Michael Grünwald

Transition path sampling simulations of structural phase transformations in nanocrystals under pressure

Dissertation

March 27, 2009

Universität Wien

To Laura and Emil.

Abstract

This dissertation describes a computer simulation study of structural transformations in nanocrystals under pressure. The mechanisms of such transformations can be quite different from the bulk, they determine the shape change during the transformation and influence the accessibility and stability of particular crystal structures. Yet the atomistic transformation details elude direct observation by experiment due to limited time and space resolution. Molecular dynamics computer simulations, on the other hand, can provide the necessary atomistic perspective but are restricted to the nanosecond timescale which requires the use of significantly higher pressures. Such deviation from experimental conditions compromises the comparability of simulation and experiment considerably. Here, we use transition path sampling, a simulation method designed to overcome the timescale problem. We introduce a new transition path sampling algorithm for nanoparticles under pressure, which features an efficient barostat of ideal gas particles. We show that the algorithm satisfies detailed balance and apply it to a structural transformation in a model of CdSe nanocrystals. Starting from a previously proposed mechanism, the algorithm quickly finds a more favorable transformation route which is characterized by a consecutive sliding of parallel crystal planes. We subject the pathways obtained with transition path sampling to a comprehensive transition state analysis based on the committor function. These calculations reveal the critical nuclei of the high pressure phase which we quantify by calculating activation enthalpies and volumes. The dependence of these quantities on crystal size shows the same scaling observed experimentally and is consistent with the peculiar shape of the critical nucleus which markedly differs from the bulk. Furthermore, a new algorithm is introduced to the transition path sampling methodology, which exploits the linear dynamics of small phase space displacements and greatly enhances the sampling efficiency of transition pathways in systems involving long and rough free energy barriers.

Zusammenfassung

Diese Arbeit beschreibt eine Computersimulationsstudie von druckinduzierten strukturellen Phasenübergängen in Nanokristallen. Die Mechanismen solcher Übergänge können sich stark von den entsprechenden Mechanismen im ausgedehnten Festkörper unterscheiden; sie bestimmen die Formänderung des Kristalls während der Transformation und beeinflussen die Erreichbarkeit und Stabilität einzelner Kristallstrukturen. Das zeitliche und örtliche Auflösungsvermögen von Experimenten ist jedoch zu gering, um die atomaren Details einer solchen Transformation direkt beobachten zu können. Computersimulationen bieten im Prinzip den notwendige atomaren Blickwinkel, sind aber auf eine Zeitskala von ein paar Nanosekunden eingeschränkt. Um in dieser kurzen Zeit eine Transformation beobachten zu können, muss ein wesentlich höherer Druck verwendet werden als im Experiment — die Aussagekraft der Simulation wird dadurch erheblich beeinträchtigt. In dieser Arbeit wird deshalb die transition path sampling Methode verwendet, mit der das Zeitskalenproblem umgangen wird und Simulationen unter experimentellen Bedingungen durchgeführt werden können. Zur Simulation von Nanokristallen unter Druck entwickeln wir einen neuen Algorithmus, in dem ein Druckbad aus idealen Gasteilchen verwendet wird. Der Algorithmus erzeugt die korrekte Wahrscheinlichkeitsdichte eines Systems bei konstantem Druck und konstanter Temperatur und wird auf eine Transformation in CdSe Nanokristallen angewandt. Unsere Simulationen identifizieren den dominanten Transformationsmechanismus und ermöglichen darüber hinaus die Bestimmung des kritischen Nukleus der Hochdruckphase durch eine umfassende statistische Analyse der simulierten Transformationspfade. Wir bestimmen Aktivierungsenthalpie und -volumen als Funktion des Drucks und vergleichen diese Größen mit dem Experiment. Die im Experiment beobachtete, annähernd lineare Abhängigkeit der Aktivierungsenthalpie vom Kristalldurchmesser stimmt gut mit unseren Simulationen überein und kann durch die spezielle längliche Geometrie des kritischen Nukleus erklärt werden. Des Weiteren wird ein neuer Algorithmus vorgestellt, der die Effizienz von transition path sampling Simulationen von Systemen mit langen und rauen freien Energiebarrieren entscheidend verbessert.

Preface

The following research articles are contained in this work, entirely or in part.

- Chapter 2:
  M. Grünwald and C. Dellago, *Transition path sampling studies of solid-solid transformations in nanocrystals under pressure*, to appear in *Trends in Computational Nanomechanics: Transcending Time and Space* (Springer Series, 2009).

- Chapter 3:
  M. Grünwald, Phillip L. Geissler, and C. Dellago, *Precision shooting: Sampling long transition pathways*, J. Chem. Phys. **129**, 194101 (2008).

- Chapter 4:
  M. Grünwald, E. Rabani and C. Dellago, *Mechanisms of the wurtzite to rocksalt transformation in CdSe nanocrystals*, Phys. Rev. Lett. **96**, 255701 (2006).

- Chapter 5:
  M. Grünwald, Phillip L. Geissler, and C. Dellago, *An efficient transition path sampling algorithm for nanoparticles under pressure*, J. Chem. Phys. **127**, 154718 (2007).

- Chapter 6:
  M. Grünwald and C. Dellago, *Transition state analysis of solid-solid transformations in nanocrystals*, submitted to J. Chem. Phys. (2009).

- Appendix B:
  M. Grünwald and C. Dellago, *Nucleation and growth in structural transformations of nanocrystals*, Nano Letters, in print (2009).

The following paper is not included.

- M. Grünwald and C. Dellago, *Ideal gas pressure bath: a method for applying hydrostatic pressure in the computer simulation of nanoparticles*, Mol. Phys. **104**, 3709 (2006).

Acknowledgements

After more than four years as a member of the Dellago group, there are quite a number of people who contributed to my work in some way or another and who I am indebted to.

First of all, Christoph Dellago has not only provided me with constant guidance, support, and encouragement in times of uncertainty, but also gave me the opportunity to learn from his directness and candor in areas off-science. He has my respect for being a great physicist and a good friend alike.

During my time in Berkeley, Phillip Geissler took over the job of diminishing my ignorance. I am thankful for his support and patience with me.

I want to thank Paul Alivisatos, Maxwell Merkle, Eran Rabani, and Josh Wittemberg for fruitful scientific discussions and collaborations. For constant financial support I am indebted to the Science College "Computational Materials Science" under the direction of Jürgen Hafner.

When spending a significant fraction of your life in a room together with the same group of people, they better be such lovely persons and good friends as Bernhard Reischl, Harald Oberhofer, Wolfgang Lechner and Jürgen Köfinger. I am certain we all will be looking back at the good times we had together.

Finally, none of this would make sense without my wife Laura, who I am deeply in love with.

Vienna, March 2009 *Michael Grünwald*

Contents

Chapter 1

Introduction

The chief goal of nanoscience is the design of materials with novel properties tailored to particular technological needs. This is achieved by exploiting the physical and chemical properties of small building blocks: When a piece of macroscopic matter is scaled down to the nanometer regime, size dependent deviations from the properties of the corresponding bulk material emerge due to the rapidly increasing fraction of surface particles. Most commonly, these deviations are of quantitative nature. For instance, the melting point of semiconductor nanocrystals decreases while their band gap and the pressure that is needed to induce structural transformations increases considerably with decreasing size [1–3].

From both a scientific and technological point of view, however, the most interesting size-related phenomena are of qualitative nature. Some of the most striking examples of such behavior can be found in solid state materials, where the emergence of fullerenes and clusters of icosahedral symmetry are examples of morphologies that exist only on the nanoscale. To explore and make technological use of any such structure, both the means of production and handling, as well as a sound understanding of its thermodynamic and kinetic stability must be developed. To this end, sophisticated methods have recently emerged for the synthesis of semiconductor nanocrystals, most notably of the material CdSe [4]. Through careful control of chemical catalysts and environmental conditions, ensembles of nanocrystals with narrow size distribution and distinct shapes and surface configurations are routinely produced and used in experiments [5–7].

Arguably one of the most in-depth studies of a size-dependent phenomenon was carried out by the group of Paul Alivisatos at UC Berkeley with their experimental investigation of the wurtzite to rocksalt transformation in CdSe nanocrystals in the size range of a few nanometers (see Fig. 1.1 for an illustration). Using diamond anvil cell techniques, they measured the size-dependence of the hysteresis curve [2, 8], determined activation energies and volumes from the pressure and temperature dependence of the rate con-

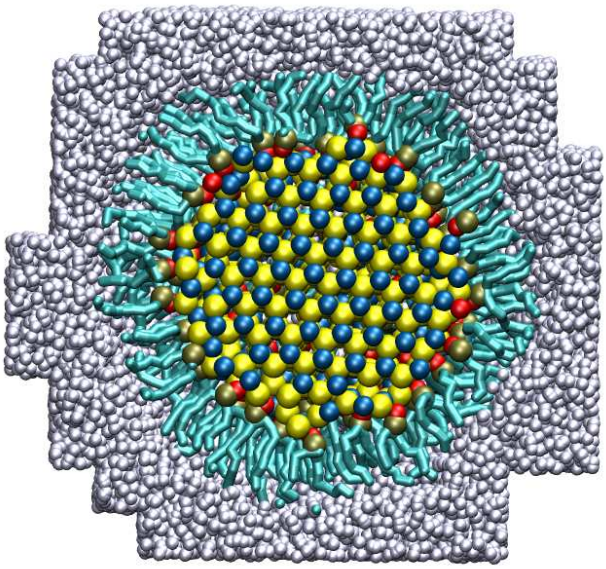

Figure 1.1: Crossection of a CdSe nanocrystal in the wurtzite structure, immersed in a pressure bath of ideal gas particles. The surface of the crystal is passivated with TOPO molecules.

stant [9, 10], investigated the shape change during transformation [11], and explored possible routes to metastability of the high pressure structure at ambient conditions [12]. This experimental evidence paints a complete picture of the thermodynamics and kinetics of nanocrystal transformations and highlights the dominant role of surface configuration and shape. Nevertheless, the specific atomistic transformation mechanism which lies at the heart of these phenomena eludes determination owing to the current limits of experimental time and space resolution. Only through the precise knowledge of the atomistic pathways, however, can differences in surface free energies between structures, changes in shape, and the metastability of particular morphologies be fully understood and, by blocking or facilitating of such pathways, possibly exploited for applications.

A number of molecular dynamics computer simulation studies have been conducted in an attempt to fill this gap [13–16]. They all follow a similar routine: A single nanocrystal is modeled *ab initio* or with an empirical pair potential, immersed in a suitable pressure medium and subjected to increasing pressure until the transformation occurs spontaneously in the simulation. The atomistic transformation mechanism can then be identified by simple visual inspection of the simulation run. However, this straightforward approach is plagued by a problem commonly encountered in the computer simulation of processes involving large kinetic barriers. The time that can

be spanned within a molecular dynamics simulation is determined by the timescale of the system's fastest molecular motion which, in order to be faithfully reproduced, limits the size of the time step used to integrate the equations of motion. For typical molecular systems, simulation runs covering the nanosecond to microsecond regime are state of the art; the relevant experiments, on the other hand, can easily last hours. The consequences of this gap are severe: To observe a transformation on the simulation timescale, the pressure has to be increased close to the point where the low pressure structure becomes mechanically unstable. Transformations occurring under these extreme conditions need not follow the same mechanistic pathways which might be observed in an experiment, where strong driving of the process is replaced by long waiting times. Thus, direct comparability of simulation and experiment is lost. Furthermore, even if the atomistic pathways were qualitatively similar in both cases, a meaningful comparison could still be only achieved through calculation of quantities that are good indicators of the actual mechanism, like the reaction rate constant, or activation energies and volumes. Again, these quantities depend sensitively on the conditions under which the transformation is observed. As a result, simulation studies of structural transformations in nanocrystals so far were reduced to a mere description of mechanisms observed using elevated pressure.

The aim of this work is to overcome this unsatisfying situation in a twofold manner. First, we circumvent the timescale problem with transition path sampling [17, 18], a set of computational techniques designed for the simulation of rare events which exploits the fact that, although activated events are *rare* on the simulation timescale, they are not necessarily *long* in duration. In fact, the transformation of a nanocrystal typically proceeds within tens of picoseconds and it is merely the time spent waiting for its spontaneous occurrence that renders it inaccessible to straightforward simulation. By eliminating the waiting times between rare events, transition path sampling allows the observation of the transformation mechanism under experimental conditions. Based on the pathways observed with transition path sampling, we then make direct contact with experiment by calculating the pressure and size dependence of activation enthalpy and volume. This is achieved by identifying the central part of transformation, the critical nucleus of the high pressure phase.

This thesis is organized as follows: In Part I we are concerned with the transition path sampling method itself. Chapter 2 gives a brief introduction to the main concepts of transition path sampling. In Chapter 3, we introduce a new algorithm which considerably enhances the sampling of transition pathways in systems that involve free energy barriers that are both long and rough. This is achieved by making use of the linear short time dynamics of small perturbations in phase space. In Part II we turn to the application of transition path sampling to the transformation of CdSe nanocrystals under pressure. Chapter 4 gives a brief account of the transformation mecha-

nisms occurring in straightforward molecular dynamics simulation. Despite involving conditions of elevated pressure, this simulation study highlights the strong dependence of the transformation mechanism on crystal size and shape and sets the stage for the transition path sampling simulations discussed in the following chapters. In Chapter 5 we extend transition path sampling to include a barostat of ideal gas particles, a method particularly designed for the simulation of nanoparticles under pressure. The algorithm is discussed in detail, and a proof of detailed balance is given as a prerequisite for the method's use in the framework of transition path sampling. The algorithm is then applied to the transformation of CdSe nanocrystals: We identify the main mechanistic route and succeed in explicitly excluding a mechanism previously suggested by an experimental study. Finally, in Chapter 6, we conduct an extensive transition state analysis of the pathways observed with transition path sampling. We identify the transition states which include the critical nuclei and discuss their peculiar dependence on crystal size and pressure. From the transition states we calculate activation enthalpies and volumes and compare them to experimental data. Our simulations explain the size trend observed in experiments and emphasize the highly anisotropic character of the nucleation event in nanocrystals, which markedly differs from comparable transformation behavior in bulk materials.

Part I

Transition path sampling

Chapter 2

Rare events in computer simulations

> *"Erst in letzter Zeit scheint einiges Licht aus einer ganz unerwarteten Ecke zu kommen, nämlich der Spielerei mit Computern, bei der jeder auch ohne große Vorkenntnisse fundamental Neues entdecken kann, falls er genügend Spürsinn und Geduld hat." Gerthsen Physik (Springer, 1997), edn. 19, p. 129*

Many interesting phenomena in nature such as first order phase transformations, biomolecular isomerizations, or transport processes in solids, are characterized by widely disparate timescales. While the waiting time for a spontaneous incidence of such a process can exceed seconds or even hours, the underlying relevant molecular motions occur on the femtosecond timescale. This fact poses a serious problem to molecular dynamics computer simulation aimed at revealing the atomistic mechanisms of such phenomena. Here, we give a review of transition path sampling, a set of computational methods designed to overcome the timescale problem.

