. The error obtained from MC bootstrapping is far smaller than the true error, and does not account for the difference between the derived and true PMF.

The reason for the inaccuracy observed in Figure 8.8 is fairly straightforward: as different regions of phase space are sampled in the x-direction, the particle may be on either side of the saddle point, due to the skew in the y-direction. The particle may spend the entire simulation time on one side of the barrier in one window, only for its neighbour to spend the entire simulation time on the other side of the barrier. The sampling error obtained from MC bootstrapping only samples subsets within the final data, and therefore does not take into account the errors the arise due to random barrier crossings. This is an extension of an already established flaw of umbrella sampling: it does not directly simulate barrier crossing events. This reveals an important caveat to umbrella sampling in general: when a poor reaction coordinate is used, there is no obvious way of knowing whether a suitable reaction coordinate is in use. Despite the fact that the simulation appears to converge and produce a PMF, the PMF it produces is incorrect. Since there is no way of immediately knowing the suitability of the reaction coordinate, a major flaw of umbrella sampling is that, when it goes wrong, there is no way of knowing it has done so.

A further caveat to umbrella sampling is that the choice of sampling along regions of phase space perpendicular to the reaction coordinate (in our case the y-direction) may also affect the final result. In our instance, we sample initial y-configurations by choosing a random point in the y-direction and equilibrating, something we expect to be similar to a more practical simu-

lation set-up. An alternative would be to not re-sample in the y-direction for each window, and instead drag the particle along the x-coordinate, something we also attempt. The results of this are shown in Figure A8, in which plots similar to Figure 8.7 are shown. Unsurprisingly, an asymmetry in the derived PMFs arises based on the starting configuration. We elaborate on this further in Section 8.3.3 in the context of the Jarzynski method.

A closer look at well-tempered metadynamics

Despite being the least susceptible to the effects of using poor reaction coordinates, umbrella sampling has been shown to have a risk attributed to it, due to the lack of knowledge of whether the reaction coordinate is suitable. Well-tempered metadynamics has been shown to have a larger degree of error associated with it. Nevertheless, there is one distinct advantage of metadynamics. Figure 8.9 tracks the shape of the returned PMF from the bias as a function of the number of Gaussian hills deposited. The PMF is again normalised so that the height of the right well is correct, meaning the height of the left well can be easily tracked, as is done in Figure 8.9. It is clear that the height is fluctuating wildly as the simulation progresses. The reason for this is as follows: the reaction coordinate is not able to directly force a barrier crossing event to take place. Instead, the particle spends an additional amount of time on one side of the barrier, resulting in an accumulation of Gaussian hills on one side of the barrier. When a crossing does take place, the Gaussian hills start to accumulate on the other side of the barrier. This is what causes the changes in direction along the y-axis in Figure 8.9, each change in direction corresponds to a crossing. The snapshots in Figure 8.9 demonstrate how convergence is not observed, due to the fluctuating shape of the returned PMF.

An important advantage of metadynamics arises from Figure 8.9. It is clear from the output of the hills that the simulation is not converging due to a poor choice of reaction coordinates. Such information is not accessible in umbrella sampling. This is an advantage of metadynamics that arises from the ability to sample barrier crossings, even in certain examples of poor reaction coordinates, such as our example. It is worth noting, that a reaction coordinate can be considered inappropriate in more ways than the having a skewed saddle point. Nevertheless, as long as the reaction coordinate does allow a crossing event to take place, its suitability can be easily assessed from a single metadynamics simulation.

A closer look at mean force integration

One important use of MFI is its ability to produce an error which accounts for the evolving bias potential. By block averaging the numerator and denominator components of equation 3.104, an estimate for the standard error for each component can be calculated. These errors can be propagated when

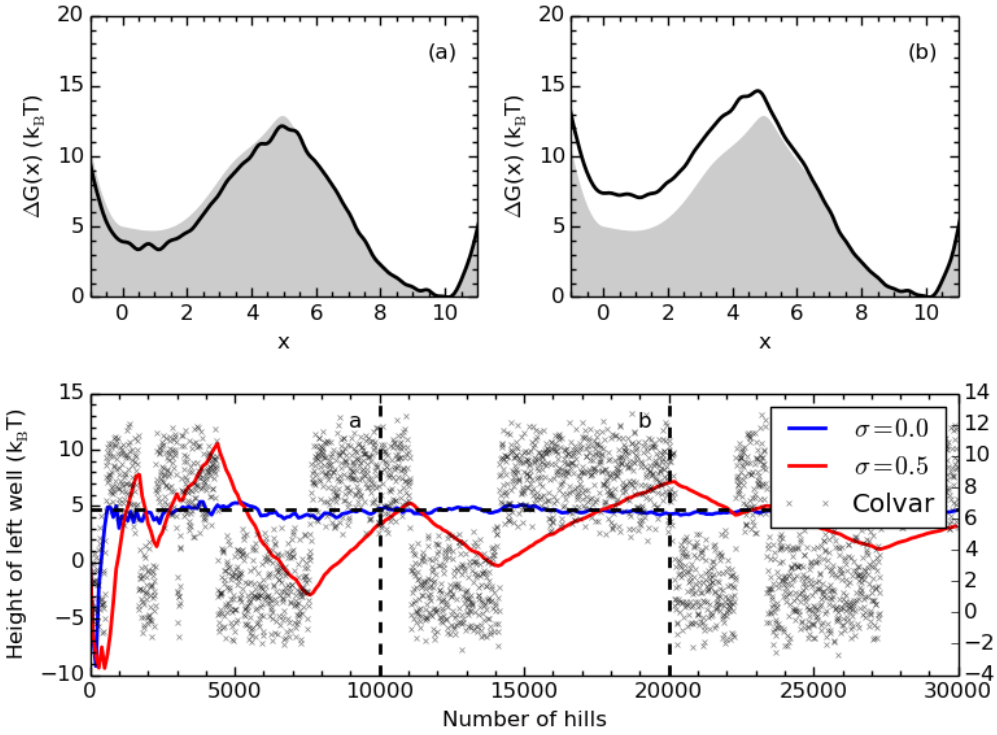

Figure 8.9: Result of a single metadynamics simulation detailing the height of the left-hand well as a function of the number of hills accumulated. The output of the collective variable x (Colvar) is also shown, as are the outputs of the bias at two time intervals, labelled a and b.

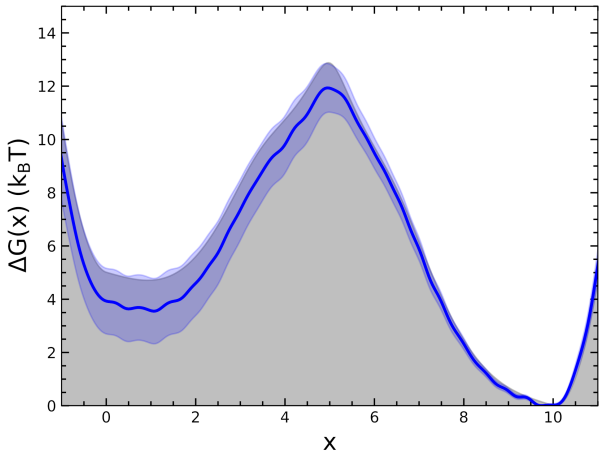

Figure 8.10: Result of a single MFI simulation with the sampling error highlighted. The simulation has produced an inaccurate free energy surface, although the error calculated by block averaging has accounted for this error.

the mean force is calculated and integrated to produce the PMF. Figure 8.10 shows an example of an MFI simulation with the estimated standard

error. Although the correct PMF has not been produced, the standard error calculated from MFI are large enough to account for the true error of the simulation. MFI therefore offers a distinct advantage over umbrella sampling and metadynamics in the sense that the error of the simulation can be quantified and the level of convergence assessed to a degree not possible with the sampling error obtained from bootstrapping.

A closer look at the Jarzynski approach

So far, the Jarzynski approach has not been discussed in terms of its failures when using a poor reaction coordinate. The reasons for this become clear upon the examination of Figure 8.11: instead of many simulations producing PMFs which fluctuate about the correct values, the Jarzynski approach produces simulations which average about an incorrect shape. This incorrect shape is dependent on the direction of travel of the particle. The reason for this discrepancy in direction is fairly intuitive: as the particle is dragged in the x-direction, it is free to move in the y-direction. Therefore, as the particle is forced towards the barrier, it travels in the y-direction away from the saddle point towards a region of lower free energy (see Figure 8.2). Rather than crossing the barrier at the saddle point, the particle instead crosses at another region of the barrier where the free energy is higher. This is why the free energy continues to increase beyond the saddle point. Furthermore, after the barrier is crossed, the particle is able to traverse along the y-direction to a region of lower free energy without moving in the x-direction. This is not picked up by the work done in the x-direction, therefore a lot of information is lost after the barrier is crossed.

Figure 8.11 shows that the standard errors produced by the simulation are larger than those produced by umbrella sampling. However, these standard errors are still very limited in accounting for the difference between estimated and exact free energy surfaces. Therefore, much like umbrella sampling, the Jarzynski approach must be used with caution. One approach for validating the derived PMF could be to run the simulation in reverse i.e. start from the other side of the barrier. However, the main advantage of the Jarzynski approach is that it can be used for processes which cannot be easily reversed.

8.3.4 Revisiting Chapters 5 and 6: does dehydration matter after all?

In Chapters 5 and 6, we derived free energy landscapes for solutes adsorbing to calcite surfaces using a combination of reaction coordinates describing the positions components of a solute and, in some cases, the hydration of the lattice sites. The validity of the combination of reaction coordinates applied is frequently discussed. For example, the need to dehydrate cations, either in solution (Chapter 5) or as part of a lattice site (Chapters 5 and 6), forms

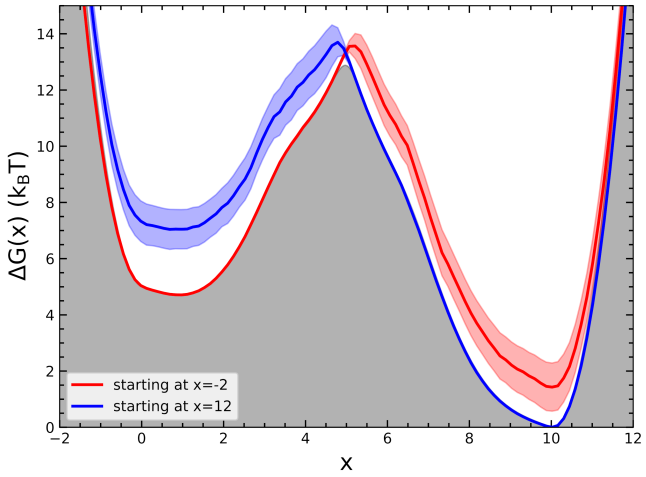

Figure 8.11: Results of two simulations applying the Jarzynski approach from two separate starting points and directions of travel. The standard errors for each simulation are also shown. The two free energy surfaces are shifted in the y-direction such that they correspond with the well that they start in. This is not to be misinterpreted to assume that the combination of the two produces the correct free energy surface.

a significant component of the discussion. By a stroke of luck, the rare event sampling method of choice was metadynamics, the method determined in this chapter to offer a far superior insight into the validity of the reaction coordinates. This allows us to revisit the results of Chapters 5 and 6, and determine more rigorously whether dehydration is important in each case.

Lone Ca ions

In Chapter 5, we concluded that dehydration of lone Ca ions was not necessary. This came down to two reasons: first, the free energy barriers from dehydration were found to be relatively small (Figure 5.3); second, the free energy surfaces obtained with and without dehydration were found to converge to very similar functions for two kink types (Figure A2). While this is a reasonable amount of evidence that dehydration is not necessary here, it would be of interest to determine whether this conclusion is supported by the history of the bias potential. We observed in Section 8.3.3 that an inappropriate choice of reaction coordinates will become apparent in the fluctuation of the bias potential over time. This can easily be examined by outputting the calculated bias potential at different intervals over the simulation and determining whether the bias potential fluctuates significantly over time.