2.1 Introduction

In the past few decades, molecular dynamics simulation has grown into a very powerful tool that today is used routinely to study the dynamics of condensed matter systems consisting of up to a few million particles with atomistic resolution. Many processes occurring in nature and technology such as the folding of a protein or the transport of a dopant through a semiconductor, however, are still beyond the reach of this methodology due to widely disparate time scales that are present in the problem. Consider, for instance, the nucleation of a crystal from the undercooled liquid. For moderate undercooling, this process typically proceeds through the formation of a critical nucleus that then grows, eventually transforming the whole sample

into the crystalline state. Since this process involves the creation of an interface between the crystallite and the metastable liquid, which is associated with a free energetic barrier, the formation of the critical nucleus is rare on the time scale of basic molecular motions. Indeed, it has been known for a long time that water, carefully cooled below the freezing point, can remain in this supercooled state for hours or even days. Thus, the time scale for nucleation exceeds the picosecond time scale of the formation and cleavage of hydrogen bonds by many orders of magnitude. Similar rare but important events, related to high energy barriers or entropic bottlenecks in phase space, can dominate the dynamics of folding proteins, chemical reactions or transport on surfaces.

Naturally, such a wide separation of time scales is a problem for molecular dynamics simulation. In this method, the equations of motion of the system are solved numerically in small time steps. The size of the time step must be selected such that even the fastest motions in the system are reproduced faithfully. In a molecular system, fast bond and angular vibrations require a time step of about 1 femtosecond. With such a time step, current computer technology permits to follow the time evolution of the system for 10^6–10^9 time steps, corresponding to total simulation times from nanoseconds to microseconds. Of course, the accessible simulation times depend on the size of the system and on the particular way of calculating the forces acting on the individual atoms. If forces are determined *ab initio* from a solution of the electronic structure problem, typical simulation times do not exceed tens of picoseconds even for moderate system sizes of 100–200 atoms. For the case of crystal formation from the supercooled liquid, this limitation in the accessible time scales means that in a molecular dynamics simulation the crystallization event simply cannot be observed.

For the computer simulator this situation is frustrating, particularly because, typically, rare events are not *slow*. Rather, if they occur, they occur rapidly. For instance, the formation of a critical crystalline nucleus proceeds rapidly, while the time spent waiting for this event may be very long. (In fact, microscopic time reversibility requires that the formation of a critical nucleus happens as quickly as its decay.) Similarly, an activated chemical reaction can proceed quickly once it is initiated, but the waiting time between subsequent reactions may be very long. To circumvent this problem caused by widely disparate time scales, several computer simulation algorithms have been devised in recent years. If the reaction mechanism is known in terms of a reaction coordinate that quantifies the reaction progress, for instance the size of the crystalline nucleus forming in the supercooled liquid, the rare event can be studied with umbrella sampling [19] or the blue moon sampling technique [20]. In these methods, an appropriate bias or constraint forces the system to visit the configurations associated with rare barrier crossing events. The detailed mechanism and rate constants of the transition can then be studied using the Bennett-Chandler approach [21, 22], in which dy-

namical trajectories are initiated from these rare configurations, expanding on the original idea of transition state theory [23, 24]. In complex molecular systems, however, *a priori* knowledge of the reaction mechanism is often not available and the methods mentioned above are not directly applicable. In such cases, methods that modify the dynamics of the system, such as metadynamics [25], temperature accelerated dynamics [26], or coarse molecular dynamics [27] can be used to explore possible mechanisms for transitions between stable states.

If both the initial and the final state are known, the transition path sampling method, an importance sampling scheme acting in trajectory space, can be used to study the transition. In contrast to other methods, in a transition path sampling simulation truly dynamical trajectories are considered whereupon both the mechanism as well as the kinetics of the transition can be determined. Alternative methods to study rare transitions between known (meta)stable states are the nudged elastic band method [28] and the string method [29,30]. For a recent review of these approaches and a discussion of their relation to the transition path sampling methodology we refer the reader to Ref. [31]. In this chapter, we will concentrate on the transition path sampling method and its application.

2.2 Transition path sampling

Transition path sampling is a computational methodology developed to study rare transitions between long-lived metastable states [17, 18]. These stable states, let us call them A and B, can be different phases of a condensed material in the case of phase transitions or different chemical species in the case of a chemical reaction. Transitions between stable states A and B are rare (otherwise we could study them with standard molecular dynamics simulation) and may involve crossings of possibly rough and unspecified free energy barriers. While transition path sampling does not require any prior knowledge of the transition mechanism, the stable states A and B between which the transition occurs must be known in advance. The central idea of transition path sampling now is to consider only short trajectories, long enough for the barrier crossing event to complete, but much shorter than the typical waiting between transitions. These short trajectories have different probabilities to be observed: trajectories fluctuating in the stable states, for instance, are more probable than reactive trajectories that cross the barrier. The statistical distribution of various trajectories is taken into account in the definition of the transition path ensemble, which assigns the appropriate probability weight to each individual trajectory. Since in transition path sampling one is interested only in transition pathways, i.e., trajectories that connect the stable states, the transition path ensemble excludes trajectories that do not start in A and end in B. The transition path ensemble is

then sampled with a Monte Carlo procedure that generates trajectories according to their statistical weight. If the sampling is ergodic, all important pathways will be found and can then be analyzed to yield information on the mechanism but also on the kinetics. In the following sections we will outline the basic principles and algorithms of transition path sampling. For further information on various aspects of transition path sampling we refer the reader to Refs. [18, 32–37].

2.2.1 The transition path ensemble

The conceptual starting point of transition path sampling is the definition of the transition path ensemble, a statistical description of all pathways connecting stable states A and B as illustrated in Fig. 2.1. Each of these trajectories has the same temporal length t and consists of an ordered sequence of $L = t/\Delta t$ microscopic states separated by a small time step Δt,

$$x(t) \equiv \{x_0, x_{\Delta t}, x_{2\Delta t}, \ldots, x_t\}. \tag{2.1}$$

Such a sequence of states may, for instance, result from a molecular dynamics or Brownian dynamics simulation carried out with a time step Δt. Each microscopic state x_τ, or time slice, along a trajectory is a complete copy of the system and, depending on the dynamics considered, includes the positions and possibly also the momenta of all particles. Subsequent time slices on a trajectory are connected by the dynamics of the system. If we denote the short time transition probability from state x_τ at time τ to state $x_{\tau+\Delta t}$ one time step later by $p(x_\tau \to x_{\tau+\Delta t})$, the probability density to observe a given trajectory is given by

$$P[x(t)] = \rho(x_0) \prod_{i=0}^{t/\Delta t-1} p(x_{i\Delta t} \to x_{(i+1)\Delta t}). \tag{2.2}$$

Here, $\rho(x_0)$ is the distribution of the initial conditions from which the trajectories start. Equation (2.2) is valid provided the dynamics of the system is Markovian, i.e., the future time evolution of the system depends only on its current state and not on its past. Most of the kinds of dynamics considered in molecular simulations, including Newtonian dynamics, Langevin dynamics and Monte Carlo dynamics, obey this condition.

The probability density of Equ. (2.2) describes the likelihood of observing a trajectory starting end ending at arbitrary microscopic states. In transition path sampling, however, one is specifically interested only in pathways that are reactive, i.e., that start in A and end in B. This condition on the pathways is included into the statistical description of pathways by multiplying the unrestricted probability density $P[x(t)]$ with the characteristic functions of regions A and B acting on the initial and final time slice of the

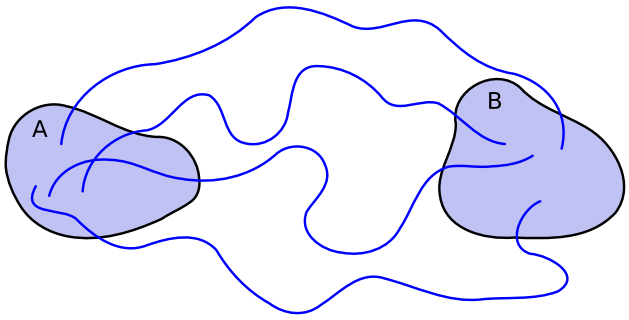

Figure 2.1: The transition path ensemble includes all trajectories of a given length that connect the stables states A and B.

path, respectively:

$$P_{AB}[x(t)] \equiv h_A(x_0)P[x(t)]h_B(x_t)/Z_{AB}(t). \tag{2.3}$$

The characteristic functions $h_A(x)$ and $h_B(x)$are defined such that they are unity if their argument is in the respective region and they vanish otherwise. Thus, $h_A(x)$ is given by

$$h_A(x) = \begin{cases} 1 & \text{if} \quad x \in A, \\ 0 & \text{if} \quad x \notin A, \end{cases} \tag{2.4}$$

and $h_B(x)$ is defined analogously. In Equ. (2.3), $P_{AB}[x(t)]$ is normalized by the factor

$$Z_{AB}(t) \equiv \int \mathcal{D}x(t)\, h_A(x_0)P[x(t)]h_B(x_t) \tag{2.5}$$

where the notation

$$\int \mathcal{D}x(t) \equiv \int \cdots \int dx_0 dx_{\Delta t} dx_{2\Delta t} \cdots dx_t, \tag{2.6}$$

familiar from path integrals, indicates an integration over all time slices of the path. The probability density $P_{AB}[x(t)]$, we call it the *transition path ensemble*, is a statistical description of all pathways of length t that connect the stable states A and B. Pathways that are not reactive are assigned a vanishing weight and thus are not members of the transition path ensemble.

The specific functional form of $P_{AB}[x(t)]$ depends on the distribution of initial conditions, the underlying dynamics and on the definition of the initial and final regions. Depending on the particular situation one considers, the distribution of initial conditions may be the microcanonical or the canonical one. Other distributions are possible as well, including non-equilibrium distributions [38, 39]. The short-time transition probabilities $p(x_\tau \to x_{\tau+\Delta t})$, which enter the expression for the transition path ensemble

in Equ. (2.3), depend on the kind of dynamics chosen to model the time evolution of the system. While for a deterministic time evolution such as Newtonian dynamics the transition probabilities are delta functions leading to a highly singular transition path ensemble [40], the transition probabilities are smooth functions for stochastic dynamics, such as the one produced by the Langevin equation [17]. Finally, care must be exercised in the definition of the stable states A and B. These regions, usually defined in configuration space, should be large enough to include all equilibrium fluctuations of the system in the stable states, but should not overlap with their mutual basins of attraction [18].

While the formal definition of the transition path ensemble poses no difficulty, its practical value hinges on ones ability to generate trajectories according to their weight in this ensemble. An efficient way to accomplish exactly that is discussed in the next section.

2.2.2 Monte Carlo in trajectory space

In a transition path sampling simulation the transition path ensemble is sampled following the basic idea of a Monte Carlo simulation. The according procedure is carried out in two basic steps. First, a new trajectory, $x^{(\mathrm{n})}(t)$, is generated from an old one, $x^{(\mathrm{o})}(t)$, for instance using the shooting algorithm described below and illustrated in Fig. 2.2. Then, the newly generated trajectory is accepted or rejected according to the relative statistical weights of the new and old trajectories. If the new trajectory is accepted, the old trajectory is replaced by the new one. Otherwise, the old one is kept. Iterating these two basic steps generates a biased random walk in trajectory space, in which trajectories are visited according to their weight in the transition path ensemble.

To ensure that the transition path ensemble is sampled correctly, one requires that detailed balance is obeyed,

$$P_{AB}[x^{(\mathrm{o})}(t)]\pi[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)] = P_{AB}[x^{(\mathrm{n})}(t)]\pi[x^{(\mathrm{n})}(t) \to x^{(\mathrm{o})}(t)]. \quad (2.7)$$

Here, $\pi[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)]$ is the probability to move from the old path $x^{(\mathrm{o})}(t)$ to the new path $x^{(\mathrm{o})}(t)$ in one Monte Carlo step. This conditions requires that the flow in trajectory space from $x^{(\mathrm{o})}(t)$ to $x^{(\mathrm{n})}(t)$ is exactly compensated by a flow of equal magnitude in the backward direction. If the transition probability π satisfies the detailed balance condition of Equ. (2.7), the Monte Carlo algorithm conserves the transition path distribution $P_{AB}[x(t)]$ and, if ergodic, results in correct sampling of reactive trajectories. For the two-step Monte Carlo procedure described above, the transition probability $\pi[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)]$ is given by the product of the probability $P_{\mathrm{gen}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)]$ to generate the new path from the old one and the probability $P_{\mathrm{acc}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)]$ to accept the newly generated path,

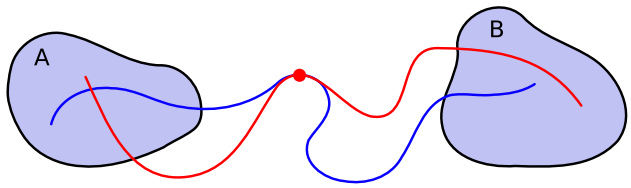

Figure 2.2: In a shooting move, a new trajectory (blue) is generated from an old one (red) by first randomly choosing a time slice of the old path, and then "shooting off" new trajectory segments forward and backward in time, starting from this shooting point. If the underlying dynamics is deterministic, the shooting point must be modified before the shooting takes place.

$$\pi[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)] = P_{\mathrm{gen}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)] P_{\mathrm{acc}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)]. \quad (2.8)$$

Inserting this particular form of the transition probability into the detailed balance condition (2.7) on obtains a condition for the acceptance probability, which can be satisfied using the celebrated Metropolis rule [41], eventually leading to

$$P_{\mathrm{acc}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)] = h_A[x_0^{(\mathrm{n})}] h_B[x_t^{(\mathrm{n})}] \times \min\left\{1, \frac{P[x^{(\mathrm{n})}(t)] P_{\mathrm{gen}}[x^{(\mathrm{n})}(t) \to x^{(\mathrm{o})}(t)]}{P[x^{(\mathrm{o})}(t)] P_{\mathrm{gen}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)]}\right\}. \quad (2.9)$$

According to this equation, which provides a general prescription for accepting or rejecting new pathways, a pathway that does not start in A and ends in B is immediately rejected. Pathways that are reactive, on the other hand, are accepted with a probability that depends both on the relative weight of the old and the new path in the transition path ensemble as well as on the ratio of the forward and backward generation probabilities.

The specific form of the acceptance probability resulting from Equ. (2.9) depends on the particular way new pathways are generated from old ones. The particular algorithm chosen to do that also controls how rapidly path space is sampled and thus determines the efficiency of the transition path sampling simulation. One path generation method that has proven particularly simple, practical, and efficient is the so-called *shooting algorithm* [40]. In this approach, one first randomly selects a time slice $x_{t'}^{(\mathrm{o})}$, the *shooting point*, from the old path. Then, this shooting point is modified, for instance by adding a random perturbation to the momenta. Starting from the modified shooting point, one integrates the equations of motion of the system forward to time t and backward to time 0 obtaining a complete new trajectory $x^{(\mathrm{n})}(t)$. While for stochastic dynamics the modification step may

be omitted, it is strictly necessary for deterministic dynamics. In the latter case, the new trajectory differs from the old one only if the shooting point is modified before integration. In both cases, the acceptance probability for the shooting move is given by

$$P_{\mathrm{acc}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)] = h_A[x_0^{(\mathrm{n})}]h_B[x_t^{(\mathrm{n})}] \min\left[1, \frac{\rho(x_{t'}^{(\mathrm{n})})}{\rho(x_{t'}^{(\mathrm{o})})}\right]. \tag{2.10}$$

Thus, non-reactive trajectories are rejected and reactive ones are accepted with a probability that depends only on the equilibrium distribution at the shooting point before and after the modification. The acceptance probability is particularly simple, if the dynamics conserves the energy and the distributions of initial conditions is the microcanonical one:

$$P_{\mathrm{acc}}[x^{(\mathrm{o})}(t) \to x^{(\mathrm{n})}(t)] = h_A[x_0^{(\mathrm{n})}]h_B[x_t^{(\mathrm{n})}]. \tag{2.11}$$

In this case, non-reactive trajectories are rejected and all reactive ones accepted.

For deterministic dynamics, the modification of the shooting point offers the possibility to tune the acceptance probability and, hence, to optimize the efficiency of the simulation. For very small perturbations of the shooting point, the new trajectory retraces the old trajectory to a large degree. Consequently, the new trajectory has a high probability to be reactive and to be accepted. In this regime, most trajectories are accepted, but since subsequent trajectories are very similar, sampling progress is slow. In contrast, very large perturbations of the shooting point lead to new trajectories that markedly differ from the old ones. Nevertheless, the sampling can be inefficient, if most of the new trajectories are non-reactive and are therefore rejected. Optimum sampling efficiency is obtained for shooting point perturbations with a magnitude somewhere between these two extreme cases. This optimum regime is often characterized by an average acceptance probability of 20-60% [42].

2.2.3 Analyzing trajectories

As a result of a transition path sampling simulation, one obtains a collection of reactive pathways that are typical representatives of the transition path ensemble. Extracting information on the transition mechanism from these pathways is, however, often non-trivial. In this section we will review several computational tools that can be used for this purpose.

A recurrent problem in molecular simulation is to identify those degrees of freedom that capture the essential physics of the process under study and to separate them from the unessential ones that merely act as random noise. For the freezing transition, for instance, it is often unclear whether the size of the crystalline nucleus is sufficient to describe the progress of the

transition or if its shape also plays an important role. While for processes occurring in low-dimensional systems with a handful of degrees of freedom, such as a chemical reaction in the gas phase, locating the saddle points on the potential energy surface often yields valuable mechanistic information, the situation is much more involved in high-dimensional systems. Consider, for instance, a chemical reaction in solution. In this case, solvent degrees of freedom may play an important role that is not easily determined with a saddle point analysis. One difficulty is that the number of saddle points grows exponentially with the number of degrees of freedom such that a complete enumeration of the saddle points becomes impractical beyond a certain system size. Perhaps more importantly, the transition of interest is typically not associated with single saddle points that the system must cross on its way from the initial to the final state. In our crystallization example, the critical nucleus does not necessarily coincide with any saddle point in the potential energy surface.

By watching the atomic motions during the transitions with a molecular viewing program on a computer, one may gain some information about the process of interest. While it is often useful and stimulating to do so, important details of the mechanism, which can be best captured in form of a reaction coordinate, may remain hidden to the eye. A reaction coordinate $q(r)$ is a function of the configuration r of the system, which quantifies the progress of the reaction. In the case of chemical reactions, for instance, bond angles or bond lengths may serve as a reaction coordinate; for a folding protein, the number of native contacts may provide a measure for the folding progress. A good reaction coordinate should tell us how far the reaction has proceeded and what is likely to happen next. The concept of the quality of a reaction coordinate can be made more precise by considering the so called *committor* [43,44], introduced by Onsager as splitting probability [45] and known as p_{fold} in the context of protein folding [46]. The committor $p_B(r)$, which can be defined in configuration space or in phase space [30], is the probability that a trajectory starting from r reaches B rather than A first. As indicated in Fig. 2.3, the committor can be calculated for a particular configuration r by initiating a number of short trajectories from that configuration and determining the fraction of trajectories that end up in B rather than A. A committor value close to 1 indicates that trajectories started from r are very likely to relax into B. While they do not necessarily lie in B itself, such configurations are strongly committed to B and can be viewed to be part of the basin of attraction of region B. Committor values close to 0, on the other hand, characterize configurations that will most likely relax into region A.

Configurations with $p_B(r) = p_A(r) = 1/2$ play a special role because they can be identified as transition states, from which both stable states are accessible with equal probability [47–51]. This statistical concept of a transition state, which generalizes the conventional definition of a transition

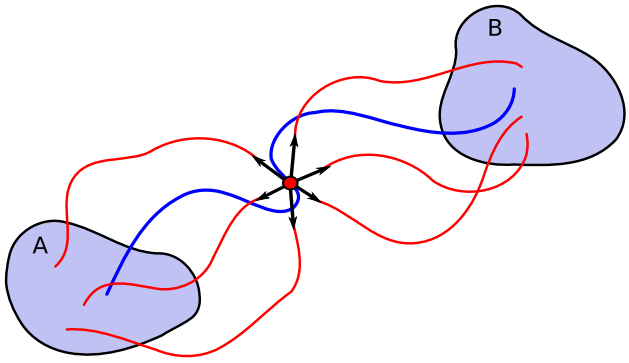

Figure 2.3: To estimate the committor p_B for a particular configuration r (red point), one starts n trajectories from r with random initial momenta and determines the number n_B of trajectories that reach B rather then A. The committor is then given by the fraction $p_B = n_B/n$.

state as saddle point on the potential energy surface, is applicable also to the complex high-dimensional systems of interest here. By determining all configurations with $p_B = 1/2$ on transition pathways one obtains the so-called *transition state ensemble*. Comparison of configurations belonging to the transition states ensemble with those from the stable states can yield important information on the transition mechanism.

As asserted above, a good reaction coordinate should provide a measure for the progress of a particular reaction. In this sense, the committor is the perfect reaction coordinate as it exactly specifies how far the reaction has proceeded and what is likely to happen next [52]. Unfortunately, the committor is very unspecific and does not automatically lead to insight into the mechanism in terms of physical variables. Furthermore, the committor is not easy to evaluate numerically such that it is impractical to use the committor, for instance, as reaction coordinate in a transition state theory calculation. However, the committor can be used as a criterion for distinguishing between good and poor reaction coordinates. If $q(r)$ is a good reaction coordinate, its value determines the progress of the reaction and the committor is completely determined by $q(r)$:

$$p_B(r) = p_B[q(r)]. \tag{2.12}$$

For a poor reaction coordinate, on the other hand, the value of the reaction coordinate does not predict the committor and the above relation does not hold.

The fact that a good reaction coordinate determines the committor can be used to test a proposed reaction coordinate $q(r)$. One may, for instance, expect that for the freezing transition of a particular material the size of the

crystalline nucleus is a good reaction coordinate. To examine the quality of this coordinate, one first computes the free energy F as a function of q. If $q(r)$ has any relation to the rare event of interest, $F(q)$ is expected to be bimodal with a barrier at $q = q^*$ separating the free energy minima corresponding to the stable states A and B. Since the value of a good reaction coordinate completely specifies the committor, all configurations with the same reaction coordinate should also have the same committor. So if one generates configurations for a fixed value of $q(r)$, for instance with a constrained molecular dynamics simulation, and computes the committor value for each of this configurations, the resulting distribution of committor values, $P(p_B)$, should be delta peaked at $p_B(q)$. In particular, configurations with a value of the reaction coordinate of $q(r) = q^*$ corresponding to the barrier top should all have the committor $p_B(q^*)$ and hence the corresponding committor distribution has a sharp peak at $p_B(q^*)$. For a good reaction coordinate, the barrier top coincides with the transition state ensemble and the peak is located at $p_B = 1/2$. In contrast, a poor reaction coordinate does not determine the value of the committor and hence does not lead to sharply peaked committor distributions. Rather, the committor distribution calculated for configurations constrained to the barrier top typically is bimodal with peaks at 0 and 1. Thus, any committor distribution without a single sharp peak is an indication of an inadequate reaction coordinate. An analysis based on committor distributions has been used to reveal the reaction coordinate of ionic dissociation [43], biomolecular isomerization [44], and the freezing transition [53].

2.2.4 Calculating rate constants

Reaction rate constants, describing the kinetics of processes involving rare events, are often measured empirically and thus provide an important way to establish close contact between molecular simulation and experiment. Since pathways harvested in a transition path sampling simulation are truly dynamical trajectories, they can be used to compute such reaction rate constants. The transition path ensemble, however, is restricted only to short trajectory segments during which the transition of interest occurs such that reaction rates cannot be directly extracted from these pathways. While the relative probabilities of different reactive trajectories are correctly described by the transition path ensemble, the information on the probability of observing a reactive event at all (as opposed to no event) is not contained in this path ensemble. Therefore, the fundamental problem in calculating reaction rate constants with transition path sampling consists in estimating the relative weight of the reactive trajectories with respect to all possible trajectories. Several approaches to do so have been suggested in the past and we will briefly survey them in this section. A more detailed review of these methods is provided in Ref. [31].

The link between the microscopic dynamics of the system and its phenomenological description in terms of reaction rate constants is provided by the time correlation function

$$C(t) \equiv \frac{\langle h_A(x_0) h_B(x_t) \rangle}{\langle h_A \rangle}. \tag{2.13}$$

Here, the angular brackets $\langle \cdots \rangle$ denote an equilibrium average. This time correlation function equals the conditional probability to observe the system in region B at time t provided it was in A at time 0. In the case of two-state kinetics, $C(t)$ approaches its asymptotic value exponentially,

$$C(t) = \langle h_B \rangle (1 - \exp(-t/\tau_{\text{rxn}})). \tag{2.14}$$

where the relaxation time τ_{rxn} is related to the forward and backward reaction rate constants k_{AB} and k_{BA} by

$$\tau_{\text{rxn}}^{-1} = k_{AB} + k_{BA}. \tag{2.15}$$

The exponential behavior of Equ. (2.14) cannot be valid for very short times. If regions A and B are not adjacent and there is a gap between them, the system will need a minimum time τ_{mol} to cross this gap. Only for times larger than τ_{mol} can exponential behavior set in as expected from the solution of the phenomenological rate equations. If there is a separation of time scales, i.e., if there is a time regime such that $\tau_{\text{mol}} < t \ll \tau_{\text{rxn}}$, the exponential growth can be approximated by a linear behavior such that

$$C(t) \approx k_{AB} t. \tag{2.16}$$

Equivalently, the time derivative $k(t) = \mathrm{d}C(t)/\mathrm{d}t$ reaches a plateau for times $\tau_{\text{mol}} < t \ll \tau_{\text{rxn}}$ [22]. Thus, knowledge of the time correlation function $C(t)$ is sufficient for a calculation of the forward reaction rate constant k_{AB}.

One transition path sampling approach for the computation of reaction rate constants consists in determining the time correlation function $C(t)$ using free energy calculation techniques [36,42]. In this method, one rewrites $C(t)$ as

$$C(t) = \frac{\int \mathcal{D}x(t) P[x(t)] h_A(x_0) h_B(x_t)}{\int \mathcal{D}x(t) P[x(t)] h_A(x_0)}, \tag{2.17}$$

and observes that $C(t)$ can be viewed as a ratio of two partition functions. In Equ. (2.17), the numerator is the partition function of all pathways starting in A at time 0 and ending in B at time t. The denominator, on the other hand, is the partition function of all pathways starting in A without and restriction on where they end. Hence, the ratio of these partition function is related to the "reversible work" $W_{AB}(t)$ required to transform the ensemble of trajectories with free endpoints into that with endpoints in B,

$C(t) = \exp[-W_{AB}(t)]$. The reversible work $W_{AB}(t)$, a free energy in trajectory space, can be calculated with standard free energy methods such as umbrella sampling [42], thermodynamic integration [54], or even Jarzynski fast switching [55]. In these calculations, one starts with a final region that encompasses the entire configuration space and then successively shrinks it to the desired size. Since the calculation of the reaction rate constant requires the calculation of the time derivative of $C(t)$, in principle several of these path free energy calculations have to be carried out for different path lengths t. This costly operation can be avoided, by calculating $C(t)$ in two steps. First, the time correlation function is calculated with a free energy procedure for one particular time t'. In the second step, the path free energy required to change the path length from t' to t is calculated. This can be done for all values of t up to a maximum time $t_{\max}$ in one single regular transition path sampling simulation [42]. Combining the results of these two calculations one obtains the correlation function $C(t)$ from 0 to $t_{\max}$ and the reaction rate constant can then be extracted from it.

An alternative transition path sampling algorithm for the calculation of reaction rate constants was proposed by Bolhuis and collaborators and named transition interface sampling (TIS) [56, 57]. In this method, pathways of variable length are used which leads to a reduced numerical effort with respect to the method described above. Transition interface sampling, however, is based on an additional assumption about correlated transitions between the stable states. The method rests on the concept of the "overall states" $\mathcal{A}$ and $\mathcal{B}$. Overall state $\mathcal{A}$ consists of points in A plus all points that originate from A in the sense that a trajectory going through such points reaches A before B if followed backwards in time. (This definition is valid only for deterministic trajectories.) Overall state $\mathcal{B}$ is defined analogously. The two overall states $\mathcal{A}$ and $\mathcal{B}$ cover the entire phase space with a possibly very complicated boundary separating them. If one now considers the time correlation function

$$\mathcal{C}(t) \equiv \frac{\langle h_{\mathcal{A}}(x_0) h_{\mathcal{B}}(x_t)\rangle}{\langle h_{\mathcal{A}}\rangle} \tag{2.18}$$

and evaluates the corresponding time derivative in the transition state theory approximation (recrossings are excluded [32]), one obtains the expression

$$k_{AB} = \frac{\langle \phi_{AB}\rangle}{\langle h_{\mathcal{A}}\rangle} \tag{2.19}$$

for the rate constant. Here, $\langle \phi_{AB}\rangle$ is the effective positive flux into region B, i.e., the average flux into B due to trajectories coming directly from A. Thus, for a trajectory connecting A with B only the first entry of the trajectory into B contributes to the effective positive flux. Since the above expression of the rate constant was obtained from an approximation based on transition state theory for the overall states $\mathcal{A}$ and $\mathcal{B}$, the underlying

approximation is that there are no correlated transitions from A to B and back, a condition that is not always satisfied. For stables states A and B defined in configuration space as is customary, chemical reactions occurring in the energy diffusion regime, for instance, may violate this assumption.

In principle, the effective positive flux $\langle \phi_{AB} \rangle$ could be calculated from a long molecular dynamics trajectory by counting the number of first entries into B occurring per time unit. Of course, rare events make this direct approach impractical. To calculate the effective positive flux, Bolhuis and collaborator have therefore developed a technique based on a sequence of non-intersecting interfaces that span the region between A and B [56, 57]. The spacing between these interfaces is selected such that a trajectory crossing interface i coming from A has a non-vanishing probability of also crossing the interface $i + 1$. The effective positive flux can then be expressed as the product of the average positive flux through the surface of A with the product of all these crossing probabilities. The effective positive flux is thus given by the average flux out of A multiplied with the probability of these exit trajectories to eventually reach B. This probability can be calculated from transition path sampling simulations carried out separately for each interface. The ensemble sampled in these simulations consists of trajectories with varying length starting in A, reaching interface i, and then going back to A or on to cross interface $i + 1$. To date, transition interface sampling has been used to calculate reaction rate constants for the freezing transition in simple liquids [53] and several biomolecular isomerizations [58, 59].