Figure 8.12 shows the bias potential outputted at different time intervals for the kink sites examined in Section 5.4 and also shown in Figure A2. Surprisingly, each free energy surfaces shows a significant degree of fluctuation over time, irrespective of whether dehydration has been applied. For the

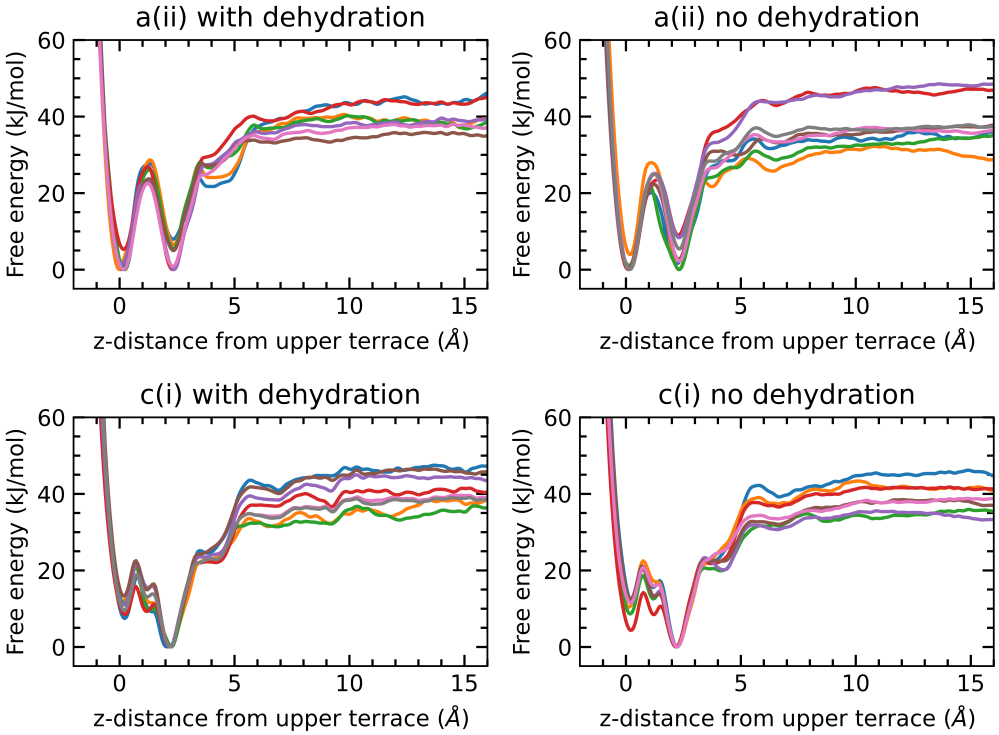

Figure 8.12: Free energy calculations for Ca ions as a function of their distance to the a(ii) and c(i) kinks with and without dehydration. These free energy surfaces are from the same simulations used to derive figure A2, except the calculated free energies have been outputted multiple times over the course of each simulation. The frequency at which the free energies were outputted was dependent on the simulation, as the total simulation times were different in each case.

a(ii) kink, the bias appears to fluctuate by a slightly greater amount when dehydration has not been applied, indicating that the lack of dehydration has impacted the convergence of the simulation. This is much less apparent for the c(i) kink, where the bias fluctuates by a very similar amount with and without dehydration. What is most surprising, however, is that the bias does fluctuate quite significantly even when dehydration is applied. The largest degrees of fluctuation are for the dissolved states, where the outputs of the bias fluctuate over more than 10 kJ/mol in each instance. It is worth noting that, in Chapters 5, we only biased the position component of the adsorbate normal to the {10.4} surface (the z-direction in simulations), while allowing movement along the perpendicular (x, y)-components within a set of constraints (see Figure 4.3). It may be that this apparent lack of convergence has arisen because we neglected the (x, y)-components of the position of the adsorbate. It appears that, while including more reaction coordinates in Chapter 5 may have been ideal, it may be more beneficial to examine other spatial components, rather than dehydration.

Asp binding to kink sites

In Chapter 6, we examined the binding of asp to the a(ii) Ca-terminated kink site, and found that the most stable configuration corresponded to one where the carboxyl groups only adsorbed to the upper terrace, similar to the bidentate configuration observed in Chapter 5. This was later generalised to all kink sites under the assumption that dehydration of the kink site isn't necessary. While this gave rise to results which ultimately were in agreement with experimental calculations, we later examined the hydration of all nearby Ca ions, and found the binding to some kink sites actually did cause the displacement of some water molecules near the a(ii) kink site (see Figure 6.12). While this wasn't an issue for the remaining a- and d-kinks, this observation does highlight that there could be potential inaccuracies in out results which could arise from dehydration. Since we also used metadynamics in Chapter 6, we are again in a position to test whether the reaction coordinates chosen were suitable. Unfortunately, in simulations carried out in Chapter 5, we did not output the value of the reaction coordinates frequently enough to apply Mean Force Integration, which is why we only examine the history of the bias potential in the previous section. However, in Chapter 6, we outputted the bias potential every 0.2 ps (corresponding to $\tau = 10$, see Equation 3.104). This allows us too apply Mean Force Integration and directly examine the error of each simulation.

Applying Mean Force Integration to simulations with multiple walkers becomes significantly more difficult for the following reason: as a single simulation runs, the bias is updated after a specified number of timesteps by appending the Gaussian hills accumulated over all other walkers to the bias. Keeping track of this is difficult, as not all simulations will run at the exact same speed. If one simulation runs slower than another, it may accumulate an additional set of hills from the other walker, whereas the other walker will not accumulate any hills from the first. This becomes an issue, as Mean Force Integration requires knowledge of the exact value of the bias potential at all times. While calculating a near estimate of the bias potential as a function of time is easily achieved by simply ordering the Gaussian hills by time, applying Mean Force Integration to this estimate of the bias will produce an inaccurate estimate of the free energy surface which will not reduce over time. For this reason, we do not calculate the free energy surface using Mean Force Integration here. We can, however, use Mean Force Integration to calculate an error as discussed in section 8.3.3. We reason that, while the free energy surfaces produced will be inaccurate, the difference will be small and will have very little effect on the magnitude of the standard error. It should be noted that this approach will not represent the exact error of the free energy derived from metadynamics, but will offer quantitative measures of convergence which can be compared with one another.

Figure 8.13 shows the free energy surfaces as calculated in Chapter 6, along with the errors calculated using Mean Force Integration, for the a(i),

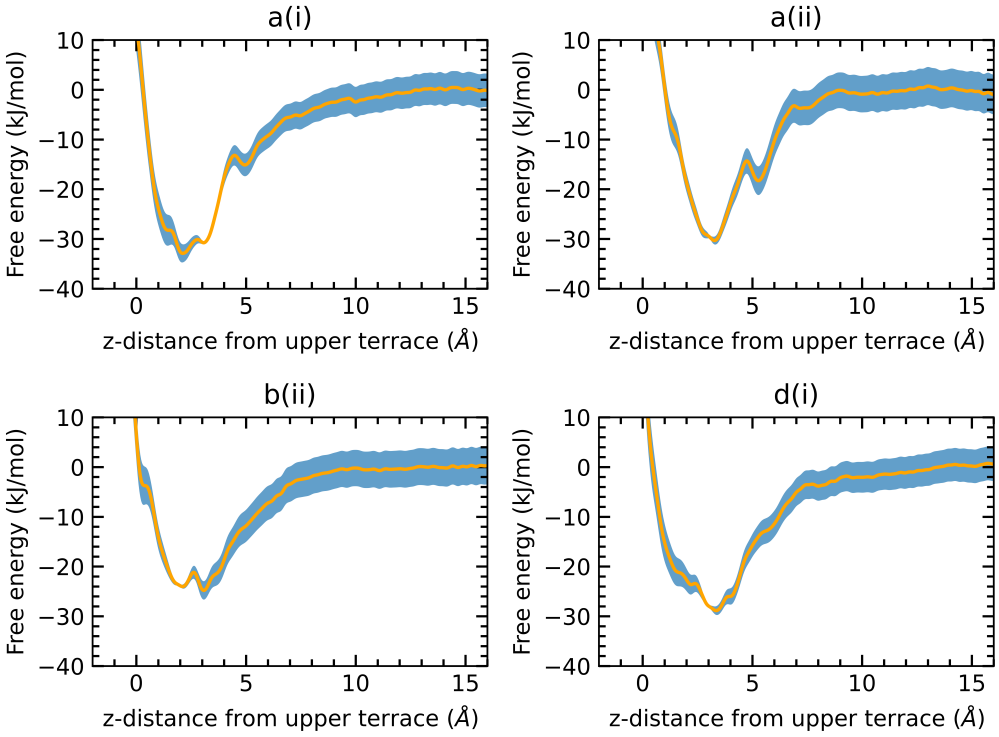

Figure 8.13: Free energy of asp as a function of its distance to the a(i), a(ii), b(ii) and d(i) kinks (shown in orange) with the error calculated using mean force integration (shown in blue). The errors for the a(ii) and b(ii) kinks are slightly larger than the other two, most likely due to the displacement of strongly-bound water at the kink site brought about by the binding of asp (see figure 6.12).

a(ii), b(ii) and d(i) kinks. We have omitted the result for the d(ii) kink for the sake of easier viewing, as the error is very similar to that of the d(i) kink, and therefore does not offer any additional value. It was found in Chapter 6 that the binding of asp to the a(ii) and b(ii) kinks resulted in the displacement of a single water molecule from the terminating Ca ion, whereas binding to the a(i), o(i) and o(ii) kinks did not displace any water molecules from the surface. This is somewhat reflected in the errors shown in figure 8.13: the errors for the a(ii) and b(ii) kinks for the unbound state (4.0 kJ/mol and 3.8 kJ/mol respectively) are marginally higher than those for the a(i), d(i) and d(ii) kinks (3.5 kJ/mol, 3.6 kJ/mol and 3.5 kJ/mol respectively). It should be noted that these values are calculated by averaging the errors between 15 and 20 Å from the surface(as is done to calculate the simulation free energy, see Figure 6.8). While it is apparent that the displacement of the water molecule causes an increase in the error of the simulation, the difference is surprisingly small. It appears that the dehydration of nearby ions has very little impact on the convergence of the metadynamics simulation. It should again be noted, however, that these errors are still relatively large, similar to what is observed in the previous section. Again, it is quite likely that these errors arise from the fact that we have only biased a single spatial component

while constraining the other spatial components within a region.

Can we trust our results?

One reoccurring theme in this section is that we have consistently found a fair degree of uncertainty of our results. We may expect to find a sampling error on the order of magnitude of 1 k_BT, but our errors are larger, indicating that convergence may also be a source of uncertainty. It is important, at this point, that we discuss the validity of our conclusions in previous chapters.

In Section 8.3.4, we found that dehydration is not completely insignificant at least for the a(ii) kink, although any uncertainty arising from this is dwarfed by a more consistent uncertainty, which we attribute to other neglected reaction coordinates (most likely other spatial components). While it is difficult to make rigorous conclusions based on the results of Figure 8.12, we consider it likely that the true errors of the simulation are greater than 1 kJ/mol. While this does mean that the adsorption energies calculated in Table 5.1 should be considered with some degree of caution, we should also note that the main conclusions of Chapter 5 are not affected by this uncertainty due to convergence. For example, we observe in Figure 8.12 that the Ca ion adsorbing to the c(i) kink always favours the bidentate configuration. When we compare this to figure 8.9 (which shows a far more significant fluctuation of the free energy landscape throughout the course of the simulation), it is clear that the uncertainty is nowhere near large enough to cast any ambiguity over whether the c(i) Ca kink favours the bidentate configuration. The conclusions in Chapter 5 (that many ions do not adsorb directly into the lattice site and that calcite growth is often multistep) therefore remain the same regardless.

In section 8.3.4, we found that the displacement of water molecules near the kink site due to the binding of asp had little impact on the convergence of the simulations. This is greatly reassuring regarding the conclusions of Chapter 6. However, we again found a significant degree of uncertainty due to convergence. We should reiterate that the MFI uncertainty does not necessarily equal the error of the simulation, however it is clear that there is some degree of uncertainty which cannot be explained by the sampling errors of the simulations. While this again means the exact calculations of adsorption free energies should be approached with caution, we also note that the experimental values for asp adsorption free energies discussed in Chapter 6 varied between -17.1 kJ/mol and -21 kJ/mol. The difference between these values alone is about the same as the errors calculated using Mean Force Integration. Considering both the experimental and computational uncertainties, the consistency between experimental and computational adsorption free energies remains intact, and our conclusions do not change. We should note, however, that some of the results in Chapter 5 may be highly sensitive to these convergence errors. The most notable examples are the morphologies calculated in Section 6.2.5. These morphologies are dependent

on the adsorption and dissolution rates of asp in the KMC model used, which are dependent on the exponents of the adsorption free energies. These reaction rates, and therefore the predicted morphologies of the growth island, are therefore highly sensitive to these convergence issues. This itself isn't an issue however, since the results of this section were largely inconclusive anyway.

8.3.5 Towards a rare event sampling method for handling poor reaction coordinates

The best way to avoid the issues described above is to identify and include the necessary collective variables for isolating reactant and product states. Knowing when to do this isn't always obvious from simulations, although we note that metadynamics-based techniques are far more able to identify when an important collective variable is missing. However, a far superior way around this would be to implement a technique which can correct for the errors arising from poor reaction coordinates.