For very long and diffusive barrier crossing processes the efficiency of transition interface sampling simulations can be considerably increased by exploiting the loss of correlations along individual pathways. This idea is used in the partial path transition interface sampling method [60]. Another method similar in spirit to the transition interface sampling algorithm is the so-called forward flux method, which can be applied also to non-equilibrium systems in which the stationary phase space distribution is unknown [61–63].

2.3 Concluding remarks

Conducting a computer simulation under near-experimental conditions (for example, using transition path sampling) is a prerequisite but not a guarantee for observing transformation pathways relevant in the real system. For structural transformations in nanocrystals, for instance, the role of different surface passivation agents, defects in the crystal lattice, or the dynamics and composition of the pressure bath is difficult to assess experimentally and even harder to model realistically in a simulation. The relevance of mechanisms observed in a computer simulation can therefore only be established by comparison with experimental data. Transition path sampling offers various techniques to do so. First, calculation of the rate constant allows direct

contact with experiments. Though computationally rather expensive, the calculation of rate constants in the framework of transition path sampling (as discussed in Section 2.2.4) does not need *a priori* knowledge of the reaction coordinate and thus is free of any bias. One drawback that makes a direct comparison with experiments via the rate constant difficult is the fact that rate constants often strongly depend on the quality of empirical potentials. With the development of ever more sophisticated pair potentials and the increasing range of problems that can be tackled *ab initio*, we nevertheless expect this approach to become fruitful for many systems of interest. A second way to establish contact with empirical data is offered by committor analysis and the determination of the transition state ensemble (as discussed in Section 2.2.3). As the central part of a rare event, the transition state not only offers direct insight into the relevant degrees of freedom governing the transformation, but can also be quantified in terms of activation energy and activation volume, quantities that are readily accessible in experiments.

In summary, transition path sampling is a versatile and efficient set of computational techniques for the study of rare events in complex systems. It has has been successfully applied to a broad range of problems from material science to molecular biology; it can be adapted to clarify the transformation details of a simple chemical reactions as well as solid-solid phase transitions, occurring in bulk and nanoscale materials. As computers advance, the range and complexity of systems to which transition path sampling can be fruitfully applied increases and a more realistic modeling of the experimental situation becomes possible in many cases.

Chapter 3

Precision shooting: Sampling long transition pathways

The kinetics of collective rearrangements in solution, such as protein folding and nanocrystal phase transitions, often involve free energy barriers that are both long and rough. Applying methods of transition path sampling to harvest simulated trajectories that exemplify such processes is typically made difficult by a very low acceptance rate for newly generated trajectories. We address this problem by introducing a new generation algorithm based on the linear short-time behavior of small disturbances in phase space. Using this "precision shooting" technique, arbitrarily small disturbances can be propagated in time, and any desired acceptance ratio of shooting moves can be obtained. We demonstrate the method for a simple but computationally problematic isomerization process in a dense liquid of soft spheres. We also discuss its applicability to barrier crossing events involving metastable intermediate states.

3.1 Introduction

Transition path sampling (TPS) is a versatile and efficient set of computational techniques for the study of rare events [17, 18, 31, 32]. It has been successfully used to reveal the microscopic mechanisms of processes as diverse as autoionization in liquid water [64], structural transformations in nanocrystalline solids [65], and folding of small proteins [66]. The purpose of this chapter is to propose a new shooting algorithm which can greatly increase the efficiency of TPS when transit times of activated trajectories greatly exceed the picosecond time scale of phase space stability.

At its core TPS is a Monte Carlo procedure enabling a random walk in the ensemble of pathways that cross a free energy barrier between two metastable states (denoted A and B). While this sampling is strongly biased towards reactive trajectories, it leaves the underlying dynamics of the sys-

tem unchanged. Thus, the result of a TPS simulation is a representative set of true dynamical pathways, weighted as if they were excerpted from an extremely long, unbiased simulation of equilibrium dynamics. Many analytical tools have been developed to extract from such a collection of trajectories useful molecular information about the process of interest [31].

The algorithm typically used to construct such a random walk is called *shooting* [18]. Here, a point along a given reactive trajectory is randomly selected and slightly changed; for instance, one might change the velocities of all particles by a small random number drawn from a symmetric distribution. Using the dynamical rules of the system, this *shooting point* is then propagated forward and backward in time to obtain a complete new trajectory. If this new trajectory still connects A with B it is accepted and used as a basis for the next shooting move; otherwise it is rejected.

The efficiency of this algorithm in exploring the transition path ensemble is based on a balance between the intrinsic instability of complex dynamical systems and the local character of the shooting move: Small disturbances grow exponentially quickly in time, leading to separation of trajectories typically within a few picoseconds. Nonetheless, if the disturbance is small, the new trajectory will be locally similar to the old one and is therefore likely to surmount the barrier between A and B; such shooting moves will be accepted frequently. Just as with conventional Monte Carlo moves in configuration space, maximum efficiency can often be obtained by adjusting the size of the disturbance to achieve an acceptance probability of roughly 40% [18].

Shooting moves are best suited for the study of systems that relax quickly (within the picosecond time scale of trajectory separation) into their product state after reaching the top of the barrier. Many interesting processes, like the nucleation of first order phase transitions or conformational change in complex molecules, proceed much more slowly from the transition state. In TPS simulations of such systems, shooting moves must be made extraordinarily subtle in order to stand a reasonable chance of connecting reactant and product states. As a matter of practice, however, disturbances cannot be made arbitrarily small due to the limited machine precision of floating point numbers. Lacking an ability to control the degree of global separation between trajectories, TPS methods are severely compromised in efficacy. The demonstrated computational advantages of importance sampling in trajectory space lose appeal when offset by the wasted effort of generating a vast excess of non-reactive paths.

In a recent paper [67], Bolhuis addressed this problem by modifying slightly the rules that propagate a system in time. Specifically, a weak stochastic component was added to the dynamics, removing the unique correspondence between a trajectory's past and its future. It thus became possible to resample only parts of an existing pathway, leading to much higher acceptance probabilities for shooting moves and, Bolhuis reports, significant improvement in sampling efficiency [67].

In this chapter, we show that it is possible to perform productive shooting moves for arbitrarily long transition paths without modifying a system's natural dynamics. Our technique for introducing and propagating extraordinarily small disturbances is based on the simple dynamics of small perturbations in phase space. We explain the method and offer a straightforward algorithm for implementation in Sec. 3.2. Use of the technique is demonstrated in Sec. 3.3 for a simple isomerization process in a dense liquid that, by construction, involves diffusive dynamics on a rugged barrier. In Sec. 3.4 we examine limitations of the method by considering reactive dynamics that pass through highly metastable, obligatory intermediate states.

3.2 Linearized dynamics of small perturbations

3.2.1 Exponential divergence of trajectories

In a TPS simulation of a system evolving with deterministic dynamics, a trajectory X of length τ consists of a number of "snapshots" $x_{i\Delta t}$, which are separated by a time step Δt,

$$X = \{x_0, x_{\Delta t}, x_{2\Delta t}, \ldots, x_\tau\} \; . \tag{3.1}$$

Here, the time slices $x_{i\Delta t}$ are full phase space vectors, detailing the positions and velocities of all particles. Subsequent time slices are related by

$$x_{(i+1)\Delta t} = \Phi(x_{i\Delta t}) \; , \tag{3.2}$$

where the function Φ propagates the system for one time step.

Consider now a shooting move, in which a small disturbance δx_0 is added to the shooting point x_0 to obtain state $y_0 = x_0 + \delta x_0$ of the shooting trajectory Y. (To simplify notation, we will assume the shooting point to be x_0, the initial state of the trajectory, throughout this section. The algorithm we will describe applies transparently to shooting points at any chosen time along the trajectory.) Usually the perturbation δx_0 affects only momentum space, but changing the positions of the particles can be useful in some cases [18]. The perturbed point y_0 is then propagated for a number of n time steps to obtain $y_t = \Phi_t(y_0)$, where $t = n\Delta t$ and Φ_t refers to the n-fold application of the time step propagator. We define the time-evolved disturbance δx_t by subtracting the old trajectory from the new one, $\delta x_t = y_t - x_t$. Due to the dynamic instability of the system, perturbations grow exponentially in time,

$$|\delta x_t| \approx |\delta x_0| \, \mathrm{e}^{\lambda t} \; . \tag{3.3}$$

Here, λ is the largest Lyapunov exponent of the system [68]. For typical fluid systems, $1/\lambda \approx 1\,\mathrm{ps}$.

We wish to control precisely the time it takes for a small perturbation to reach a size of order 1, at which point the new trajectory will be essentially separated from the old one. This time determines the probability that the new trajectory will be reactive and therefore acceptable. Because of subsequent exponential growth, $|\delta x_0|$ must be decreased by many orders of magnitude to increase the separation time of trajectories by even a few picoseconds (see Figs. 3.1 and 3.2). With the standard double precision format for representing floating point numbers on a computer, however, the smallest number that can be added to 1.0 to give a result distinguishable from 1.0 is of the order of 10^{-15}, and numerical results become unreliable at values of $|\delta x_0|$ well above this limit. (We assume throughout this chapter that a system of units has been chosen such that typical numerical values of coordinates and momenta are of order 1.) Especially when the total length of the transition path is significantly longer than a few picoseconds, the limited range of practical displacement sizes constitutes a severe sampling problem: Shooting moves will only be accepted from points in the vicinity of the barrier top; otherwise, new trajectories will simply return to the stable state they came from and be rejected. As the system may stay near the *a priori* unknown barrier top only for a small fraction of the total transition time, sampling can break down completely. In these cases, implementing shooting displacements of arbitrarily small size would be very helpful.

3.2.2 Dynamics in the linear regime

We propose to solve this problem by using perturbation theory to follow the time evolution of the *displacement* vector δx_0 itself, up to the point where it grows large enough to allow accurate evaluation of the sum $x_t + \delta x_t$. Expanding y_t around x_0, we obtain

$$\delta x_t = y_t - x_t = \Phi_t(x_0 + \delta x_0) - \Phi_t(x_0) = \frac{\partial \Phi_t(x_0)}{\partial x_0}\, \delta x_0 + \mathcal{O}\left[(\delta x_0)^2\right] \ . \quad (3.4)$$

For small displacements $|\delta x_0| < 10^{-15}$, the linear approximation is for all practical purposes exact on the scale of a single time step,

$$\delta x_t = S\, \delta x_0 \ , \quad (3.5)$$

where the matrix S is given by

$$S = \frac{\partial \Phi_t(x_0)}{\partial x_0} \ . \quad (3.6)$$

To integrate δx_0 forward in time according to equation (3.5), the equations of motion for the matrix S could in principle be solved numerically [71]. Doing so in practice would be cumbersome, requiring calculation of all second derivatives of the interaction potential with respect to particle positions. We

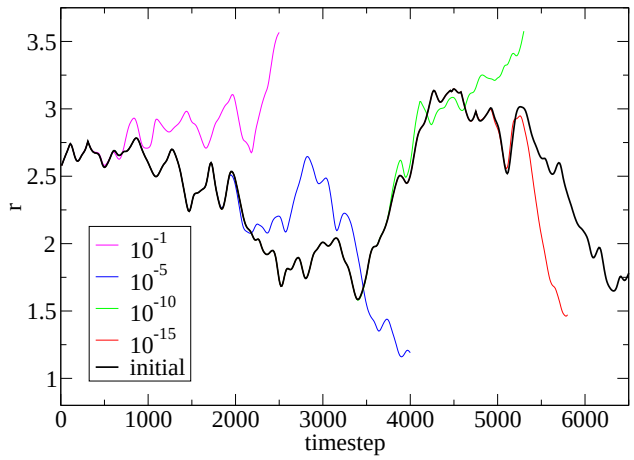

Figure 3.1: Distance r, in reduced units, between two particles in a liquid of soft spheres as a function of time along a reference trajectory and four shooting trajectories. (Here and throughout the chapter, if not explicitly mentioned, the model system consists of 108 WCA [69] particles in their liquid state. The total energy per particle is 1.0 and the density is 0.75, in reduced units. The equations of motion are integrated with the velocity Verlet algorithm [70] with a time step of $\Delta t = 0.002$.) At time zero, displacements of various size are added to particle velocities in the reference trajectory. Because phase space disturbances grow exponentially in time, decreasing the shooting displacement by successive orders of magnitude results in only linear increase in the time that elapses before trajectories separate. Shooting displacements smaller than 10^{-15} can not be resolved in double precision; the resulting shooting trajectories will exactly retrace the initial one.

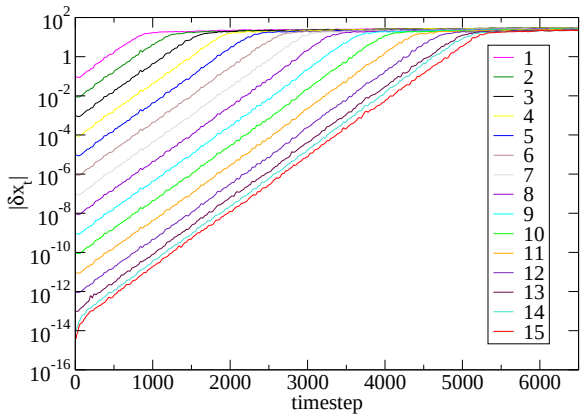

Figure 3.2: Time evolution of small displacements δx_t in a liquid of 108 soft spheres for disturbances of various sizes, $|\delta x_0| = 10^{-\alpha}$ (legend values indicate values of α). All displacements grow exponentially with the same rate, up to the time where trajectories separate. The maximum possible value of $|\delta x_t|$ is determined by the dimensions of the simulation box. Note that adjacent lines are equidistant in the linear regime, except for displacement sizes smaller than 10^{-12}. Although these smallest displacements yield trajectories that can in practice be distinguished from the base trajectory, limited numerical precision introduces rounding errors that degrade computational estimates of linear divergence.

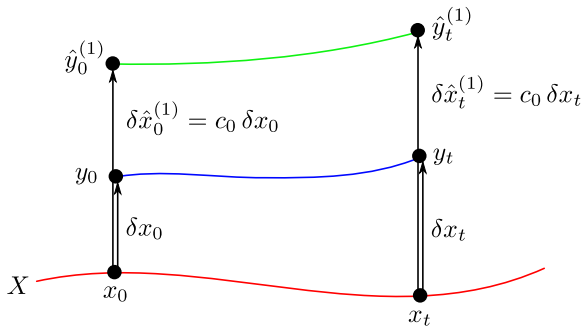

Figure 3.3: Time evolution of shooting moves from a point x_0 on the reference trajectory X. The two shooting displacements δx_0 and $\delta\hat{x}_0^{(1)}$, as vectors in high-dimensional phase space, point in the same direction but have different magnitudes, $\delta\hat{x}_0^{(1)} = c_0\,\delta x_0$. At a later time t short enough that first-order perturbation theory remains valid, these displacements remain proportional, $\delta x_t^{(1)} = c_0\,\delta x_t$. The displacement of interest δx_t, no matter how small, can thus be constructed in the linear regime simply by dividing $\delta\hat{x}_t^{(1)}$ by the original scaling factor c_0.

propose a much simpler approach for advancing δx_t, inspired by methods for computing Lyapunov exponents in systems whose interaction potentials lack well-defined second derivatives. Our implementation is illustrated in Fig. 3.3.

Instead of integrating the small perturbation δx_0, we follow the time evolution of a related perturbation $\delta\hat{x}_0^{(1)}$, which *is* large enough to be added at the shooting point and propagated in the usual way, i.e., by integrating Newton's equation of motion for $\hat{y}_0^{(1)} = x_0 + \delta\hat{x}_0^{(1)}$. We use a superscript for $\hat{y}^{(i)}$ and $\delta\hat{x}^{(i)}$ because in the following we will consider a family of different perturbed trajectories $\hat{Y}^{(i)} = \{\hat{y}_0^{(i)}, \ldots, \hat{y}_\tau^{(i)}\}$, with $\hat{y}_t^{(i)} = x_t + \delta\hat{x}_t^{(i)}$. Exploiting the linearity described by Eq. (3.5), we choose $\delta\hat{x}_0^{(1)}$ to be in the same direction (in the high-dimensional phase space) as δx_0,

$$\delta\hat{x}_0^{(1)} = c_0\, \delta x_0 \ , \tag{3.7}$$

where c_0 is a scalar constant. If $\delta\hat{x}_0^{(1)}$ is also small enough to justify the linear approximation of Eq. (3.5),

$$\delta\hat{x}_t^{(1)} = S\, \delta\hat{x}_0^{(1)} \ , \tag{3.8}$$

then the initial relationship between δx_0 and $\delta\hat{x}_0^{(1)}$ holds also at a later time t,

$$\delta x_t = S\, \delta x_0 = \frac{1}{c_0}\, S\, \delta\hat{x}_0^{(1)} = \frac{1}{c_0}\, \delta\hat{x}_t^{(1)} \ . \tag{3.9}$$

In the linear regime it is thus possible to follow the time evolution of arbitrarily small displacements δx_0 by monitoring larger, proportional displacements.

The linear approximation for the "helper" displacement $\delta\hat{x}_t^{(1)}$ in Eq. (3.8) will of course remain valid for only a short time $t_{\text{lin}}^{(1)}$, typically less than 1 ps. Our interest in the trajectory $\hat{Y}^{(1)}$, however, is only as a proxy for the evolution of smaller displacements that cannot be represented explicitly. As $\delta\hat{x}_t^{(1)}$ approaches the boundary of the linear regime, $t \lesssim t_{\text{lin}}^{(1)}$, we may therefore switch our attention to a different helper trajectory $\hat{Y}^{(2)}$, one whose displacement is initially too small to be of practical use but by the time $t_{\text{lin}}^{(1)}$ grows large enough to be represented explicitly. The new displacement $\delta\hat{x}_t^{(2)} = \hat{y}_t^{(2)} - x_t$ can be obtained at any time $t < t_{\text{lin}}^{(1)}$ simply by scaling $\delta\hat{x}_t^{(1)}$ appropriately, $\delta\hat{x}_t^{(2)} = \delta\hat{x}_t^{(1)}/c_1$. Because it is initially smaller than $\delta\hat{x}_t^{(1)}$, it will remain in the linear regime for a longer time, $t_{\text{lin}}^{(2)} > t_{\text{lin}}^{(1)}$. For times $t > t_{\text{lin}}^{(1)}$ we therefore proceed by integrating standard equations of motion for $\hat{Y}^{(2)}$ until it approaches the boundary of the linear approximation. At that point we repeat the procedure, scaling back the displacement to switch attention to yet another helper trajectory.

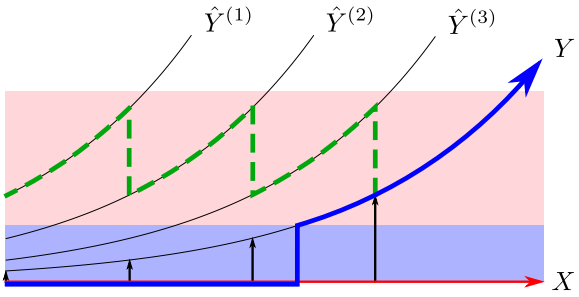

Figure 3.4: Precision shooting algorithm for generating a trial trajectory Y (thick curve) whose initial displacement from the base trajectory X (thin straight line) is extraordinarily small. Points in the dark field (whose extent is $\approx 10^{-15}$) cannot be numerically distinguished from the reference trajectory. Until the trial trajectory exits this region, its time evolution is calculated by proxy using "helper" trajectories $\hat{Y}^{(i)}$ (thin curves). Displacements of trial (vertical arrows) and helper trajectories from the base trajectory are related by proportionality as long as they remain within the linear regime, represented by the light field. To preserve this simple relationship, the displacement of the helper trajectory $\hat{Y}^{(i)}$ from the reference trajectory is scaled back when it threatens to leave the region of linear dynamics, effectively switching the system to the next helper trajectory $\hat{Y}^{(i+1)}$. By following the systems dynamics along many sections of helper trajectories (dashed curve) and by keeping track of the rescaling factors, one can accurately construct the state of the trial trajectory once it becomes distinguishable from the reference trajectory. The result of this shooting move is a trial trajectory Y that is numerically identical to the base trajectory X over a certain length of time and then emerges from it in the correct way.

In effect we monitor a single displacement from the reference trajectory whose magnitude is periodically scaled down such that the linear approximation is always valid. In this way we can monitor the time evolution of an arbitrarily small disturbance. The corresponding trajectory will be numerically indistinguishable from the reference trajectory as long as the displacement's magnitude is smaller than $\sim 10^{-15}$. At later times the displaced trajectory can be distinguished, and its dynamics can be safely computed by integrating equations of motion in the usual way.

3.2.3 Algorithm

This insight suggests the following algorithm, which implements a shooting trajectory Y, whose initial deviation δx_0 from the base-trajectory X is smaller than the precision limit. This is done by monitoring "helper" trajectories $\hat{Y}^{(j)}$, which are obtained by repeated rescaling. For an illustration

of this algorithm see Fig. 3.4.

1. At the shooting point x_0, add a displacement $\delta\hat{x}_0^{(1)}$ of fixed size $|\delta\hat{x}_0^{(1)}| = \sigma$. The displacement $\delta\hat{x}_0^{(1)}$ is parallel to δx_0 and larger by a factor of c_0.

2. Propagate the point $\hat{y}_0^{(1)} = x_0 + \delta\hat{x}_0^{(1)}$ forward in time for n time steps, corresponding to a time interval $t = n\Delta t$.

3. Compute the factor $c_1 = |\delta\hat{x}_t^{(1)}|/|\delta\hat{x}_0^{(1)}| = |\hat{y}_t^{(1)} - x_t|/|\delta\hat{x}_0^{(1)}|$ quantifying the divergence from the reference trajectory. Switch to a new helper trajectory by setting $\hat{y}_t^{(2)} = x_t + \delta\hat{x}_t^{(1)}/c_1 = x_t + \delta\hat{x}_t^{(2)}$. Store the factor c_1.

4. Propagate the new displacement forward in time by integrating the equations of motion for n steps beginning from $\hat{y}_t^{(2)}$. Calculate and store the factor $c_2 = |\delta\hat{x}_{2t}^{(2)}|/|\delta\hat{x}_t^{(2)}|$.

5. Iterate step 4, each time beginning from $\hat{y}_{(j-1)t}^{(j)}$ and rescaling by the factor $c_j = |\delta\hat{x}_{jt}^{(j)}|/|\delta\hat{x}_{(j-1)t}^{(j)}|$. At every iteration compute the displacement of interest, $\delta x_{jt} = C_j^{-1}\,\delta\hat{x}_{jt}^{(j)}$, where $C_j = \Pi_{k=0}^{j-1} c_k$ is the product of all factors used for rescaling so far. To store the current point along the actual shooting trajectory, compute $y_{jt} = x_{jt} + \delta x_{jt}$. As long as $|\delta x_{jt}| \lesssim 10^{-15}$, y_{jt} will be numerically identical to x_{jt}.

6. If $|\delta x_{jt}| > \sigma$, the actual shooting displacement is large enough to be treated in the usual way: Set $y_{jt} = x_{jt} + \delta x_{jt}$ and integrate equations of motion from this point without further rescalings (ceasing iteration of step 4).

Note that for shooting moves conducted at points other than x_0, the procedure must be repeated backward in time to obtain a complete shooting trajectory. In the following we discuss the accuracy of this scheme and give recommendations for choosing values of σ and n.

3.2.4 Validity of the linear approximation

The above algorithm is exact only if the linear approximation of Eq. (3.5) holds, and if calculations are carried out with infinite numerical precision. For perturbations of finite size, and given practical computational limitations, deviations from this approximation occur. The question thus arises, how accurate an approximation is this approach for propagating small disturbances? More specifically, to what extent do helper displacements remain proportional to the actual shooting displacements of interest? One could certainly imagine that the fast growth of small non-linearities rapidly erodes

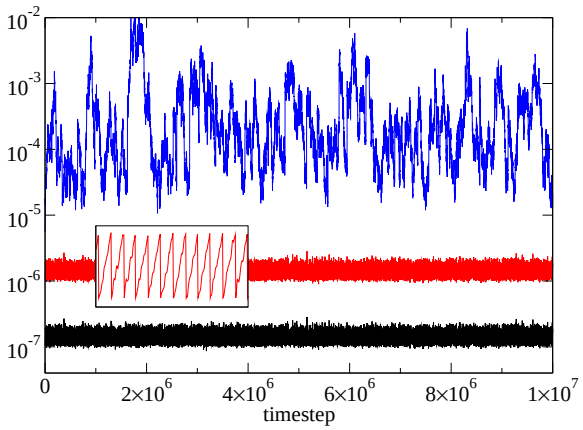

Figure 3.5: Time evolution of the magnitude of two shooting displacements for a fluid of WCA particles. Displacements of size $|\delta\hat{x}_0^{(1)}| = 10^{-7}$ (bottom curve) and $|\delta\hat{x}_0^{(2)}| = 10^{-6}$ (middle curve) are initially proportional (pointing in the same direction in phase space). Both are rescaled to their initial size every 100 time steps (see inset for a magnified view) to preserve this linear relationship. Deviation from proportionality is quantified by the relative error $\epsilon(t)$ (top curve) defined in Eq. (3.10).

the linear relationship on which we depend. Here we present evidence from computer simulations that proportionality of small displacements can hold in practice over very long time scales.

Figure 3.5 shows the time evolution of two proportional disturbances. The initial displacement vectors point in the same direction of phase space but have different magnitudes, $|\delta\hat{x}_0^{(1)}| = 10^{-7}$ and $|\delta\hat{x}_0^{(2)}| = 10^{-6}$. The respective shooting trajectories were propagated independently, and the displacements from the base trajectory were rescaled to their initial length every 100 time steps. For a perfectly linear time evolution, these displacements remain proportional at all later times. In practice the vectors $\delta\hat{x}_t^{(1)}$ and $\delta\hat{x}_t^{(2)}$ will develop a nonzero angle due to non-linearities. To quantify this deviation from parallel alignment, we define

$$\epsilon(t) = \frac{|\delta\hat{x}_t^{(1)} - \delta\hat{x}_t^{(2)}/10|}{|\delta\hat{x}_t^{(1)}|}\,. \tag{3.10}$$

As shown in Fig. 3.5, the relative error $\epsilon(t)$ does not grow above a low level even for very long simulation runs.

This long time stability of aligned disturbances holds over a broad range of displacement sizes between 10^{-10} and 10^{-3}. Values of $\epsilon(t)$ can be somewhat larger than for the specific case plotted in Fig. 3.5 but on average do not grow larger than 10^{-3} for any case. The error is insensitive to the choice

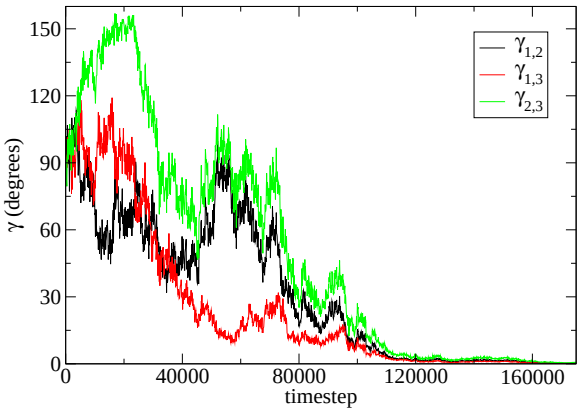

Figure 3.6: Relative orientations of three phase space displacement vectors ($\delta\hat{x}_t^{(1)}$, $\delta\hat{x}_t^{(2)}$, and $\delta\hat{x}_t^{(3)}$) for a WCA fluid. We plot the angle $\gamma_{i,j} = \cos^{-1}(\delta\hat{x}_t^{(i)} \cdot \delta\hat{x}_t^{(j)} / |\delta\hat{x}_t^{(i)}||\delta\hat{x}_t^{(j)}|)$ between each pair of displacements as a function of time. Displacements differ in initial direction and size ($|\delta\hat{x}_0^{(1)}| = 10^{-8}$, $|\delta\hat{x}_0^{(2)}| = 10^{-8}$, $|\delta\hat{x}_0^{(3)}| = 10^{-7}$) and are rescaled to their original sizes at different intervals ($n^{(1)} = 100$, $n^{(2)} = 200$, $n^{(3)} = 500$).

of the rescaling interval $n\Delta t$, as long as the displacements stay smaller than approximately 10^{-2}, where the linear regime breaks down. For displacements smaller than 10^{-10}, rounding errors become problematic, and phase space displacements do not remain parallel to a good approximation. One might expect the region of long time stability to extend to even lower levels, closer to the precision limit of 10^{-15}. However, the total achievable accuracy in a computer simulation depends, among other factors, on the details of the integrator Φ, the size of the time step, and the dimensionality of the system, and can lie well above the precision limit of 10^{-15}. We find that for the particular system used here, the effective precision level is about 10^{-12}. The interplay of non-linear terms in the time evolution and numerical precision is further discussed in Appendix A.

In the light of these observations, a value of $\sigma = 10^{-6}$ as initial magnitude for helper displacements seems appropriate. This choice lies midway between the upper limit of the linear regime (approximately 10^{-2}), and the point where rounding errors become dominant (approximately 10^{-10}). As the accuracy of the rescaling scheme is quite insensitive to the frequency of rescalings, many equally good choices of n are possible. As a starting point, a value of n which leads to rescalings every time the displacements have doubled their size is advisable.