This could exist in two forms. One means could be to introduce some kind of memory to a metadynamics-based approach. Figure 8.5 demonstrates that adding memory to a simulation is able to improve the calculated crossing rate. Applying a similar approach to a bias potential may be possible. This could be used to overcome the difficulties discussed in this chapter without the need for searching for additional reaction coordinates. Another approach could be to post-process the output of e.g. a metadynamics simulation in an attempt to estimate the true PMF. For example, one interesting observation to be made about the results of these simulations is that the Jarzynski approach produces an accurate estimate on the side of the saddle point at which it starts. The method starts to fail in the vicinity of the saddle point. A similar point can be made about the output of metadynamics simulations at different points in time. This isn't to say that either approach could be used to overcome the issues presented by poor reaction coordinate. For that to work, the free energy difference between reactant and product states would have to be known. This cannot be obtained through any of the methods studied here, although the exact position of the saddle point could give an estimate of where the 'forward' and 'backward' free energy plots meet. Calculating the position of the saddle point may be possible from the output of a metadynamics simulation such as the one in Figure 8.9, in which a visual distinction can be made between reactant and product states from the output of the collective variable as a function of time. It may be possible to use this data to estimate the position of the saddle point. We do not offer a means to do so, neither do we propose a method which involves a memory-dependent bias. However, the work carried out throughout this chapter provides a basis on which such a method could be proposed.

8.4 Conclusions

This chapter began by modelling a discrete barrier-crossing process in two dimensions, and examining the same process over a coarse-grained, one-dimensional barrier. To do so, we developed a toy PMF, with a tunable parameter determining the skew of the saddle point separating reactant and product states. Taking the x-component of the PMF as a reaction coordinate, we were able to see important detail neglected in the one dimensional PMF in the form of a reduction of the free energy barrier height when the saddle point was skewed. Deriving Markovian reaction rates along the x-direction resulted in a overestimated crossing rate. However, introducing non-Markovianity by deriving reaction rates based on previous configurations provided a correction to the barrier crossing rate. Our results have demonstrated that coarse graining introduces non-Markovian kinetics into the coarse-grained system. The extent of the effects on non-Markovianity are dependent on the extent to which the free energy barrier height is underestimated.

We proceeded to turn to a scenario seen in practice: one in which rare event sampling methods are used to attempt to recover the one-dimensional surface where the reaction coordinate is unsuitable, i.e. the coarse-grained kinetics are non-Markovian. The methods applied were metadynamics, umbrella sampling, mean force integration and the Jarzynski method. Apart from the Jarzynski method, all methods produced a PMF which varied significantly about the true PMF based on the initial velocity seed. In every case, the method breaks down for the same reason we observe non-Markovianity: the reaction coordinate choice neglects an important component of the multi-dimensional saddle point, which results in a smaller free energy barrier on the coarse-grained system. The rare event sampling method, which attempts to recreate this smaller barrier, is insufficient to allow a crossing to take place.

We additionally demonstrated that the output of a single metadynamics simulation was sufficient to assess the suitability of the reaction coordinate choice. Applying mean force integration additionally allows for the quantification of the error of the simulation. Umbrella sampling, as well as the Jarzynski approach, is disadvantageous in this regard, as the sampling errors do not account for the true error of the simulation, and are more likely to produce incorrect results without giving any indication of their validity. This result is of general interest to the molecular dynamics and rare event sampling communities. Extreme caution should be taken when using certain methods, particularly umbrella sampling, as the validity of the collective variable choice is not apparent. Metadynamics-based approaches are significantly less problematic in this regard.

We also put the results in the previous sections to practice by applying them to the work carried out in Chapters 5 and 6. We used the output of metadynamics simulations, as well as Mean Force Integration, to determine whether ignoring dehydration of lone Ca ions and calcite kink sites a sensible

choice. Intriguingly, we found that, while dehydration generally had little impact on the convergence of simulations, there was still a degree of uncertainty evident in the convergence of metadynamics simulations as well as the error calculating using Mean Force Integration. While these results had little impact on the conclusions of this thesis, we note that convergence is often a generally under-appreciated feature of rare event sampling which should be addressed when possible. Applying the techniques discussed in this chapter may be very useful in determining whether the reaction coordinate choices are appropriate.

We finally note that it may be possible to develop a rare event sampling technique which does not require such a careful selection of reaction coordinates in order to recover an accurate estimate of the PMF. This offers a potential scope for future work in this field. Whatever form such a technique will take, it is sure to be based on non-Markovianity.

9 CONCLUSIONS

9.1 Summary

Chapter 1 outlined the long-term ambitions of materials science, and provided an introduction to the concept of biomineralisation and why materials scientists aim to exploit it. Nature teaches us that an exquisite level of control over the crystallisation process is achievable, and that biological systems can harness these processes to produce materials with bespoke morphological and mechanical properties despite only having access to the chemical composition and ambient conditions of their natural environment. This Chapter concluded with the three main objectives of this thesis: First, we wished to use molecular simulation to model the interactions between adsorbates (including Ca, CO_3 and many impurities such as amino acids) with the kink sites in which the majority of growth events take place; second, we wished to use these results in tandem with coarse grained models to predict macroscopic observables of calcite grown in the presence of impurities; finally, we wished to use molecular simulation to model case studies of biogenic calcite in order to elucidate the atomistic mechanisms behind their superior properties.

Chapter 2 provided a comprehensive discussion of calcite, the most prevalent carbonate on Earth, and a hugely abundant biomineral, present in biological systems such as many forms of marine life. In Chapter 2, the computational methods typically used to study calcite mineralisation were outlined, and a review of the progress and challenges faced in understanding its growth was carried out. Additionally, impurities were discussed in terms of their effects on the growth and the morphological and mechanical properties of calcite. A review of the experimental and computational techniques used to study impurities and calcite was carried out. It was noted that many of the impacts of impurities on the calcite growth process are well characterised through experiment, although their computational counterparts are severely lacking, due to a lack of thermodynamically consistent force fields and understanding of the kink sites which ultimately facilitate the growth of calcite. Additionally, studies of the mechanical impact of impurities on calcite are limited to straightforward synthetic examples, or biological case studies too complex for current molecular simulation resources and techniques.

In Chapter 5, we addressed the issues faced in understanding the kink growth process by using molecular dynamics and rare event sampling techniques to determine adsorption free energies and binding configurations of Ca and CO_3 units adsorbing to calcite kink sites. Most crucially, we find that not all kink sites follow the classical picture of units adsorbing into their kink lattice sites via a single-step process. Instead, many ions reside in

a bidentate configuration, named as such for their tendency to reside above the step and form two interactions with units on the upper terrace. These bidentate configurations were observed to a larger extent for CO_3 units, despite the greater adsorption free energies found for CO_3. The configurations in which atoms reside above the step are analogous to what is observed for kink nucleation, in which CO_3 units prefer to interact only with the step. Having identified that the prevalence of the bidentate configuration implies that the growth process must involve multiple adsorbates, we proceeded to simulate a two-step kink growth process, and found that the terminating CO_3 ion was only able to transition from its bidentate configuration into its lattice site upon the adsorption of an additional Ca ion. We have shown in Chapter 5 that the growth of kink sites is significantly more complex than expected, and that kink-terminating ions do not always reside in their lattice site. This additionally has consequences for the study of impurities and their interaction with calcite, since the natural configuration of the kink site must be taken into account.

Having laid the groundwork for the study of solutes and their interaction with kink sites, Chapter 6 concerned the interaction of various functional groups with calcite kink sites, in addition to terrace and step sites. We began with a case study which found that positively-charged polyamine groups and amine-rich amino acids (lysine and arginine) were able to occlude in calcite at high concentrations, contrary to the view that negatively-charged functional groups were the primary driver of adsorption of peptides and amino acids to calcite. By isolating the relevant primary and secondary amine groups, we found a strong interaction to negatively-charged calcite terrace, step and kink sites. In particular, the NH_3^+ group had a very strong interaction with CO_3-terminated kinks. We also observed a significantly smaller interaction between the arginine side chain and all surface sites, leading to the conclusion that lysine binds primarily through its side-chain amine group, whereas arginine doesn't. Our simulations, in combination with experimental results have changed our understanding of how the adsorption and occlusion of biomolecules is facilitated. These results also provoke additional questions into whether amine groups or carboxyl groups facilitate the occlusion of other amino acids into calcite.

In the second part of Chapter 6, we proceeded to study the interaction of aspartate with calcite kink sites, with an ultimate aim to derive the adsorption free energy of aspartate and compare with experimental findings. After eliminating most complex configurations, we used rare event sampling to determine the adsorption free energies and configurations of asp. We initially found that asp does not displace water out of the kink site, and instead adopts a configuration akin to the bidentate configuration observed in Chapter 5 in which the amino acid only interacts through its carboxyl group with Ca ions on the upper terrace. We subsequently proceeded to determine adsorption free energies and configurations for all kink types, and found a similar strength and mechanism of interaction for most kink sites, save for

a single exception where the terminating Ca-ion prefers to adopt its bidentate configuration. The adsorption free energies were used to parametrise a KMC model, which was used to calculate the step velocity inhibition as a function of aspartate concentration. It was ultimately found that the acute step velocities were hindered more significantly than obtuse step velocities, in line with previous expectations. However, one caveat is that we were unable to reproduce experimental morphologies using our KMC simulations. Finally, we used the step velocities to calculate the normal growth rate of the crystal, and derived an adsorption free energy of -19 kJ/mol, a value within the range of empirically derived values. This second part of Chapter 6 is particularly significant, as it validates the atomic potentials used by reproducing experimental values to a quantitative degree. We have determined the likely mechanisms through which aspartate interacts with calcite, something which may be extended to other amino acids in future. The work carried out in this chapter is also the first of its kind, in that it comprehensively studies the interaction between an impurity and all kink sites. This work has provided significant progress in how impurities and their tuning of calcite growth can be studied with molecular simulation. However, many open questions remain. For example, we do not know why aspartate (as well as many other impurities) does not produce a morphology change on obtuse steps, despite significantly inhibiting their growth. Additionally, questions remain regarding the extent of the influence of amine groups on the binding of amino acids.

Having studied the effects of impurities on the growth of calcite, Chapter 7 concerns the impact of impurities on the mechanical properties of calcite. We began with a case study in which non-charged, polymer-coated nanoparticles were able to incorporate into calcite at high concentrations without altering the single-crystallinity of the host crystal. The lattice distortions caused by the nanoparticles were unusual in that an elongation in the c-axis was produced, while a compression along the a-axis was produced. We used molecular dynamics to understand this phenomenon by approximating the weakly-interacting nanoparticles as holes within the lattice. By calculating stress distributions and examining lattice spacings, we found that the nanoparticles caused a compression in the crystal along the a-axis, while itself contracting along the c-axis without compressing the host crystal. The average lattice spacing was used to calculate the average strain of the material with respect to pure calcite, and the calculated signs and magnitudes were in agreement with experimental values. Our results showed that the anomalous, axis-dependent strain of the material is a result of the calcite molecular structure, rather than any properties of the nanoparticle itself.

In the second section of Chapter 7, we turned our attention to a case study of the brittlestar *Ophiocoma wendtii*, in which magnesium-rich coherent nanoparticles were found to improve the fracture toughness of calcitic lenses within the arm of the brittlestar. It was thought that the nanoprecipitates improve the structure by inducing a compressive stress on the host

crystal due to the tensile stress within the nanoprecipitate. We used molecular dynamics to calculate stress distributions within the lattice, and directly simulated a crack-propagation event to identify toughening mechanisms. We identified three sources: first, the suggested mechanism of prestressing the lattice is confirmed through simulation; second, the inhomogeneous stress distribution within the lattice deflects any propagating cracks away from the lowest energy surface towards the nanoprecipitate, resulting in an increase yield stress; third, due the the highly inhomogeneous stress distribution within the nanoprecipitate, propagation directly through the nanoprecipiate produces a highly disordered crack which does not cleanly propagate along the {10.4} cleavage plane.

In Chapter 8, we created a toy system consisting of a two-dimensional free energy surface with an adjustable parameter controlling the skew of a single saddle point. By taking a single component as a reaction coordinate, we were able to tune the amount of detail neglected when coarse-graining onto a single coordinate. When attempting to numerically recover the free energy barrier crossing rate in one-dimension, we found a large discrepancy between the correct result and the coarse-grained result as more detail was neglected. However, the introduction of non-Markovian kinetics partially corrected the coarse-grained reaction rate. The extent of the correction was universally found to be greater when the memory introduced to the system was increased. Our calculations have confirmed that the coarse-graining of a system imposes non-Markovian kinetics on the system, which can have a significant contribution to the free energy barrier crossing rate.