The fact that $\epsilon(t)$ does not show any systematic long-time growth in Fig. 3.5 seems surprising. After all, no constraint is imposed on the direc-

tion of the displacement vectors. Why does an accumulation of errors not eventually lead to decoupling and $\epsilon(t) \approx 1$? Stability of the precision shooting algorithm is in fact a simple and direct consequence of the collective dynamics of displacements in the linear regime. In Figure 3.6 we plot the angles between three periodically rescaled shooting displacement vectors of different size and random initial direction. Eventually, they all rotate into the same direction, which is associated with the largest Lyapunov exponent λ of the system. The time scale on which the directions of different displacement vectors converge is on the order of $1/\Delta\lambda$, where $\Delta\lambda$ is the difference between the first and second largest Lyapunov exponents [72]. It is because of this convergence, that the difference vector $\delta\hat{x}_t^{(1)} - \delta\hat{x}_t^{(2)}/c$ between two proportional displacements with initially identical direction will stay small.

We point out that this property constitutes a main difference of our method over the stochastic scheme introduced by Bolhuis [67] and similar algorithms. Consider, for instance, the following simple algorithm that can be viewed as a smooth version of the stochastic scheme by Bolhuis:

- Choose a shooting point $x_{\mathrm{s}\Delta\mathrm{t}}$.
- A fixed number of timesteps n earlier and later, at the points $x_{(s+n)\Delta t}$ and $x_{(s-n)\Delta t}$, add a displacement of 10^{-15} to one velocity component of one particle.
- Integrate the points $x_{(s+n)\Delta t}$ and $x_{(s-n)\Delta t}$ forward and backward in time, respectively, to get a complete new trajectory.

Just like the precision shooting algorithm, this simple scheme results in a shooting trajectory that is numerically identical to its base trajectory for a certain period of time. However, the emerging separation between base and shooting trajectories will not be consistent with a shooting move conducted at x_0, but rather with two uncorrelated shooting moves at $x_{(s+n)\Delta t}$ and $x_{(s-n)\Delta t}$. Our algorithm, on the other hand, correctly reproduces the correlated forward and backward dynamics of a displacement introduced at $x_{s\Delta t}$.

3.3 A simple test system

We demonstrate the precision shooting algorithm on a simple isomerization process of a solvated diatomic molecule in three dimensions.

Our test system consists of 389 particles interacting *via* the WCA [69] potential. We use conventional reduced units, with particle mass and potential parameters σ and ϵ all set to unity. Particles #1 and #2 do not interact *via* the WCA potential, but are bonded through a one-dimensional

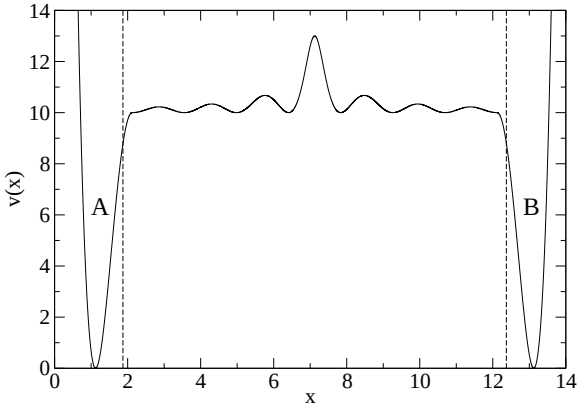

Figure 3.7: Potential energy $v(x)$ of interaction between particles #1 and #2, comprising the diatomic molecule in our model isomerization process, plotted as a function of the difference $x = x_2 - x_1$ between their x-coordinates. The dashed lines mark the boundaries of the minima A and B.

potential with two deep minima separated by a rough barrier (see Fig. 3.7):

$$v(x) = \begin{cases} h_1 \left[1 - q(x)^2/w^2\right]^2 & \text{if } q(x) < 0\,, \\ h_1 \left[1 - (q(x) - b)^2/w^2\right]^2 & \text{if } q(x) > b\,, \\ h_1 + h_2 \frac{\cos^2[a(q(x)-b/2)]}{\sqrt{1+ga^2(q(x)-b/2)^2}} & \text{else}\,. \end{cases} \tag{3.11}$$

Here, $q(x) = x - (r_c + w)$, $x = x_2 - x_1$ is the difference between the x-component of the position of the bonded particles, $r_c = 2^{1/6}$ is the cutoff of the WCA potential, $w = 1$ determines the width of the minima, $b = 10$ and $h_1 = 10$ are the length and height of the barrier in between, respectively, and the constants $h_2 = 3$, $a = 7\pi/b$, and $g = 2$ determine the shape of the barrier. The potential and its derivative are continuous by construction.

To speed computation, we borrow a trick from Bolhuis' work [67]: Particle #2 is considered to lie always to the right of particle #1, hence $x > 0$. This choice, together with the one-dimensionality of $v(x)$, allows us to choose a simulation box with dimensions $14.4 \times 6 \times 6$. The resulting particle density is 0.75, the total energy per particle is 1.0, and the temperature is 0.45, as gauged by average kinetic energy. We use the velocity Verlet algorithm [70] to integrate the equations of motion with a time step of 0.002.

We are interested in sampling the transition of the dimer from the "contracted" minimum A at $x_\mathrm{A} = r_\mathrm{c}$ to the "extended" minimum B at $x_\mathrm{B} = r_\mathrm{c} + b + 2w$. The dimer is defined to be in state A for $x < x_\mathrm{A} + 0.75w$ and in state B for $x > x_\mathrm{B} - 0.75w$ (see Fig. 3.7). Because the system is dense, and the barrier is both long and rough, relaxation from the transition state into either stable minimum is quite protracted.

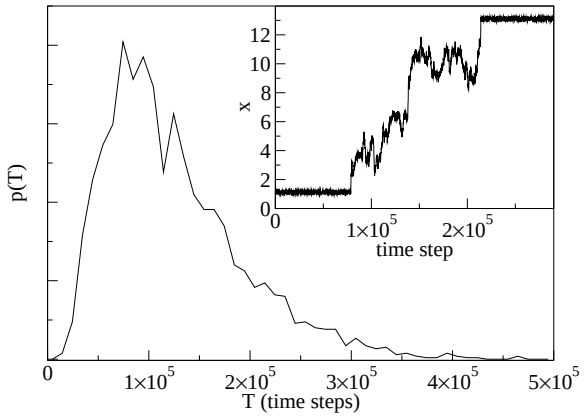

Figure 3.8: Distribution of transition times T for our model isomerization process, as gauged from 2500 trajectories initiated at the barrier top. Inset: Difference $x = x_2 - x_1$ between the x-coordinates of particles #1 and #2 as a function of time for a typical trajectory.

In conducting TPS simulations it is important that sampled trajectories are not shorter than typical spontaneous barrier-crossing events [18]. We determine this typical duration for our simple model system by initiating many straightforward molecular dynamics simulations with the dimer bond length set at $x = r_c + w + b/2$, corresponding to the middle of the barrier. Integrating the equations of motion forward and backward in time yields a representative sample of the transition path ensemble. For a particular trajectory, the transition time T is the time the system spends between regions A and B. The resulting distribution of transition times is plotted in Fig. 3.8. For TPS simulations, we choose a total trajectory length of 3×10^5 time steps, long enough to include 98% of the natural transition path ensemble. The bias of our sampling to short transitions is therefore minor.

Although the artificial potential energy landscape studied here does not directly represent any physical system of interest, it nevertheless shares with many real systems features that lead to long transition pathways and make straightforward application of TPS methods ineffective. In our view roughness of the barrier region is an important ingredient. Models featuring long but flat barriers, such as that of Ref. [67], should not in fact pose any severe problems for path sampling via the standard shooting move. Assuming that motion atop such a flat barrier is diffusive in nature, and that time evolution from the edge of the barrier proceeds into the adjacent minimum with near certainty, then a trajectory initiated on the barrier will relax into stable state A with probability $p_A = 1 - y/b$, where y is the initial distance from A and b is the width of the barrier. Similarly, the probability of relaxing first

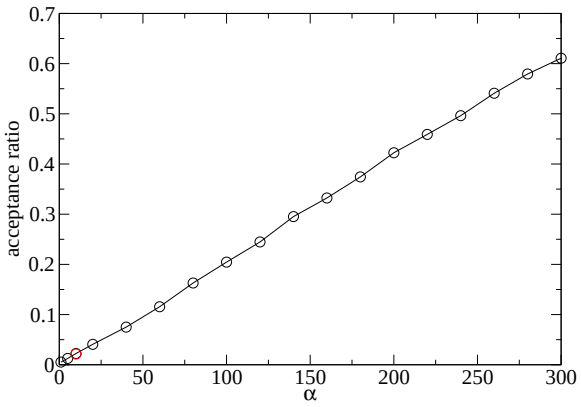

Figure 3.9: Fraction of shooting moves for our model isomerization process that are accepted in long TPS simulations. Acceptance ratios are shown for shooting displacements of various sizes, $\Delta p = 10^{-\alpha}$, implemented using the standard shooting algorithm for values of α of 1, 5, and 10, and the precision shooting algorithm for $\alpha \geq 10$. For $\alpha = 10$, the result obtained from the precision shooting algorithm (red) is effectively indistinguishable from the one obtained with standard shooting.

into state B is $p_{\mathrm{B}} = y/b$. A standard shooting move from the barrier region then yields a reactive trajectory with probability

$$P_{\mathrm{acc}} = \frac{1}{b} \int_0^b \mathrm{d}y \, p_{\mathrm{A}} p_{\mathrm{B}} = \frac{1}{6} \,. \tag{3.12}$$

This value of the acceptance rate should correspond to near optimal sampling of the transition path ensemble [18]. A problematically low acceptance rate would only arise if one were to sample trajectories of insufficient length, i.e., paths shorter than typical spontaneous transitions.

In our TPS simulations, only momenta (and not particle positions) are disturbed in the shooting moves, with each particle's momentum changed in each direction by an amount drawn from a Gaussian distribution of standard deviation Δp (followed by rescaling of all momenta to enforce energy conservation [17]). We conduct standard shooting moves with values of Δp of 10^{-1}, 10^{-5}, and 10^{-10}, as well as precision shooting moves with Δp ranging in size from 10^{-10} to 10^{-300}. The latter are implemented using helper displacements with $\Delta p = 10^{-7}$ and are rescaled every time they reach twice their original size. For the system size studied here, the initial magnitudes of the resulting displacement vectors are larger than the corresponding value of Δp by a factor of roughly 34 on average. For instance, $\Delta p = 10^{-7}$ results in displacement vectors with an initial size of about $|\delta x_0| = 3.4 \times 10^{-6}$. For each set of sampling parameters, we attempt 50,000 Monte Carlo moves in

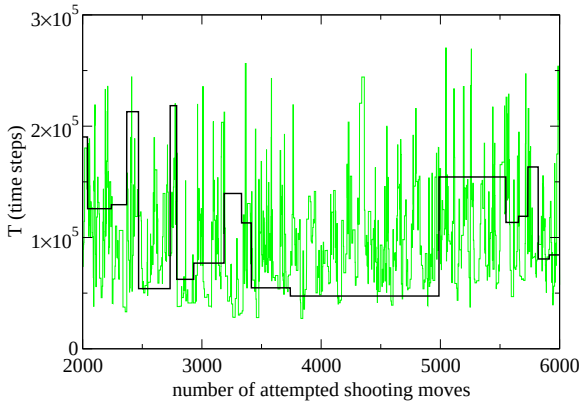

Figure 3.10: Variation in isomerization transition time T over the course of long TPS runs. The bold black line shows results for a simulation using shooting displacements with $\Delta p = 10^{-1}$, while the thin green line corresponds to $\Delta p = 10^{-100}$.

trajectory space. Roughly half of these trial moves are generated by shooting. The other half are generated by a procedure called "shifting" [18], in which short trajectory segments are added to and subtracted from the ends of an existing path.

Figure 3.9 shows the fraction of attempted shooting moves that are accepted in TPS simulations of the diatomic isomerization with a rough barrier. While standard shooting moves are accepted with low frequency, any desired acceptance ratio can be obtained by using the precision shooting technique. Figure 3.10 shows changes in transition time T over the course of two TPS runs with shooting displacements of $\Delta p = 10^{-1}$ and $\Delta p = 10^{-100}$. A dramatic difference in the efficiency of generating qualitatively different trajectories for the two cases is evident.

To assess the improvement in sampling efficiency achieved with precision shooting, we quantify the computational effort necessary to generate statistically independent transition pathways. More specifically, we calculate the autocorrelation function

$$c(n) = \frac{\langle \delta T(0) \delta T(n) \rangle}{\langle \delta T^2 \rangle}, \tag{3.13}$$

where $\delta T(n) = T(n) - \langle T \rangle$ is the deviation of the transition time after the n-th shooting move from its average $\langle T \rangle$, as calculated from all collected trajectories [18, 67]. Rapid decay of $c(n)$ indicates an efficient sampling of trajectories. Figure 3.11 shows the logarithm of $c(n)$ for different shooting displacement magnitudes along with the "decorrelation time" ν, defined as the number of successive shooting moves after which the correlation function decays to a value less than 1/2. The maximal sampling efficiency is achieved

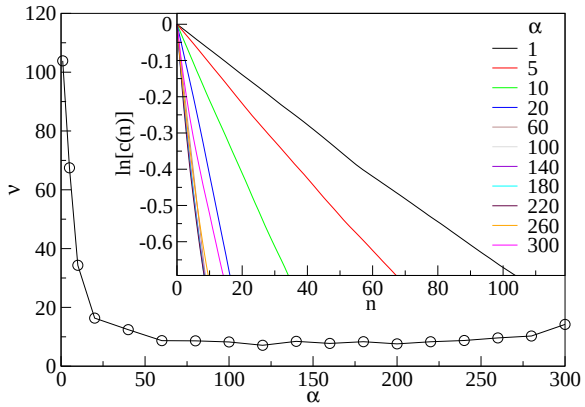

Figure 3.11: Number of shooting moves required to generate a statistically independent isomerization trajectory using shooting displacements of various sizes, with $\Delta p = 10^{-\alpha}$. Inset shows decay of correlation $c(n)$ in the transition time T following n attempted shooting moves. The "decorrelation time" ν is defined as the value of n beyond which $c(n) < 1/2$.

for shooting displacements with $\Delta p \approx 10^{-100}$. Improvement over the largest displacement we considered ($\Delta p = 10^{-1}$) is more than ten-fold. Following Bolhuis [67], we also investigate as a measure of decorrelation changes in the bond length x midway in time through the crossing event. Decay of correlations in this quantity, and the implied dependence of sampling efficiency on shooting displacement size, mirror those reported for the transition time T.

As Fig. 3.11 illustrates, sampling is comparably efficient for a broad range of displacement sizes between $\Delta p = 10^{-60}$ and $\Delta p = 10^{-260}$. In this regime, the efficiency gain due to increased acceptance rates for smaller shooting moves is compensated almost exactly by the efficiency loss due to increased similarity between the shooting trajectory and its base trajectory. Using equation (3.3), the time T_{id} over which a shooting trajectory with displacement size $10^{-\alpha}$ cannot be resolved from its base trajectory can be approximated by the time required for the displacement size to reach 10^{-15},

$$T_{\mathrm{id}} \approx \frac{1}{\lambda_1} \ln \frac{10^{-15}}{10^{-\alpha}} = (\alpha - 15) \frac{\ln 10}{\lambda_1} \,. \tag{3.14}$$

For our system, $1/\lambda_1 \approx 150\,\Delta t$ and therefore $T_{\mathrm{id}} \approx 10^5$ time steps for the smallest displacement with $\Delta p = 10^{-300}$. Even for this small displacement size, T_{id} is only 30% of the total trajectory length L and efficient sampling is still possible. If δx_0 is decreased further, T_{id} will become comparable to L and sampling efficiency will decrease accordingly. (To extend the precision shooting algorithm to shooting displacements smaller than 10^{-308}, the

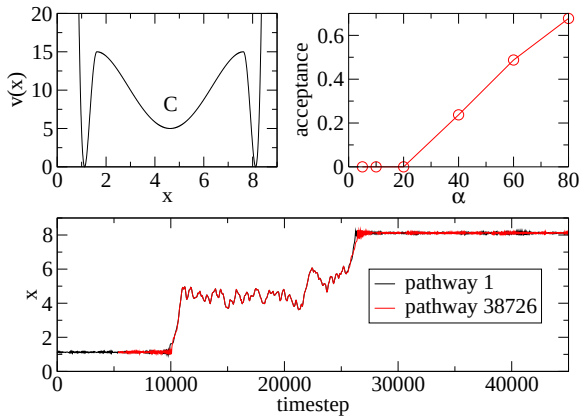

Figure 3.12: TPS simulation of isomerization dynamics which must proceed through a deep intermediate energy minimum. Top left: Interaction potential between particles #1 and #2. Top right: Acceptance ratio of shooting moves as a function of displacement size, $\Delta p = 10^{-\alpha}$. Bottom: Bond length x as a function of time for the starting trajectory and for a pathway obtained after many shooting and shifting moves. The two trajectories have been shifted in time to highlight their similarity in the vicinity of the intermediate state.

smallest representable number in double precision, exponents can be conveniently stored separately as integer numbers.) For a displacement with $\Delta p = 10^{-60}$, the largest size that leads to optimum efficiency, T_{id} amounts to only 5% of the total trajectory length.

3.4 Metastable intermediate states

By extending the time span over which a shooting trajectory tracks its base trajectory, the algorithm proposed in this work can substantially increase the efficiency of TPS simulations that suffer from poor acceptance of shooting moves. The method is fully consistent with deterministic dynamics and faithfully reproduces the divergent behavior of arbitrarily small displacements in phase space. We emphasize, however, that the method does not solve all problems whose primary symptom is a low shooting acceptance rate. Most importantly, it does not overcome challenges associated with metastable intermediate states. In this section we explore this difficulty in the context of diatomic isomerization.

In order to explore the consequences of metastable intermediates, we have modified the diatomic potential $v(x)$ to include a deep minimum mid-

way between contracted and extended states (see Fig. 3.12),

$$v(x) = \begin{cases} h_1 \left[1 - q(x)^2/w_1^2\right]^2 & \text{if } q(x) < 0\,, \\ h_1 \left[1 - (q(x) - 2w_2)^2/w_1^2\right]^2 & \text{if } q(x) > 2w_2\,, \\ h_1 - h_2 \left[1 - (q(x) - w_2)^2/w_2^2\right]^2 & \text{else}\,. \end{cases} \tag{3.15}$$

Here, $q(x) = x - (r_c + w_1)$, $w_1 = 0.5$, $w_2 = 3$, $h_1 = 15$, and $h_2 = 10$. Limited by machine precision, standard shooting moves fail completely in this case: Even shooting moves initiated near the intermediate minimum C rapidly separate from their base trajectories and with high probability do not escape to stable state A or B. Only with the precision shooting technique, using a displacement size smaller than 10^{-20}, are we able to conduct successful shooting moves. This success does not indicate, however, that trajectory space is sampled efficiently: A comparison of the first trajectory[1] with a pathway obtained after many thousands of shooting moves shows that those parts of the trajectory spent within the intermediate are not resampled at all (see Fig. 3.12); they are numerically identical.

Transitions involving strongly metastable intermediates are in fact fundamentally problematic for TPS methods, unless the dynamics of intermediates' appearance and disappearance can be identified as distinct kinetic substeps. If the typical time spent in C is manageable in a computer simulation, then the intermediate does not pose a problem even to the standard shooting move. If, on the other hand, the free energy barriers delimiting the intermediate state are large compared to typical thermal excitations, then escaping C will itself be a rare event. In such cases, typical transitions from A to B require at least two unlikely fluctuations (activating entry and exit of each intermediate state), well separated in time. Any shooting move that perceptibly modifies dynamics between these rare fluctuations will be rejected with high probability. Precision shooting can readily generate subtly modified pathways that remain reactive but cannot be expected to effectually switch between reactive trajectories that follow substantially different courses through the intermediate state. As TPS leaves a system's natural dynamics unchanged, it can eliminate only the largest time scale associated with a rare event. Without resorting to methods that prescribe in some sense the detailed route between stable states, one can overcome the challenge of metastable intermediates with TPS only by subdividing transition dynamics into several steps, each of which involves a single dynamical bottleneck.

[1] A first trajectory is constructed in the following way: From the border of state A, trajectories are shot into the intermediate state, which is divided into small windows along the direction of the coordinate x. Starting with the first of these windows, trajectories are accepted if they cross the border to the next one. After accepting a few such trajectories, the simulation moves on to the next window, eventually leading to a trajectory that crosses the intermediate from A to B. Note that a similar procedure is used in forward flux sampling [61].

Part II

Structural transformations of nanocrystals

Chapter 4

Mechanisms of the wurtzite to rocksalt transformation in CdSe nanocrystals

We study the pressure-driven phase transition from the four-coordinate wurtzite to the six-coordinate rocksalt structure in CdSe nanocrystals with molecular dynamics computer simulations. With an ideal gas as pressure medium, we apply hydrostatic pressure to spherical and faceted nanocrystals ranging in diameter from 25 Å to 62 Å. In spherical crystals, the main mechanism of the transformation involves the sliding of (100) planes, but depending on the specific surface structure we also observe a second mechanism proceeding through the flattening of (100) planes. In faceted crystals, the transition proceeds via a five-coordinated hexagonal structure, which is stabilized at intermediate pressures due to dominant surface energetics.

4.1 Introduction

The thermodynamic properties of nanosized particles can differ significantly from those of the corresponding bulk materials due to the large surface to volume ratio. In particular, the kinetics and the mechanism of first order phase transitions are strongly affected by the surface free energetics of such systems. In a series of recent experiments, Alivisatos and coworkers [9–11, 73–75] have demonstrated that the pressure-induced transition from the four-coordinate wurtzite structure to the six-coordinate rocksalt structure in CdSe is strongly influenced by crystal size. The transition pressure increases by about a factor of two as the size of the sample is decreased from macroscopic dimensions to the nanometer scale [73], and also the kinetics of this highly activated process change markedly as function of size [9, 10].

A considerable amount of work has been concerned with the wurtzite

to rocksalt transition and possible intermediate structures in the bulk [15, 76–83], where, for CdSe, the transformation occurs at approximately 2.5 GPa [84, 85]. In a recent molecular dynamics study of bulk CdSe, Shimojo *et al.* identified two main mechanisms [15]. The first mechanism was previously proposed by Tolbert and Alivisatos and involves the flattening-out of parallel (100) planes [73], the second mechanism is realized through sliding of parallel (100) planes along the [010] direction. Using transition path sampling methods [17], Zahn, Grin and Leoni recently showed that the second mechanism is highly preferred and the transition does not involve a concerted motion of atoms, but occurs via nucleation and growth [83].

In the nanocrystal, surface effects may significantly alter the transition mechanism. In this Letter, we use molecular dynamics simulation to study the structural transformation of CdSe nanocrystals of various shapes and sizes. Our simulations reveal that transition mechanisms are strongly shape-dependent: Most *spherical* nanocrystals transform through the bulk mechanism of sliding (100) planes [83], but the other mechanism, unfavored in the bulk, also occurs. For *faceted* nanocrystals the transformation proceeds through an intermediate five-coordinate structure, that is unstable in the bulk but is stabilized by surface effects in the nanoparticle.

4.2 Simulation details

In all our simulations we use the empirical pair potential for CdSe developed by Rabani [85], designed to reproduce the lattice and elastic constants of bulk CdSe as well as the bulk wurtzite to rocksalt transition pressure of 2.5 GPa [83, 85].

A crucial point in the simulation of nanoparticles under pressure is the choice of pressure medium. As such we use an ideal gas of non-interacting particles that interact with the crystal atoms through the soft-sphere pair potential $u(r) = \epsilon\,(\sigma/r)^{12}$. As the equation of state of the ideal gas is known analytically, the pressure can be easily tuned by controlling the density of the pressure bath. We set $\epsilon = 1$ kJ/mol and $\sigma = 3.0$ Å, large enough to prevent infiltration of gas particles into the nanocrystal. For an efficient calculation of the forces needed in the molecular dynamics simulation we use cell lists [86] and a cutoff of 2σ for the interactions between ideal gas particles and the crystal atoms. To reduce the number of ideal gas particles required to exert a given pressure, these particles fill only a thin layer around the crystal. The volume that is occupied by the gas consists of all cells, already defined for the cell lists, that can hold possible interaction partners of the crystal atoms. This minimal volume of cells is updated every time step and particles leaving the volume are no longer considered. The loss of gas particles is compensated by randomly introducing new gas particles on the cell walls confining the volume with statistics appropriate for an ideal gas at

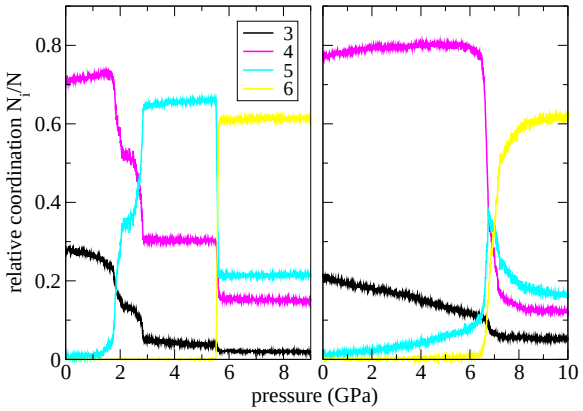

Figure 4.1: Relative coordination N_i/N as a function of pressure for a *faceted* $Cd_{847}Se_{847}$ crystal (left) and a *spherical* $Cd_{1456}Se_{1457}$ crystal (right). The wurtzite to rocksalt transition can be easily identified through the change from predominant four-coordination to six-coordination. In faceted crystals (left), the stable five-coordinated h-MgO structure is observed at intermediate pressures. For spherical crystal shapes (right), the transformation from four-coordination to six-coordination occurs directly and the intermediate h-MgO structure is never observed.

temperature T and pressure P. When, by movement of crystal atoms, a new cell is added to the gas atmosphere, it is filled with a number of particles drawn from an appropriate Poisson distribution; when a cell is removed from the gas atmosphere, the particles in this cell are no longer considered. In this method, the gas serves as a barostat as well as a thermostat. The exerted pressure is hydrostatic (we have numerically verified that in our simulations the forces tangential to the crystal surfaces vanish on average) and can be simply controlled by adjusting the number of ideal gas particles.

To study the effect of surface structure and crystal size on the transition mechanism, we use nanocrystals of two different shapes, ranging in radius from 12.6 Å to 31.5 Å, consisting of 300 to 4500 atoms. Spherical nanocrystals are cut from a large wurtzite lattice with centers randomly distributed over the wurtzite unit cell. The resulting nanocrystals have disordered surfaces and, due to the randomly chosen positions of the center, slightly differ in atom number even for equal sizes. Faceted nanocrystals with well defined surface structure were obtained by cleaving the bulk lattice along equivalent (100) wurtzite planes and at (001) and $(00\bar{1})$ planes perpendicular to the [001] direction of the c-axis. Similar to crystals used in experiments [87], the aspect ratio c/a was chosen to vary around a value of 1.25, slightly elongated along the c-axis. The resulting nanocrystals have stable low-index surfaces and can be seen as a short version of the nanorods studied in recent

experiments [75].

All of our simulations are carried out at $T = 300$ K and follow the same scheme: A single crystal is initially equilibrated for 15 ps at zero pressure. Then the pressure is increased every 10 ps in steps of 0.25 GPa until a maximum pressure of 8–11 GPa is reached, depending on crystal size and shape. The equations of motion are integrated using the velocity Verlet algorithm [70] with a time step of 2 fs. The mass of the ideal gas particles is 10 amu. The longest of a total of 75 simulation runs have a length of 450 ps, and about 350.000 gas particles are required to apply a maximum pressure of 11 GPa.

4.3 Shape-dependent transformation mechanism

Apart from analyzing the simulation runs by visual inspection, we calculate the number of atoms N_i with i nearest neighbours. Here, atoms are defined to be nearest neighbors if they are closer than 3.3 Å, the location of the first minimum of the radial distribution function. The transition is monitored by plotting the fractions N_i/N of i-coordinated atoms as a function of pressure (see Fig. 4.1). Whereas spherical crystals directly transform from the wurtzite to the rocksalt structure, as indicated by the sudden change from four- to six-coordination (Fig. 4.1, right), faceted crystals take a five-coordinated structure at intermediate pressures (Fig. 4.1, left). The latter transition is depicted in Fig. 4.2 A. Starting on one side of the crystal, the puckered wurtzite (001) layers are leveled out, leading to a compression of the whole crystal along the wurtzite c-axis. The resulting structure is similar to hexagonal BN and was named h-MgO, as it was first discovered as a metastable phase of MgO [77]. Recently, it was considered as a metastable intermediate in the wurtzite to rocksalt transition in bulk CdSe [15]. Its high energy, however, rules it out as a possible stable intermediate structure for the bulk transformation [15]. To confirm this, we calculate the enthalpy of bulk CdSe in constant-pressure Monte Carlo simulations (see Fig. 4.3) [85]. The enthalpy of h-MgO is never the lowest, rendering this structure unaccessible in the bulk. During the transition, the (001) planes of the crystal are flattened out and the surface atoms are shifted from a three-coordinate to a four-coordinate environment. This favorable reorganization of the surface stabilizes the five-coordinate h-MgO structure in the nanocrystal sufficiently with respect to the bulk to make this structure accessible during the transition. As expected, we observe a strong decrease in the wurtzite to h-MgO transition pressure with decreasing crystal size [1]: Whereas in the largest

[1] While the transition pressure to the intermediate h-MgO structure decreases with decreasing crystal size, the phase transition pressure for the overall wurtzite to rocksalt transition, obtained by averaging the up- and downstroke transition pressures, shows the opposite trend, as observed in the experiments.