We proceeded in Chapter 8 to characterise the failure of rare event sampling techniques when the reaction coordinate choice neglects important detail of the system thermodynamics. We found that all techniques studied ultimately produced inaccurate results when attempting to recover the free energy landscape. Crucially, however, we found that metadynamics-based approaches were far superior for determining whether the reaction coordinate choice is suitable, due to the history-dependent bias potential producing different outputs at different times. Methods such as umbrella sampling do not give any indication as to whether an appropriate set of reaction coordinates have been used, and can produce an incorrect free energy landscape which appears converged and accurate. In future, studies adopting umbrella sampling techniques should also provide evidence that the reaction coordinate choice is suitable, or risk recovering significantly inaccurate results. Alternatively, using metadynamics-based approaches gives an indication of the error either through the history of the bias potential or by calculating the error of the simulation using methods such as mean force integration.

9.2 Outlook

For the most part, we have fulfilled the objectives of this thesis. In Chapter 5, we successfully modelled the interactions of Ca and CO_3 units with the majority of calcite kink sites, thus identifying that calcite growth is not, as previously thought, a series of single-step attachment events. We were also able in Chapter 6 to determine the binding configurations and adsorption free energies of an amino acid, plus several functional groups, to multiple kink sites. We were also broadly successful in using our results for aspartate to parametrise a coarse-grained (KMC) model. This model predicted the adsorption free energy of aspartate so be within the range of experimental results, thus validating our previous calculations of adsorption free energies to kink sites. Finally, in Chapter 7, we successfully applied molecular simulation to the case study of the brittlestar *Ophiocoma wendtii*. We were able to identify several mechanisms through which the Mg-calcite is able to increase the fracture toughness of calcite, something that would not be achievable through experimentation.

Perhaps the largest open questions remaining upon completion of this thesis have arisen from Chapter 6. While we have made significant progress in understanding how impurities interact with calcite growth sites, the work in Chapter 6 has highlighted the fact that many question remain. The morphological impact of aspartate on obtuse steps is an example of this: despite producing step velocity inhibitions consistent with reality, our simulations consistently produced a rounding of the obtuse steps: something not observed in reality. There are many possible explanations for this. One such explanation is that we only parametrise our KMC model against free energy calculations for kink sites, and approximate attachment to corner sites and pseudofaces as attachments to kink sites. Since pseudofaces and corners are observed both in reality and in simulations, ignoring both could have a significant impact on the morphologies produced from KMC simulations. Further calculations of adsorption free energies to corner sites and pseudofaces could provide further insight and be used to parametrise a superior KMC model. Another notable issue is that many unanswered questions remain regarding the growth of calcite itself. In Chapter 5, we were forced to ignore several CO_3-terminated kink sites due to the number of reaction coordinates required and therefore the large simulation times required. Nevertheless, running such calculations would still be possible given enough time. Furthermore, these calculations could be extended to corner sites to obtain more information on whether the terminating ions adopt a lattice or bidentate configuration.

One particular theme connecting chapters 5 and 6 is the favouring of configurations in which adsorbates only interact with the upper terrace without displacing water from the kink site. This is particularly noticeable for aspartate, which is found to produce a minimal dehydration of the surface. This leads to questions as to whether aspartate may occlude with water. This also

runs counter to previous expectations that impurities such as asp are able to improve mass transport to the surface via its dehydration. More work is clearly required in this field. One area of further study would be to calculate binding free energy to step and terrace sites and observe whether any dehydration takes place for these surfaces.

Another concern arising from Chapter 6 is the large discrepancy observed in the strengths of the interactions between functional groups and calcite surface ions. In our study of amine groups, using the AMBER force fields, we found a strong interaction with negatively-charged calcite surface sites. However, with a different set of intermolecular interactions for aspartate (Raiteri et al. 2012), we find that binding is dominated by carboxyl groups, and that amine groups exhibited next to no binding. We know from experimental occlusion rates that amine groups must have a stronger interaction than those we find in Section 6.2, but a weaker one than that observed in Section 6.1. Clearly, more work is required in developing force fields which can accurately model the interaction between amine groups and calcite. Some success has already been found from the development of polarisable force fields for certain functional groups. This is something which could be extended to modelling amino acids in their entirety.

Bibliography

[1] J. R. Young and K. Henriksen, "Biomineralization within vesicles: the calcite of coccoliths," *Reviews in Mineralogy and Geochemistry*, vol. 54, no. 1, pp. 189–215, 2003.

[2] K. Müller and L. D. Brown, "Location of saddle points and minimum energy paths by a constrained simplex optimization procedure," *Theoretica Chimica Acta*, vol. 53, no. 1, pp. 75–93, 1979.

[3] I. Polishchuk, A. A. Bracha, L. Bloch, D. Levy, S. Kozachkevich, Y. Etinger-Geller, Y. Kauffmann, M. Burghammer, C. Giacobbe, J. Villanova, *et al.*, "Coherently aligned nanoparticles within a biogenic single crystal: a biological prestressing strategy," *Science*, vol. 358, no. 6368, pp. 1294–1298, 2017.

[4] J. C. Weaver, G. W. Milliron, A. Miserez, K. Evans-Lutterodt, S. Herrera, I. Gallana, W. J. Mershon, B. Swanson, P. Zavattieri, E. DiMasi, and D. Kisailus, "The stomatopod dactyl club: a formidable damage-tolerant biological hammer," *Science*, vol. 336, no. 6086, pp. 1275–1280, 2012.

[5] A. H. Heuer, D. J. Fink, V. J. Laraia, J. L. Arias, P. D. Calvert, K. Kendall, G. L. Messing, J. Blackwell, P. C. Rieke, D. H. Thompson, A. P. Wheeler, A. Veis, and A. I. Caplan, "Innovative materials processing strategies: a biomimetic approach," *Science*, vol. 255, no. 5048, pp. 1098–1105, 1992.

[6] L. M. Gordon, M. J. Cohen, K. W. MacRenaris, J. D. Pasteris, T. Seda, and D. Joester, "Amorphous intergranular phases control the properties of rodent tooth enamel," *Science*, vol. 347, pp. 746–750, 2015.

[7] S. Weiner and P. M. Dove, "An overview of biomineralization processes and the problem of the vital effect," *Reviews in Mineralogy and Geochemistry*, vol. 54, no. 1, pp. 1–29, 2003.

[8] Y. Politi, T. Arad, E. Klein, S. Weiner, and L. Addadi, "Sea urchin spine calcite forms via a transient amorphous calcium carbonate phase," *Science*, vol. 306, no. 5699, pp. 1161–1164, 2004.

[9] C. E. Killian, R. A. Metzler, Y. Gong, I. C. Olson, J. Aizenberg, Y. Politi, F. H. Wilt, A. Scholl, A. Young, A. Doran, *et al.*, "Mechanism of calcite co-orientation in the sea urchin tooth," *Journal of the American Chemical Society*, vol. 131, no. 51, pp. 18404–18409, 2009.

[10] J. D. Currey, "Mechanical properties of mother of pearl in tension," *Proc. R. Soc. B*, vol. 196, no. 1125, pp. 443–463, 1977.

[11] F. Barthelat, H. Tang, P. D. Zavattieri, C. M. Li, and H. D. Espinosa, "On the mechanics of mother-of-pearl: A key feature in the material hierarchical structure," *J. Mech. Phys.*, vol. 55, pp. 306–337, 2007.

[12] S. Weiner and H. D. Wagner, "The material bone: Structure-mechanical function relations," *Annu. Rev. Mater. Sci.*, vol. 28, pp. 271–297, 1998.

[13] P. Fratzl and R. Weinkamer, "Nature's hierarchical materials," *Prog. Mater. Sci.*, vol. 52, pp. 1263–1334, 2007.

[14] W.-T. Hou and Q.-L. Feng, "Morphologies and growth model of biomimetic fabricated calcite crystals using amino acids and insoluble matrix membranes of *Mytilus edulis*," *Crystal Growth & Design*, vol. 6, no. 5, pp. 1086–1090, 2006.

[15] W. D. Carlson, "The calcite–aragonite equilibrium: effects of Sr substitution and anion orientational disorder," *American Mineralogist*, vol. 65, no. 11-12, pp. 1252–1262, 1980.

[16] H. Blatt, "Origin of sedimentary rocks," *Soil Science*, vol. 115, no. 5, p. 400, 1973.

[17] G.-T. Zhou, Q.-Z. Yao, S.-Q. Fu, and Y.-B. Guan, "Controlled crystallization of unstable vaterite with distinct morphologies and their polymorphic transition to stable calcite," *European Journal of Mineralogy*, vol. 22, no. 2, pp. 259–269, 2010.

[18] W. A. Deer, "Rock-forming minerals," Geological Society of London, 1978.

[19] H. H. Teng, P. M. Dove, and J. J. De Yoreo, "Kinetics of calcite growth: surface processes and relationships to macroscopic rate laws," *Geochimica et Cosmochimica Acta*, vol. 64, no. 13, pp. 2255–2266, 2000.

[20] J. J. De Yoreo and P. G. Vekilov, "Principles of crystal nucleation and growth," *Reviews in Mineralogy and Geochemistry*, vol. 54, no. 1, pp. 57–93, 2003.

[21] W.-K. Burton, N. Cabrera, and F. Frank, "The growth of crystals and the equilibrium structure of their surfaces," *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 243, no. 866, pp. 299–358, 1951.

[22] A. Gratz, P. Hillner, and P. Hansma, "Step dynamics and spiral growth on calcite," *Geochimica et Cosmochimica Acta*, vol. 57, no. 2, pp. 491–495, 1993.

[23] G. I. Taylor, "The mechanism of plastic deformation of crystals. part i.—theoretical," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 145, no. 855, pp. 362–387, 1934.

[24] R. Demichelis, A. Schuitemaker, N. A. Garcia, K. B. Koziara, M. De La Pierre, P. Raiteri, and J. D. Gale, "Simulation of crystallization of biominerals," *Annual Review of Materials Research*, vol. 48, pp. 327–352, 2018.

[25] M. De La Pierre, R. Demichelis, U. Wehrmeister, D. E. Jacob, P. Raiteri, J. D. Gale, and R. Orlando, "Probing the multiple structures of vaterite through combined computational and experimental raman spectroscopy," *The Journal of Physical Chemistry C*, vol. 118, no. 47, pp. 27493–27501, 2014.

[26] L. S. Parvaneh, D. Donadio, and M. Sulpizi, "Molecular mechanism of crystal growth inhibition at the calcium oxalate/water interfaces," *The Journal of Physical Chemistry C*, vol. 120, no. 8, pp. 4410–4417, 2016.

[27] G. Wulff, "Zur frage der geschwindigkeit des wachstums und der auflö der kristallflä zeitschr," *F. Kristallog*, vol. 34, pp. 449–530, 1901.

[28] H. Eyring, "Structure and properties of solid surfaces," *Journal of Chemical Education*, vol. 31, no. 9, p. 501, 1954.

[29] S. Parker, "Prediction of mineral crystal structures," *Solid State Ionics*, vol. 8, no. 3, pp. 179–186, 1983.

[30] S. Parker, C. Catlow, and A. Cormack, "Prediction of mineral structure by energy minimisation techniques," *Journal of the Chemical Society, Chemical Communications*, no. 17, pp. 936–938, 1983.

[31] S. Parker, C. Catlow, and A. Cormack, "Structure prediction of silicate minerals using energy-minimization techniques," *Acta Crystallographica Section B: Structural Science*, vol. 40, no. 3, pp. 200–208, 1984.

[32] N. H. de Leeuw and S. C. Parker, "Surface structure and morphology of calcium carbonate polymorphs calcite, aragonite, and vaterite: an atomistic approach," *The Journal of Physical Chemistry B*, vol. 102, no. 16, pp. 2914–2922, 1998.

[33] J. Titiloye, N. De Leeuw, and S. Parker, "Atomistic simulation of the differences between calcite and dolomite surfaces," *Geochimica et Cosmochimica Acta*, vol. 62, no. 15, pp. 2637–2641, 1998.

[34] S. Piana, M. Reyhani, and J. D. Gale, "Simulating micrometre-scale crystal growth from solution," *Nature*, vol. 438, no. 7064, pp. 70–73, 2005.

[35] J. Li and M. F. Doherty, "Steady state morphologies of paracetamol crystal from different solvents," *Crystal Growth & Design*, vol. 17, no. 2, pp. 659–670, 2017.

[36] M. W. Anderson, J. T. Gebbie-Rayet, A. R. Hill, N. Farida, M. P. Attfield, P. Cubillas, V. A. Blatov, D. M. Proserpio, D. Akporiaye, B. Arstad, *et al.*, "Predicting crystal growth via a unified kinetic three-dimensional partition model," *Nature*, vol. 544, no. 7651, pp. 456–459, 2017.

[37] A. R. Hill, P. Cubillas, J. T. Gebbie-Rayet, M. Trueman, N. de Bruyn, Z. Al Harthi, R. J. Pooley, M. P. Attfield, V. A. Blatov, D. M. Proserpio, *et al.*, "CrystalGrower: a generic computer program for Monte Carlo modelling of crystal growth," *Chemical science*, vol. 12, no. 3, pp. 1126–1146, 2021.

[38] S. Kerisit and S. C. Parker, "Free energy of adsorption of water and metal ions on the {1014} calcite surface," *Journal of the American Chemical Society*, vol. 126, no. 32, pp. 10152–10161, 2004.

[39] R. Kristensen, S. Stipp, and K. Refson, "Modeling steps and kinks on the surface of calcite," *The Journal of Chemical Physics*, vol. 121, no. 17, pp. 8511–8523, 2004.

[40] D. Spagnoli, S. Kerisit, and S. C. Parker, "Atomistic simulation of the free energies of dissolution of ions from flat and stepped calcite surfaces," *Journal of Crystal Growth*, vol. 294, no. 1, pp. 103–110, 2006.

[41] M. Wolthers, D. Di Tommaso, Z. Du, and N. H. de Leeuw, "Calcite surface structure and reactivity: molecular dynamics simulations and macroscopic surface modelling of the calcite–water interface," *Physical Chemistry Chemical Physics*, vol. 14, no. 43, pp. 15145–15157, 2012.

[42] M. Wolthers, D. Di Tommaso, Z. Du, and N. H. de Leeuw, "Variations in calcite growth kinetics with surface topography: molecular dynamics simulations and process-based growth kinetics modelling," *CrystEngComm*, vol. 15, no. 27, pp. 5506–5514, 2013.

[43] M. De La Pierre, P. Raiteri, A. G. Stack, and J. D. Gale, "Uncovering the atomistic mechanism for calcite step growth," *Angewandte Chemie International Edition*, vol. 56, no. 29, pp. 8464–8467, 2017.

[44] P. Raiteri, J. D. Gale, D. Quigley, and P. M. Rodger, "Derivation of an accurate force-field for simulating the growth of calcium carbonate from aqueous solution: A new model for the calcite- water interface," *The Journal of Physical Chemistry C*, vol. 114, no. 13, pp. 5997–6010, 2010.

[45] P. Raiteri, J. D. Gale, D. Quigley, and P. M. Rodger, "Derivation of an accurate force-field for simulating the growth of calcium carbonate from aqueous solution: a new model for the calcite–water interface," *J. Phys. Chem. C*, vol. 114, pp. 5997–6010, 2010.

[46] W. Stumm, J. J. Morgan, and J. I. Drever, "Aquatic chemistry," *Journal of Environmental Quality*, vol. 25, no. 5, p. 1162, 1996.

[47] A. E. Nielsen and J. M. Toft, "Electrolyte crystal growth kinetics," *Journal of Crystal Growth*, vol. 67, no. 2, pp. 278–288, 1984.

[48] D. Di Tommaso, E. Ruiz-Agudo, N. H. de Leeuw, A. Putnis, and C. V. Putnis, "Modelling the effects of salt solutions on the hydration of calcium ions," *Physical Chemistry Chemical Physics*, vol. 16, no. 17, pp. 7772–7785, 2014.

[49] M. Kowacz, C. Putnis, and A. Putnis, "The effect of cation: anion ratio in solution on the mechanism of barite growth at constant supersaturation: Role of the desolvation process on the growth kinetics," *Geochimica et Cosmochimica Acta*, vol. 71, no. 21, pp. 5168–5179, 2007.

[50] M. Andersson, S. Dobberschütz, K. K. Sand, D. Tobler, J. J. De Yoreo, and S. Stipp, "A microkinetic model of calcite step growth," *Angewandte Chemie*, vol. 128, no. 37, pp. 11252–11256, 2016.

[51] J. A. Koskamp, S. E. Ruiz-Hernandez, D. Di Tommaso, A. M. Elena, N. H. De Leeuw, and M. Wolthers, "Reconsidering calcium dehydration as the rate-determining step in calcium mineral growth," *The Journal of Physical Chemistry C*, vol. 123, no. 44, pp. 26895–26903, 2019.

[52] L. N. Lammers, K. Kulasinski, P. Zarzycki, and D. J. DePaolo, "Molecular simulations of kinetic stable calcium isotope fractionation at the calcite-aqueous interface," *Chemical Geology*, vol. 532, p. 119315, 2020.

[53] M. De La Pierre, P. Raiteri, and J. D. Gale, "Structure and dynamics of water at step edges on the calcite {1014} surface," *Crystal Growth & Design*, vol. 16, no. 10, pp. 5907–5914, 2016.

[54] Y. Liang, D. R. Baer, J. M. McCoy, J. E. Amonette, and J. P. Lafemina, "Dissolution kinetics at the calcite-water interface," *Geochimica et Cosmochimica Acta*, vol. 60, no. 23, pp. 4883–4887, 1996.

[55] Y. Liang and D. Baer, "Anisotropic dissolution at the $CaCO_3$ (1014)—water interface," *Surface Science*, vol. 373, no. 2-3, pp. 275–287, 1997.

[56] J. De Yoreo, L. Zepeda-Ruiz, R. Friddle, S. Qiu, L. Wasylenki, A. Chernov, G. Gilmer, and P. Dove, "Rethinking classical crystal growth models through molecular scale insights: consequences of kink-limited kinetics," *Crystal Growth & Design*, vol. 9, no. 12, pp. 5135–5144, 2009.

[57] N. De Leeuw, S. Parker, and J. Harding, "Molecular dynamics simulation of crystal dissolution from calcite steps," *Physical Review B*, vol. 60, no. 19, p. 13792, 1999.

[58] A. G. Stack and M. C. Grantham, "Growth rate of calcite steps as a function of aqueous calcium-to-carbonate ratio: independent attachment and detachment of calcium and carbonate ions," *Crystal Growth & Design*, vol. 10, no. 3, pp. 1409–1413, 2010.

[59] J. N. Bracco, M. C. Grantham, and A. G. Stack, "Calcite growth rates as a function of aqueous calcium-to-carbonate ratio, saturation index, and inhibitor concentration: Insight into the mechanism of reaction and poisoning by strontium," *Crystal Growth & Design*, vol. 12, no. 7, pp. 3540–3548, 2012.

[60] K. Sand, D. Tobler, S. Dobberschutz, K. Larsen, E. Makovicky, M. Andersson, M. Wolthers, and S. Stipp, "Calcite growth kinetics: Dependence on saturation index, Ca: CO_3 -activity ratio, and surface atomic structure," *Crystal Growth & Design*, vol. 16, no. 7, pp. 3602–3612, 2016.

[61] L. C. Nielsen, D. J. De Paolo, and J. J. De Yoreo, "Self-consistent ion-by-ion growth model for kinetic isotopic fractionation during calcite precipitation," *Geochimica et Cosmochimica Acta*, vol. 86, pp. 166–181, 2012.

[62] R. Darkins, I. J. McPherson, I. J. Ford, D. M. Duffy, and P. R. Unwin, "Critical step length as an indicator of surface supersaturation during crystal growth from solution," *Crystal growth & design*, vol. 22, no. 2, pp. 982–986, 2022.

[63] J. D. Weeks and G. H. Gilmer, "Dynamics of crystal growth," *Advances in Chemical Physics*, vol. 40, pp. 157–228, 1979.

[64] A. Chernov, L. Rashkovich, and P. Vekilov, "Steps in solution growth: dynamics of kinks, bunching and turbulence," *Journal of Crystal Growth*, vol. 275, no. 1-2, pp. 1–18, 2005.

[65] H. H. Teng, P. M. Dove, C. A. Orme, and J. J. De Yoreo, "Thermodynamics of calcite growth: baseline for understanding biomineral formation," *Science*, vol. 282, no. 5389, pp. 724–727, 1998.

[66] V. Voronkov, "Structure of a crystal surface and kossel's model," in *Growth of Crystals*, pp. 1–17, Springer, 1976.

[67] A. G. Stack, P. Raiteri, and J. D. Gale, "Accurate rates of the complex mechanisms for growth and dissolution of minerals using a combination of rare-event theories," *Journal of the American Chemical Society*, vol. 134, no. 1, pp. 11–14, 2012.

[68] R. Darkins, *Computational insight into the molecular mechanisms that control the growth of inorganic crystals*. PhD thesis, UCL (University College London), 2017.

[69] A. A. Chernov, "The spiral growth of crystals," *Soviet Physics Uspekhi*, vol. 4, no. 1, p. 116, 1961.

[70] K. Y. Foo and B. H. Hameed, "Insights into the modeling of adsorption isotherm systems," *Chemical Engineering Journal*, vol. 156, no. 1, pp. 2–10, 2010.

[71] G. Montanari, L. Lakshtanov, D. Tobler, K. Dideriksen, K. Dalby, N. Bovet, and S. Stipp, "Effect of aspartic acid and glycine on calcite growth," *Crystal Growth & Design*, vol. 16, no. 9, pp. 4813–4821, 2016.

[72] R. Stepic, L. Jurkovic, K. Klementyeva, M. Ukrainczyk, M. Gredicak, D. M. Smith, D. Kralj, and A.-S. Smith, "Adsorption of aspartate derivatives to calcite surfaces in aqueous environment," *Crystal Growth & Design*, vol. 20, no. 5, pp. 2853–2859, 2020.

[73] J. E. Vance, "Growth and perfection of crystals.," *Journal of the American Chemical Society*, vol. 81, no. 13, pp. 3489–3490, 1959.

[74] W. Van Enckevort and A. Van den Berg, "Impurity blocking of crystal growth: a Monte Carlo study," *Journal of Crystal Growth*, vol. 183, no. 3, pp. 441–455, 1998.

[75] K. J. Davis, P. M. Dove, and J. J. De Yoreo, "The role of Mg as an impurity in calcite growth," *Science*, vol. 290, no. 5494, pp. 1134–1137, 2000.

[76] L. E. Wasylenki, P. M. Dove, D. S. Wilson, and J. J. De Yoreo, "Nanoscale effects of strontium on calcite growth: An in situ AFM study in the absence of vital effects," *Geochimica et Cosmochimica Acta*, vol. 69, no. 12, pp. 3017–3027, 2005.

[77] J. Astilleros, C. Pina, L. Fernandez-Daz, and A. Putnis, "Molecular-scale surface processes during the growth of calcite in the presence of manganese," *Geochimica et Cosmochimica Acta*, vol. 66, no. 18, pp. 3177–3189, 2002.

[78] S. Elhadj, J. De Yoreo, J. Hoyer, and P. Dove, "Role of molecular charge and hydrophilicity in regulating the kinetics of crystal growth," *Proceedings of the National Academy of Sciences*, vol. 103, no. 51, pp. 19237–19242, 2006.

[79] F. C. Meldrum and S. T. Hyde, "Morphological influence of magnesium and organic additives on the precipitation of calcite," *Journal of Crystal Growth*, vol. 231, no. 4, pp. 544–558, 2001.

[80] C. Orme, A. Noy, A. Wierzbicki, M. McBride, M. Grantham, H. Teng, P. Dove, and J. De Yoreo, "Formation of chiral morphologies through selective binding of amino acids to calcite surface steps," *Nature*, vol. 411, no. 6839, pp. 775–779, 2001.

[81] S. Elhadj, E. Salter, A. Wierzbicki, J. De Yoreo, N. Han, and P. Dove, "Peptide controls on calcite mineralization: Polyaspartate chain length affects growth kinetics and acts as a stereochemical switch on morphology," *Crystal Growth & Design*, vol. 6, no. 1, pp. 197–201, 2006.

[82] K. M. Towe and G. R. Thompson, "The structure of some bivalve shell carbonates prepared by ion-beam thinning," *Calcified Tissue Research*, vol. 10, no. 1, pp. 38–48, 1972.

[83] A. Berman, L. Addadi, Å. Kvick, L. Leiserowitz, M. Nelson, and S. Weiner, "Intercalation of sea urchin proteins in calcite: study of a crystalline composite material," *Science*, vol. 250, no. 4981, pp. 664–667, 1990.

[84] S. Borukhin, L. Bloch, T. Radlauer, A. H. Hill, A. N. Fitch, and B. Pokroy, "Screening the incorporation of amino acids into an inorganic crystalline host: the case of calcite," *Advanced Functional Materials*, vol. 22, no. 20, pp. 4216–4224, 2012.

[85] K. Rae Cho, Y.-Y. Kim, P. Yang, W. Cai, H. Pan, A. N. Kulak, J. L. Lau, P. Kulshreshtha, S. P. Armes, F. C. Meldrum, *et al.*, "Direct observation of mineral–organic composite formation reveals occlusion mechanism," *Nature Communications*, vol. 7, no. 1, pp. 1–7, 2016.

[86] Y.-Y. Kim, R. Darkins, A. Broad, A. N. Kulak, M. A. Holden, O. Nahi, S. P. Armes, C. C. Tang, R. F. Thompson, F. Marin, *et al.*, "Hydroxyl-rich macromolecules enable the bio-inspired synthesis of single crystal nanocomposites," *Nature Communications*, vol. 10, no. 1, pp. 1–15, 2019.

[87] O. Nahi, A. Broad, A. N. Kulak, H. M. Freeman, S. Zhang, T. D. Turner, L. Roach, R. Darkins, I. J. Ford, and F. C. Meldrum, "Positively charged additives facilitate incorporation in inorganic single crystals," *arXiv preprint arXiv:2201.03298*, 2022.

[88] B. Marzec, D. C. Green, M. A. Holden, A. S. Coté, J. Ihli, S. Khalid, A. Kulak, D. Walker, C. Tang, D. M. Duffy, *et al.*, "Amino acid assisted incorporation of dye molecules within calcite crystals," *Angewandte Chemie International Edition*, vol. 57, no. 28, pp. 8623–8628, 2018.

[89] O. Nahi, A. N. Kulak, A. Broad, Y. Xu, C. OŚhaughnessy, O. J. Cayre, S. J. Day, R. Darkins, and F. C. Meldrum, "Solvent-mediated enhancement of additive-controlled crystallization," *Crystal Growth & Design*, vol. 21, no. 12, pp. 7104–7115, 2021.

[90] U. G. Wegst, H. Bai, E. Saiz, A. P. Tomsia, and R. O. Ritchie, "Bioinspired structural materials," *Nature Materials*, vol. 14, no. 1, pp. 23–36, 2015.

[91] A. J. Andersson, F. T. Mackenzie, and N. R. Bates, "Life on the margin: implications of ocean acidification on mg-calcite, high latitude and cold-water marine calcifiers," *Marine Ecology Progress Series*, vol. 373, pp. 265–273, 2008.

[92] Y.-Y. Kim, J. D. Carloni, B. Demarchi, D. Sparks, D. G. Reid, M. E. Kunitake, C. C. Tang, M. J. Duer, C. L. Freeman, B. Pokroy, *et al.*, "Tuning hardness in calcite by incorporation of amino acids," *Nature Materials*, vol. 15, no. 8, pp. 903–910, 2016.

[93] J. Titiloye, S. Parker, and S. Mann, "Atomistic simulation of calcite surfaces and the influence of growth additives on their morphology," *Journal of Crystal Growth*, vol. 131, no. 3-4, pp. 533–545, 1993.

[94] A. Pavese, M. Catti, G. D. Price, and R. Jackson, "Interatomic potentials for CaCO 3 polymorphs (calcite and aragonite), fitted to elastic and vibrational data," *Physics and Chemistry of Minerals*, vol. 19, pp. 80–87, 1992.

[95] T. Gilbert, "Soft-sphere model for closed-shell atoms and ions," *The Journal of Chemical Physics*, vol. 49, no. 6, pp. 2640–2642, 1968.

[96] U. Aschauer, D. Spagnoli, P. Bowen, and S. C. Parker, "Growth modification of seeded calcite using carboxylic acids: Atomistic simulations," *Journal of Colloid and Interface Science*, vol. 346, no. 1, pp. 226–231, 2010.

[97] H. Nada, "Difference in the conformation and dynamics of aspartic acid on the flat regions, step edges, and kinks of a calcite surface: a molecular dynamics study," *The Journal of Physical Chemistry C*, vol. 118, no. 26, pp. 14335–14345, 2014.

[98] J. Aufort, A. Schuitemaker, R. Green, R. Demichelis, P. Raiteri, and J. D. Gale, "Determining the adsorption free energies of small organic molecules and intrinsic ions at the terrace and steps of calcite," *Crystal Growth & Design*, vol. 22, no. 2, pp. 1445–1458, 2022.

[99] A. S. Cote, R. Darkins, and D. M. Duffy, "Deformation twinning and the role of amino acids and magnesium in calcite hardness from molecular simulation," *Phys. Chem. Chem. Phys.*, vol. 17, pp. 20178–20184, 2015.

[100] P. P. Ewald, "Die berechnung optischer und elektrostatischer gitterpotentiale," *Annalen der PPhysics and chemistry of mineralshysik*, vol. 369, no. 3, pp. 253–287, 1921.

[101] J. W. Eastwood, R. W. Hockney, and D. Lawrence, "P3m3dp—the three-dimensional periodic particle-particle/particle-mesh program," *Computer Physics Communications*, vol. 19, no. 2, pp. 215–261, 1980.

[102] E. Braun, S. M. Moosavi, and B. Smit, "Anomalous effects of velocity rescaling algorithms: the flying ice cube effect revisited," *Journal of Chemical Theory and Computation*, vol. 14, no. 10, pp. 5262–5272, 2018.

[103] D. J. Evans and B. L. Holian, "The Nosé–Hoover thermostat," *The Journal of Chemical Physics*, vol. 83, no. 8, pp. 4069–4074, 1985.

[104] S. Melchionna, G. Ciccotti, and B. Lee Holian, "Hoover NPT dynamics for systems varying in shape and size," *Molecular Physics*, vol. 78, no. 3, pp. 533–544, 1993.

[105] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, "The weighted histogram analysis method for free-energy calculations on biomolecules. i. the method," *Journal of Computational Chemistry*, vol. 13, no. 8, pp. 1011–1021, 1992.

[106] A. Grossfield, "Wham: the weighted histogram analysis method, version 2.0.11."

[107] V. Marinova and M. Salvalaglio, "Time-independent free energies from metadynamics via mean force integration," *The Journal of Chemical Physics*, vol. 151, no. 16, p. 164115, 2019.

[108] G. E. Crooks, "Nonequilibrium measurements of free energy differences for microscopically reversible markovian systems," *Journal of Statistical Physics*, vol. 90, no. 5, pp. 1481–1487, 1998.

[109] C. Jarzynski, "Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach," *Physical Review E*, vol. 56, no. 5, p. 5018, 1997.

[110] H. Y. Tang and I. J. Ford, "Free energies of molecular clusters determined by guided mechanical disassembly," *Physical Review E*, vol. 91, no. 2, p. 023308, 2015.

[111] L. Maragliano, A. Fischer, E. Vanden-Eijnden, and G. Ciccotti, "String method in collective variables: Minimum free energy paths and isocommittor surfaces," *The Journal of Chemical Physics*, vol. 125, no. 2, p. 024106, 2006.

[112] E. Weinan, W. Ren, and E. Vanden-Eijnden, "Simplified and improved string method for computing the minimum energy paths in barrier-crossing events," *Journal of Chemical Physics*, vol. 126, no. 16, p. 164103, 2007.

[113] G. Henkelman, B. P. Uberuaga, and H. Jónsson, "A climbing image nudged elastic band method for finding saddle points and minimum energy paths," *The Journal of Chemical Physics*, vol. 113, no. 22, pp. 9901–9904, 2000.

[114] E. Vanden-Eijnden *et al.*, "Towards a theory of transition paths," *Journal of statistical physics*, vol. 123, no. 3, pp. 503–523, 2006.

[115] E. Vanden-Eijnden, "Transition path theory," in *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, pp. 453–493, Springer, 2006.

[116] A. F. Voter, "Introduction to the kinetic Monte Carlo method," in *Radiation Effects in Solids*, pp. 1–23, Springer, 2007.

[117] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, "LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales," *Comp. Phys. Comm.*, vol. 271, p. 108171, 2022.

[118] Y. Wu, H. L. Tepper, and G. A. Voth, "Flexible simple point-charge water model with improved liquid-state properties," *The Journal of Chemical Physics*, vol. 124, no. 2, p. 024503, 2006.

[119] P. S. Branicio and D. J. Srolovitz, "Local stress calculation in simulations of multicomponent systems," *J. Comput. Phys.*, vol. 228, pp. 8467–8479, 2009.

[120] R. Darkins, M. L. Sushko, J. Liu, and D. M. Duffy, "Stress in titania nanoparticles: an atomistic study," *Phys. Chem. Chem. Phys.*, vol. 16, pp. 9441–9447, 2014.

[121] G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, "Plumed 2: New feathers for an old bird," *Computer Physics Communications*, vol. 185, no. 2, pp. 604–613, 2014.

[122] P. Raiteri, A. Laio, F. L. Gervasio, C. Micheletti, and M. Parrinello, "Efficient reconstruction of complex free energy landscapes by multiple walkers metadynamics," *The Journal of Physical Chemistry B*, vol. 110, no. 8, pp. 3533–3539, 2006.

[123] M. Hong and H. H. Teng, "Implications of solution chemistry effects: Direction-specific restraints on the step kinetics of calcite growth," *Geochimica et Cosmochimica Acta*, vol. 141, pp. 228–239, 2014.

[124] M. N. Joswiak, M. F. Doherty, and B. Peters, "Ion dissolution mechanism and kinetics at kink sites on NaCl surfaces," *Proceedings of the National Academy of Sciences*, vol. 115, no. 4, pp. 656–661, 2018.

[125] R. Impey, P. Madden, and I. McDonald, "Hydration and mobility of ions in solution," *The Journal of Physical Chemistry*, vol. 87, no. 25, pp. 5071–5083, 1983.

[126] S. Piana, F. Jones, and J. D. Gale, "Assisted desolvation as a key kinetic step for crystal growth," *Journal of the American Chemical Society*, vol. 128, no. 41, pp. 13568–13574, 2006.

[127] Y.-Y. Kim, L. A. Fielding, A. N. Kulak, O. Nahi, W. Mercer, E. R. Jones, S. P. Armes, and F. C. Meldrum, "Influence of the structure of block copolymer nanoparticles on the growth of calcium carbonate," *Chemistry of Materials*, vol. 30, no. 20, pp. 7091–7099, 2018.

[128] J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, and D. A. Case, "Development and testing of a general AMBER force field," *Journal of Computational Chemistry*, vol. 25, no. 9, pp. 1157–1174, 2004.

[129] J. Wang, W. Wang, P. A. Kollman, and D. A. Case, "Antechamber: an accessory software package for molecular mechanical calculations," *J. Am. Chem. Soc*, vol. 222, p. U403, 2001.

[130] P. Raiteri, R. Demichelis, J. D. Gale, M. Kellermeier, D. Gebauer, D. Quigley, L. B. Wright, and T. R. Walsh, "Exploring the influence of organic species on pre-and post-nucleation calcium carbonate," *Faraday Discussions*, vol. 159, no. 1, pp. 61–85, 2012.

[131] R. Momper, M. Nalbach, K. Lichtenstein, R. Bechstein, and A. Kühnle, "Stabilization of polar step edges on calcite (10.4) by the adsorption of congo red," *Langmuir*, vol. 31, no. 26, pp. 7283–7287, 2015.

[132] H. H. Teng, P. M. Dove, and J. J. De Yoreo, "Reversed calcite morphologies induced by microscopic growth kinetics: insight into biomineralization," *Geochimica et Cosmochimica Acta*, vol. 63, no. 17, pp. 2507–2512, 1999.

[133] S. Mann, *Biomineralization: principles and concepts in bioinorganic materials chemistry*, vol. 5. Oxford University Press on Demand, 2001.

[134] L. B. Gower, "Biomimetic model systems for investigating the amorphous precursor pathway and its role in biomineralization," *Chemical reviews*, vol. 108, no. 11, pp. 4551–4627, 2008.

[135] J. Aizenberg, A. J. Black, and G. M. Whitesides, "Control of crystal nucleation by patterned self-assembled monolayers," *Nature*, vol. 398, no. 6727, pp. 495–498, 1999.

[136] G. Falini, S. Albeck, S. Weiner, and L. Addadi, "Control of aragonite or calcite polymorphism by mollusk shell macromolecules," *Science*, vol. 271, no. 5245, pp. 67–69, 1996.

[137] C.-C. Chen, C.-C. Lin, L.-G. Liu, S. V. Sinogeikin, and J. D. Bass, "Elasticity of single-crystal calcite and rhodochrosite by brillouin spectroscopy," *American Mineralogist*, vol. 86, no. 11-12, pp. 1525–1529, 2001.

[138] J. Aizenberg, A. Tkachenko, S. Weiner, L. Addadi, and G. Hendler, "Calcitic microlenses as part of the photoreceptor system in brittlestars," *Nature*, vol. 412, no. 23, pp. 819–822, 2001.

[139] E. Seknazi, S. Kozachkevich, I. Polishchuk, N. B. Stein, J. Villanova, J.-P. Suuronen, C. Dejoie, P. Zaslansky, A. Katsman, and B. Pokroy, "From spinodal decomposition to alternating layered structure within single crystals of biogenic magnesium calcite," *Nature Communications*, vol. 10, no. 1, pp. 1–9, 2019.

[140] A. Guinier, "Structure of age-hardened aluminium-copper alloys," *Nature*, vol. 142, pp. 569–570, 1938.

[141] G. D. Preston, "The diffraction of x-rays by age-hardening aluminium copper alloys," *Proc. R. Soc. A*, vol. 167, p. 526, 1938.

[142] W. Verestek, A.-P. Prskalo, M. Hummel, P. Binkele, and S. Schmauder, "Molecular dynamics investigations of the strengthening of Al-Cu alloys during thermal ageing," *Physical Mesomechanics*, vol. 20, no. 3, pp. 291–304, 2017.

[143] C. Singh and D. Warner, "Mechanisms of Guinier–Preston zone hardening in the athermal limit," *Acta Materialia*, vol. 58, no. 17, pp. 5797–5805, 2010.

[144] S. P. Ringer and K. Hono, "Microstructural evolution and age hardening in aluminium alloys: Atom probe field-ion microscopy and transmission electron microscopy studies," *Mater. Charact.*, vol. 44, pp. 101–131, 2000.

[145] D. M. Duffy, "Coherent nanoparticles in calcite," *Science*, vol. 358, no. 6368, pp. 1254–1255, 2017.

[146] A. Pavese, M. Catti, S. C. Parker, and A. Wall, "Modelling of the thermal dependence of structural and elastic properties of calcite caco3," *Phys. Chem. Miner.*, vol. 23, pp. 89–93, 1996.

[147] E. Lenzi, L. Evangelista, M. Lenzi, H. Ribeiro, and E. C. de Oliveira, "Solutions for a non-Markovian diffusion equation," *Physics Letters A*, vol. 374, no. 41, pp. 4193–4198, 2010.

[148] A. Mura, M. S. Taqqu, and F. Mainardi, "Non-Markovian diffusion equations and processes: analysis and simulations," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 21, pp. 5033–5064, 2008.

[149] P. E. Smith, "The alanine dipeptide free energy surface in solution," *The Journal of Chemical Physics*, vol. 111, no. 12, pp. 5568–5579, 1999.

[150] B. M. Pettitt and M. Karplus, "Conformational free energy of hydration for the alanine dipeptide: thermodynamic analysis," *The Journal of Physical Chemistry*, vol. 92, no. 13, pp. 3994–3997, 1988.

[151] A. C. Lasaga, "Transition state theory," *Rev. Mineral.;(United States)*, vol. 8, 1981.

[152] R. Becker and W. Döring, "Kinetische behandlung der keimbildung in übersättigten dämpfen," *Annalen der Physik*, vol. 416, no. 8, pp. 719–752, 1935.

[153] R. Henery, "The generalized Langevin equation and the fluctuation-dissipation theorems," *Journal of Physics A: General Physics*, vol. 4, no. 5, p. 685, 1971.

[154] M. Boguná, L. F. Lafuerza, R. Toral, and M. Á. Serrano, "Simulating non-Markovian stochastic processes," *Physical Review E*, vol. 90, no. 4, p. 042108, 2014.

[155] N. Masuda and L. E. Rocha, "A Gillespie algorithm for non-Markovian stochastic processes," *SIAM Review*, vol. 60, no. 1, pp. 95–115, 2018.

Appendix

A.1 Derivations

A.1.1 Equivalence of equations 8.5 and 8.6

Equation 8.6 states that,

$$\Gamma_{x_i \to x_{i+1}} = \left\langle \mathcal{T}_{x_i \to x_{i+1}} \right\rangle^{-1} \tag{A1}$$

where $\mathcal{T}_{x_i \to x_{i+1}}$ is the time taken for an $x_i \to x_{i+1}$ reaction to take place. Now, consider these transitions occurring over many different y-values, each with its own average time taken for an $x_i \to x_{i+1}$ reaction to take place. We can denote this average time taken as $\left\langle \mathcal{T}_{x_i \to x_{i+1}, y_j} \right\rangle$ where y_j is the y-coordinate. Equation 8.6 is therefore the average of $\left\langle \mathcal{T}_{x_i \to x_{i+1}, y_j} \right\rangle$ over all j values. However, one further detail is that the number of times an $x_i \to x_{i+1}, y_j$ reaction rate is recorded in a simulation is proportional to the probability of finding the y-coordinate at y_j, given that the x-coordinate is at x_i. This is denoted by $\mathcal{P}_{j|i}$ in the main text. The number of times a reaction will occur in a simulation is also proportional to the average $x_i \to x_{i+1}$ reaction rate given a y-coordinate of y_j, $\Gamma_{x_i \to x_{i+1}, y_j}$. Equation 8.6 therefore simplifies to

$$\Gamma_{x_i \to x_{i+1}} = \left(\frac{\sum_j \left\langle \mathcal{T}_{x_i \to x_{i+1}, y_j} \right\rangle \mathcal{P}_{j|i} \Gamma_{x_i \to x_{i+1}, y_j}}{\sum_{j'} \mathcal{P}_{j'|i} \Gamma_{x_i \to x_{i+1}, y_{j'}}} \right)^{-1} \tag{A2}$$

Now, note that $\Gamma_{x_i \to x_{i+1}, y_j}$ is simply the inverse of $\left\langle \mathcal{T}_{x_i \to x_{i+1}, y_j} \right\rangle$ by definition. Also note that $\sum_{j'} \mathcal{P}_{j'|i} = 1$. Equation A2 therefore simplifies to

$$\Gamma_{x_i \to x_{i+1}} = \left(\frac{1}{\sum_{j'} \mathcal{P}_{j'|i} \Gamma_{x_i \to x_{i+1}, y_{j'}}} \right)^{-1} \tag{A3}$$

which is equation 8.5.

A.2 Figures

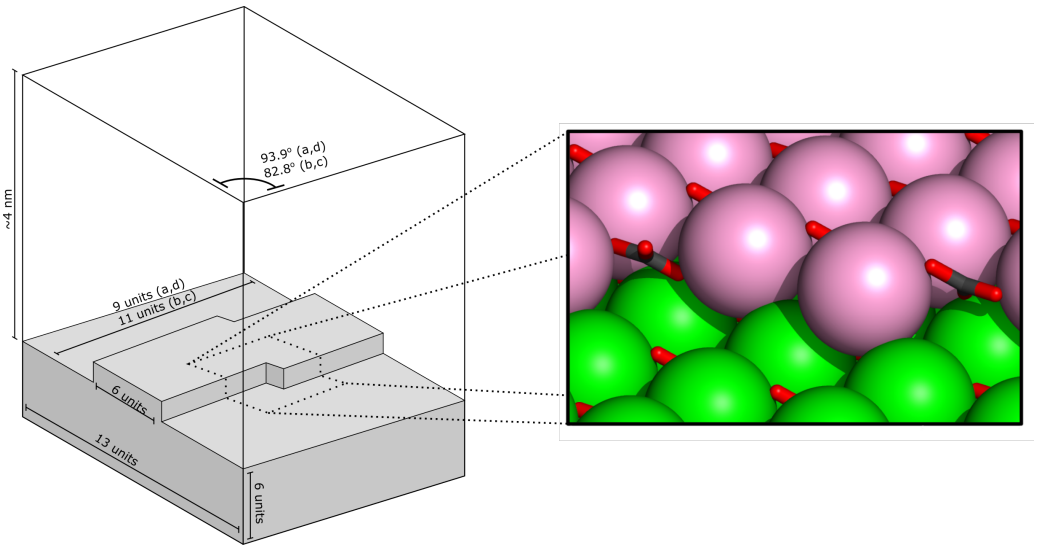

Figure A1: Schematic of the simulation cell setup designed to isolate a particular kink site in simulations. The cell dimensions are also labelled.

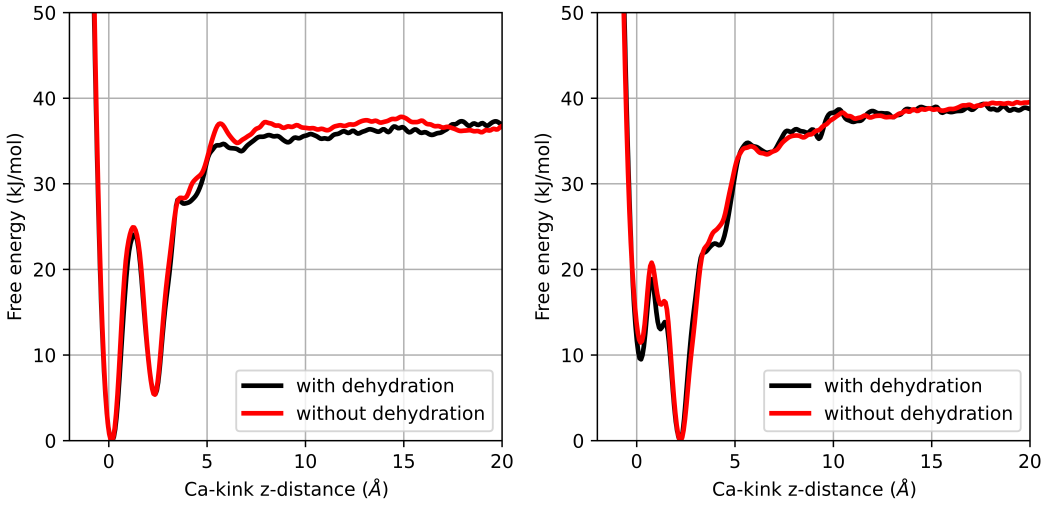

Figure A2: Free energy surfaces for Ca ions as a function of distance from their respective kink sites for the a(ii) (left) and c(i) (right) kinks. A good agreement is found between simulations where dehydration is included as a reaction coordinate, and where it isn't.

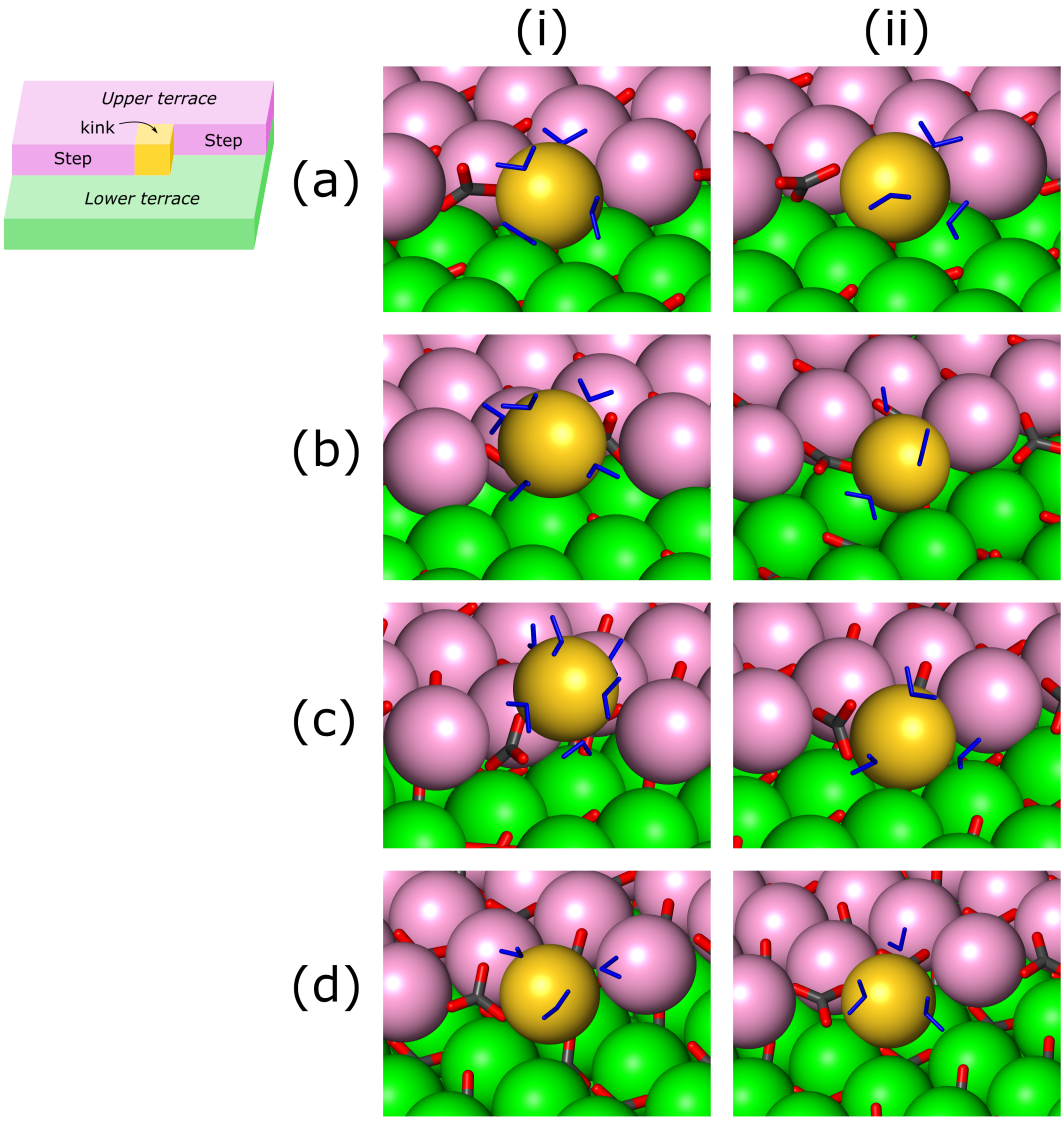

Figure A3: Snapshots from simulations showing the strongest binding configurations for every Ca terminated kink. The water molecules coordinated with the Ca kink are also shown.

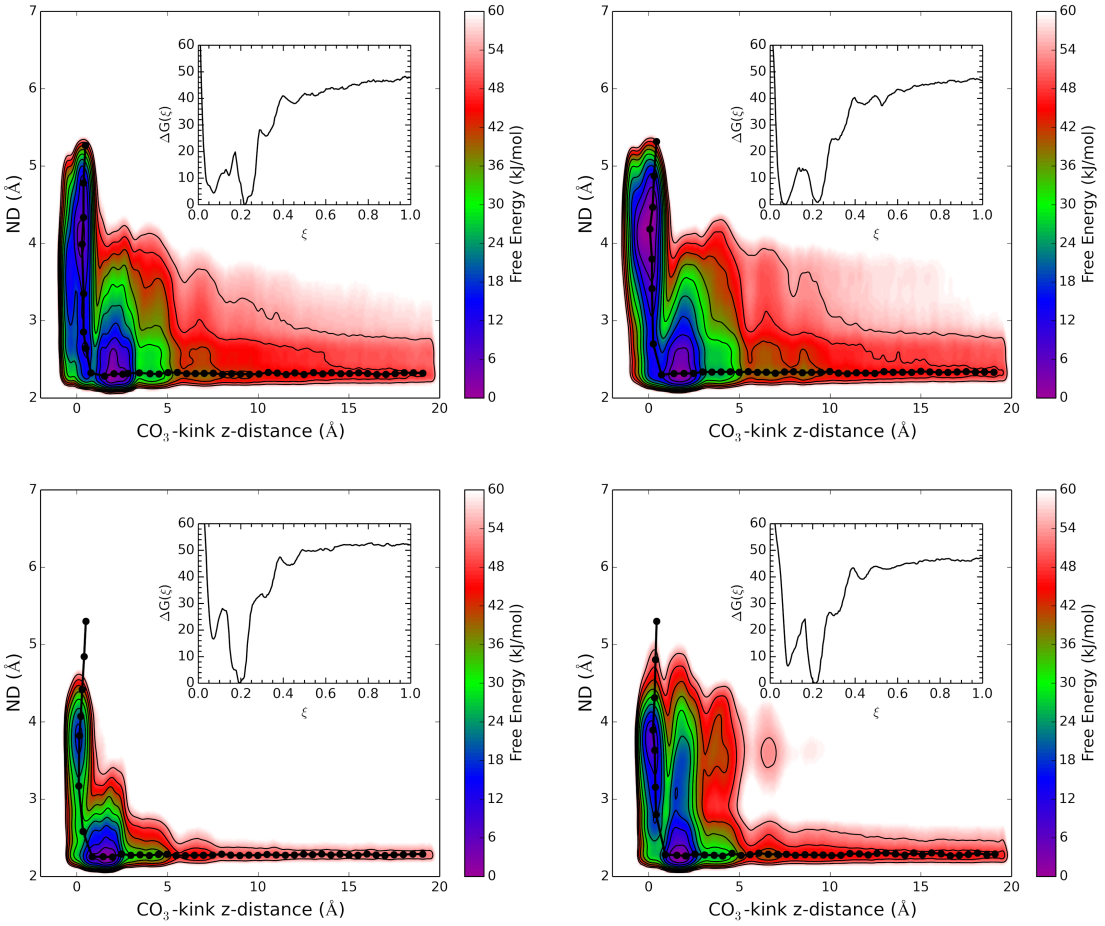

Figure A4: Free energy surfaces for CO_3-terminated kinks as a function of CO_3-kink z-distance and ND (see equation 5.1) for a(i) (top-left), a(ii) (top-right), d(i) (bottom-left), and d(ii) (bottom-right) kinks. The connected black dots trace the MFEP, which is also shown in the insets.

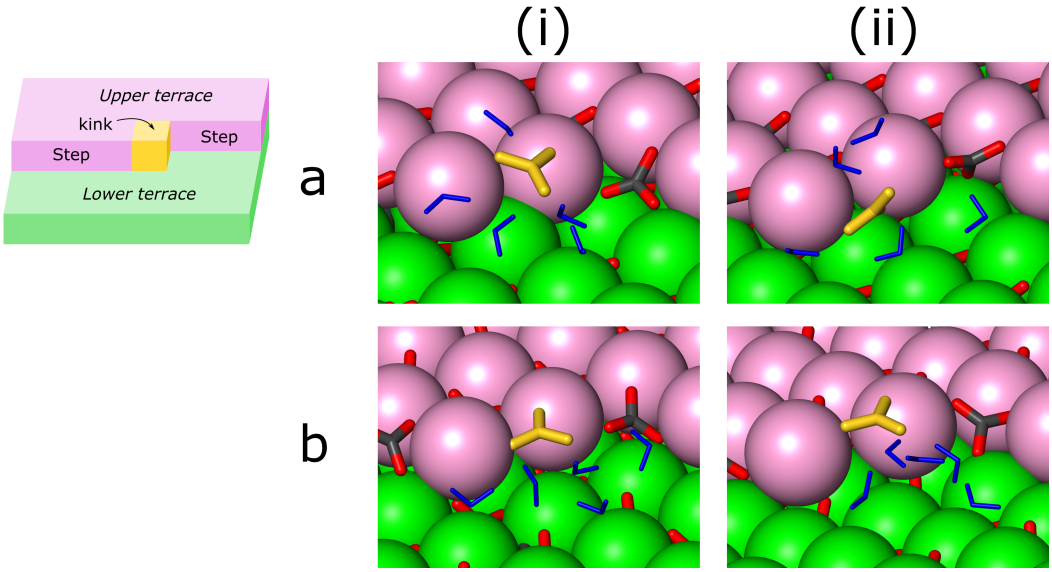

Figure A5: Snapshots from simulations showing the strongest binding configurations for every CO_3 terminated kink. The nearest five water molecules to the kink site are also shown.

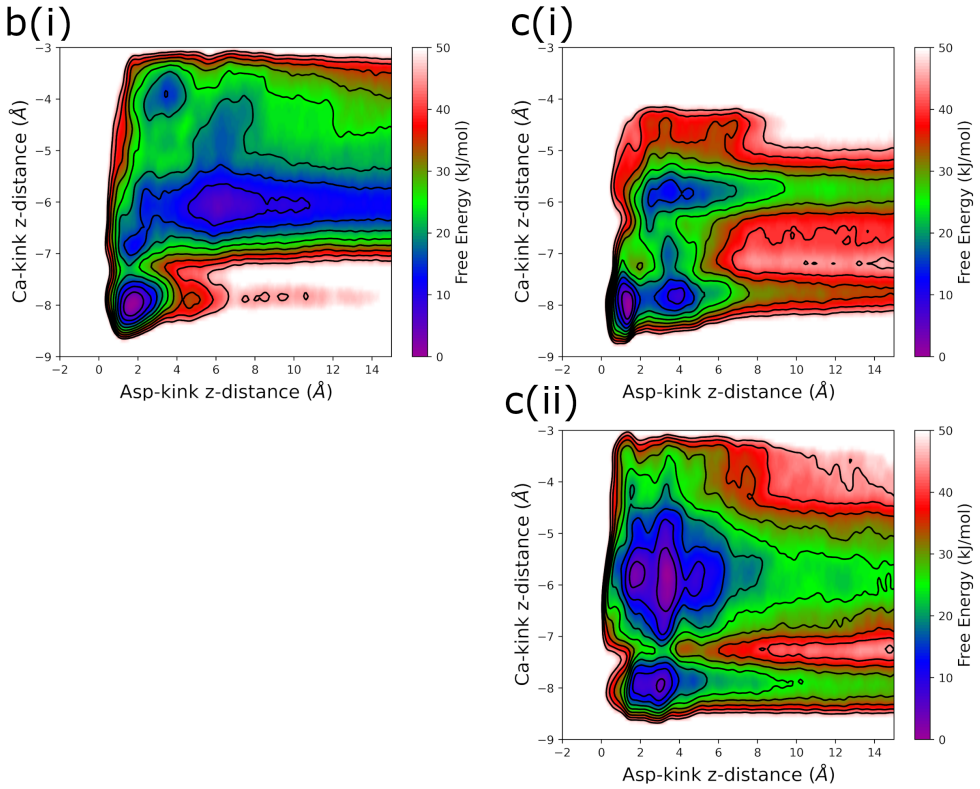

Figure A6: Two-dimensional free energy surfaces as a function of the asp-kink z-distance (x-axis) and the terminating Ca-kink z-distance (y-axis). The y-axis is integrated out using equation 3.90 to produce Figure 6.8.

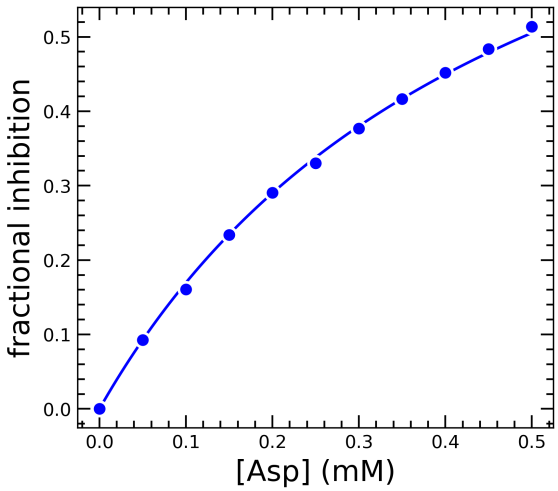

Figure A7: Fractional inhibition of the normal growth rate. The Langmuir equation is used to calculate ΔG_{ads} as shown in equation 6.5.

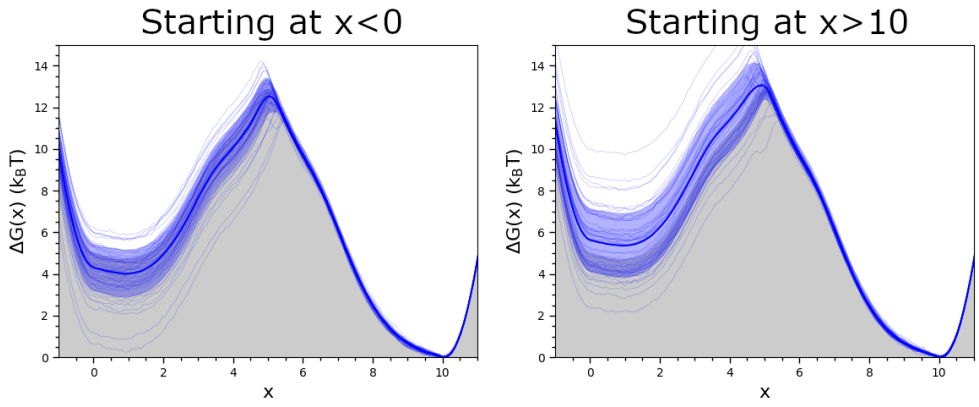

Figure A8: Derived PMF from 50 umbrella sampling simulations. Unlike Figure 8.7, the initial configurations are chosen by dragging the particle along the x-direction while allowing the particle to relax in the y-direction. This method produces an offset of the average of the derived PMF dependent on the side of the barrier at which the simulation is started.