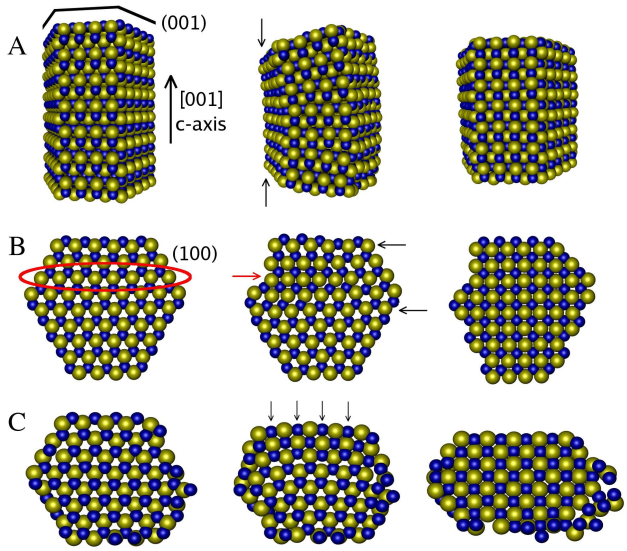

Figure 4.2: (**A**) Wurtzite to *h*-MgO transition mechanism in a *faceted* $Cd_{847}Se_{847}$ crystal. The transition occurs through a compression of the crystal in the direction of the wurtzite c-axis, the puckered (001) layers, indicated in the figure, are flattened out. (**B**) *h*-MgO to rocksalt transition mechanism in a *faceted* $Cd_{847}Se_{847}$ crystal, seen down the wurtzite c-axis. The transition starts on the surface with the displacement of the outer atoms of a (100) plane into the crystal (red arrow), forming the rocksalt nucleus. Next-nearest neighbour planes preferably slide in alternate directions (black arrows). (**C**) Wurtzite to rocksalt transition mechanism in a *spherical* $Cd_{366}Se_{359}$ crystal, seen down the wurtzite c-axis. Starting on the surface, atoms across six-membered hexagonal rings come together to form the rocksalt structure, flattening out the wurtzite (100) planes.

crystals the *h*-MgO structure is stable over a small pressure range only, we observe an immediate wurtzite to *h*-MgO transition at zero pressure for a $Cd_{72}Se_{72}$ crystal. The critical size, at which the *h*-MgO structure becomes unstable with respect to the wurtzite structure, increases with increasing pressure.

The *h*-MgO to rocksalt transition is shown in Fig. 4.2 B. The mechanism involves the sliding of parallel (100) planes and is equal to the mechanism observed in the bulk [83], except for the compression along the wurtzite c-axis, which has already occurred during the wurtzite to *h*-MgO transition. The transition nucleates at the surface where the outer atoms of a (100) plane move along the [010] direction into the crystal, transforming the hexagonal 120° bond angle along the (100) plane into the cubic 90° angle. Cadmium and selenium atoms along the plane form cubic bonds with

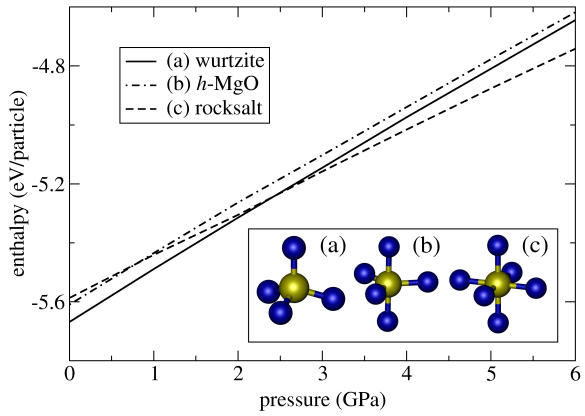

Figure 4.3: Enthalpy as a function of pressure for bulk CdSe in the wurtzite, h-MgO and rocksalt structure obtained from Monte Carlo simulations at constant pressure and temperature. Inset: Atomic configuration for the three structures.

the next selenium and cadmium atoms along the two neighbouring (100) planes, initiating the rocksalt nucleus. This displacement and change of angle quickly propagates along the (100) plane, at the same time starting the transition in the adjacent (100) planes. This sliding plane motion then moves through the whole crystal. The shearing directions of the planes, parallel or antiparallel, vary and determine the overall shape change of the crystal during the transformation. In some of the larger faceted crystals, we observe the formation of grain boundaries during the h-MgO to rocksalt phase transition. If the crystal is sufficiently large, the transition can occur in different parts of the crystal simultaneously. Grain boundaries then form where two rocksalt structures with different orientation meet.

The majority of spherical crystals transforms through the sliding of (100) planes mechanism observed in the h-MgO to rocksalt transition of faceted crystals. However, the transition occurs directly from the wurtzite to the rocksalt structure and the h-MgO structure is not observed. The compression of the wurtzite lattice in the direction of the c-axis takes place at the same time as the shearing motion of (100) planes. In faceted crystals with flat (001) surfaces terminating the crystal in the c-direction, there is no transition motion in the [001] direction during the h-MgO to rocksalt transition, leading to a flat, low-index (001) rocksalt surface. The situation is different in spherical crystals. As the spherical shape is not suitable to accommodate low-index rocksalt faces, the transition motion is more complicated. In an effort to produce low-energy rocksalt surfaces, wurtzite (100) planes can, in addition to the shearing motion, move in the [001] direction during the transition. Moreover, the formation of lattice defects is found in spherical

crystals of all sizes. These include grain boundaries between rocksalt domains of different orientation, as well as dislocations. Whereas only one 14 Å crystal displays a grain boundary after the transition, over 80% of the largest sized crystals transform with lattice defects.

In a single, spherical $Cd_{366}Se_{359}$ nanocrystal we observed a second transition mechanism (Fig. 4.2 C). Here, the transition does not involve a shearing motion but a flattening-out of (100) planes. Starting on one side of the crystal, atoms across six-membered hexagonal rings come together to form the rocksalt structure and at the same time level differences in the c-direction are flattened out. A transition through this mechanism requires a significant overall shape change of the crystal, transforming a sphere into an oblate ellipsoid. This mechanism was previously proposed and later discarded by Alivisatos *et al.* [11,73], because the significant shape change accompanying the transition through this mechanism could not be confirmed by experimental data [11]. The exact prerequisites for the two different mechanisms to happen are unclear, but seem to be strongly dependent on the specific surface structure. Although we observe the alternative mechanism in this pureness only in one case, many 14.0 Å radius crystals and even a few larger crystals, spherical and faceted, display a mixture of both mechanisms.

Although our simulations show many similarities with actual experiments on CdSe nanocrystals, there are also important differences. The *h*-MgO structure is not observed in experiments and the transition mechanism proposed in agreement with experimental data involves the sliding of (001) planes [11], not (100) planes as in our simulations. However, a possibly metastable *h*-MgO structure might have a lifetime too short to be resolved in the experiments. Also, crystals used in experiments are covered with surfactants, which strongly influence surface energies [88]. The presence of such a surface passivation layer may block the transition path to the *h*-MgO structure and render it unaccesible in the experiments, favoring a different transition mechanism. Computationally quite a challenge, a detailed description of surface passivation nevertheless has to be included in future simulations. Recently, progress has been made towards a description of surfactants through simple force fields [89].

Another difference between simulation and experiment lies in the accessible time scales. Although our rate of pressurization is lower than in comparable simulation studies [13,14,90–92], it is still many orders of magnitude larger than in experiments, where observation times of several minutes up to hours are common [11]. This has two effects: First, sudden changes in pressure can lead to structural instabilities, possibly changing the statistical weight of different transition mechanisms. Second, on the short timescale available in simulations, the use of significantly higher pressures is necessary to observe the transition. This again may lead to different transition mechanisms. Despite of the large difference in timescale, the transitions in our simulations happen at pressures comparable to those found in ex-

periments. This indicates that transition barriers are quite different in the two cases, again pointing to the important role of surface passivation. To study the transition at the thermodynamic transition pressure, transition path sampling methods will be applied in future work [17, 83].

The observation of grain boundaries in our simulations is another striking difference to the experiments, where nanocrystals transform free of defects over many pressure cycles [10, 73]. The main mechanism in the formation of these defects is surface-related: When the simple sliding-planes motion would introduce too many surface defects in the final rocksalt structure, as in the case of the spherical crystals, different domains of rocksalt are formed to minimize the surface free energy by the formation of low-index rocksalt surfaces. In faceted crystals, the sliding-planes motion alone leads to well defined rocksalt surfaces. The formation of lattice defects in this case would destroy these surfaces, introducing high-energy steps and edges. Nevertheless, grain boundaries can form in faceted crystals if, due to the high degree of over-pressurization, the transformation is triggered in different parts of the crystal simultaneously. Recently, a simulation study of GaAs nanocrystals showed results similar to our work [90, 91]: Larger spherical crystals are found to transform with grain boundaries, whereas faceted crystals, more often display single-domain behaviour.

4.4 Summary

In conclusion, we show that the structural transformation in CdSe nanocrystals is dominated by surface effects. The transformation nucleates at the surface and proceeds via the bulk mechanism of sliding (100) planes, but depending on the details of the surface structure, other mechanisms involving the flattening of (100) planes or the formation of a five-coordinate intermediate structure are possible. Moreover, our simulations show that in nanocrystals it is possible to stabilize structures that are unstable in the bulk by controlling the overall shape and surface structure of the particles.

Chapter 5

An efficient transition path sampling algorithm for nanoparticles under pressure

We apply transition path sampling to the simulation of nanoparticles under pressure. As a barostat we use a bath of ideal gas particles that form a stochastically updated atmosphere around the nanoparticle. We justify this algorithm by showing that it preserves the distribution of an ideal gas at constant temperature and pressure by satisfying detailed balance. Based on this result, we present a simple and efficient transition path sampling scheme for the study of activated processes in nanoparticles under pressure. As a first application, we investigate the h-MgO to rocksalt transformation in faceted CdSe nanocrystals. Starting from an artificial mechanism involving a uniform motion of all atoms, trajectories quickly converge towards the dominant mechanism of nucleation and growth along parallel (100) planes.

5.1 Introduction

The physical and chemical properties of nanoscale matter are dominated by surface effects and can deviate from the properties of the bulk material in both quantitative and qualitative ways. For instance, pressure induced structural phase transformations show strong dependence on particle size: nanocrystals transform at higher pressures and show much stronger hysteresis than the corresponding bulk materials [10, 73, 93–96]. Furthermore, the enhanced role of surface free energy can generate unique structural transformation mechanisms and stabilize crystal structures that are unstable in the bulk [97] .

Molecular dynamics computer simulations can play a key role in identifying the atomistic details of pressure-induced structural transformations in nanocrystals [14,16,90–92,97–99] as well as in the bulk [15,80,83]. An impor-

tant issue in the simulation of nanoparticles under pressure is the method through which hydrostatic pressure is applied. While methods exist that allow pressurization of a nanoparticle in vacuum [100], an explicit pressure medium must be included to model the experimental situation more closely. Two main methods have been proposed so far: In the first, the medium consists of Lennard-Jones particles and the pressure is controlled by a variant of the extended Lagrangian approach [14, 16, 90, 91, 98, 100], changing the dynamics of the particles. In the second method, the pressure medium is made up of soft sphere particles, for which the equation of state is known [92,100]. Here, the pressure is controlled by a single parameter of the interaction potential, although the volume accessible to the pressure medium, a necessary parameter for exactly tuning the pressure, can only be estimated. In both methods, special care must be taken in choosing the parameters of the pressure medium: for the exerted pressure to be hydrostatic, the dynamics of the pressure medium must not slow down significantly even at pressures of tens of GPa.

In our simulations we use a recently developed ideal gas barostat [97, 101], where the nanoparticle is surrounded by a thin atmosphere of non-interacting particles that are stochastically introduced on the surface of the atmosphere, with appropriate statistics for an ideal gas at the desired temperature and pressure. In this approach, the pressure is easily controlled by adjusting the number of gas particles in the simulation. Furthermore, there is no need for an additional thermostat because the temperature of the system is set by the gas particles' velocity distribution. As there are no interactions in the gas, the exerted pressure is ideally hydrostatic at all times.

The ideal gas barostat seems intuitively correct and can be implemented efficiently in a simulation. Yet, its usefulness can only be established by showing that it indeed preserves the distribution of a system at constant pressure and temperature. While such a demonstration is straightforward for simple algorithms, complications arise in our case due to the stochastic nature of the pressure bath. As the number of gas particles in the simulation is not constant and particles may be added or removed from the simulation in different ways, a careful analysis of the different steps of the algorithm is necessary.

In this chapter, we show that the ideal gas barostat, appropriately implemented, satisfies detailed balance and therefore preserves the correct equilibrium distribution of a nanocrystal immersed in an ideal gas at constant temperature and pressure. This result allows us to implement a simple transition path sampling scheme [17,18] for the simulation of activated processes in nanoparticles under pressure. We apply the algorithm to study the recently observed mechanism of the h-MgO to rocksalt transformation [97] in a model of faceted CdSe nanocrystals at the bulk transition pressure of 2.5 GPa. Starting from a trajectory displaying a transformation mechanism

previously suggested by Tolbert and Alivisatos [73], the simulation quickly finds its way to the primary mechanism, which involves the sliding of parallel (100) planes.

The chapter is organized in the following way. In section 5.2 we review the details of the ideal gas pressure bath and rewrite the algorithm in a way suitable for further analytical treatment. Section 5.3 contains the proof of detailed balance for this algorithm. Based on this result, we present an efficient transition path sampling algorithm for nanoparticle and pressure bath in section 5.4. A discussion of our simulation results on CdSe nanocrystals is given in section 5.5.

5.2 Ideal gas pressure bath

Our pressure bath consists of particles that do not interact with each other, but interact with atoms in the nanoparticle by a repulsive soft sphere potential of the form

$$u(r) = \begin{cases} \epsilon \left[(r/\sigma)^{-12} - (r_{\text{cut}}/\sigma)^{-12} \right] & \text{if } r < r_{\text{cut}} \,, \\ 0 & \text{if } r \geq r_{\text{cut}} \,. \end{cases} \tag{5.1}$$

We set $\epsilon = 1\,\text{kJ/mol}$ and choose σ large enough to prevent gas particles from penetrating the nanoparticle (for CdSe nanocrystals a value of 3 Å is sufficient); we cut the potential at a value of $r_{\text{cut}} = 2\sigma$. To compute the forces between gas and crystal, we use the cell list method [70] with cells of sidelength $l_{\text{cell}} = r_{\text{cut}} + 0.1\,\text{Å}$.

The gas particles fill a thin atmosphere around the nanocrystal. Using the cell list method, it is convenient to define this atmosphere as consisting of all cells that can hold possible interaction partners of crystal atoms. Thus, the outer boundary of the ideal gas atmosphere consists of flat rectangular parts, as illustrated in Fig. 5.1. Gas particles that leave this atmosphere are no longer considered while new particles are introduced on the surface of the atmosphere with the correct statistics for an ideal gas at pressure P and temperature T. The equations of motion of all particles are integrated using the Velocity Verlet algorithm [70]. Note that the stochastic treatment of the pressure bath bears some resemblance to an existing method for modeling solvent dynamics [102].

The following algorithm differs in some details from the one presented in an earlier paper [101], which for subtle reasons does not precisely satisfy detailed balance (for a discussion of this issue see section 5.3.4). However, we have verified that results obtained using the old version [97, 101] remain unchanged when using the correctly balanced algorithm enumerated below and depicted in Fig. 5.2:

1. Propagate the positions of the crystal atoms for one time step (first step of the Velocity Verlet algorithm).

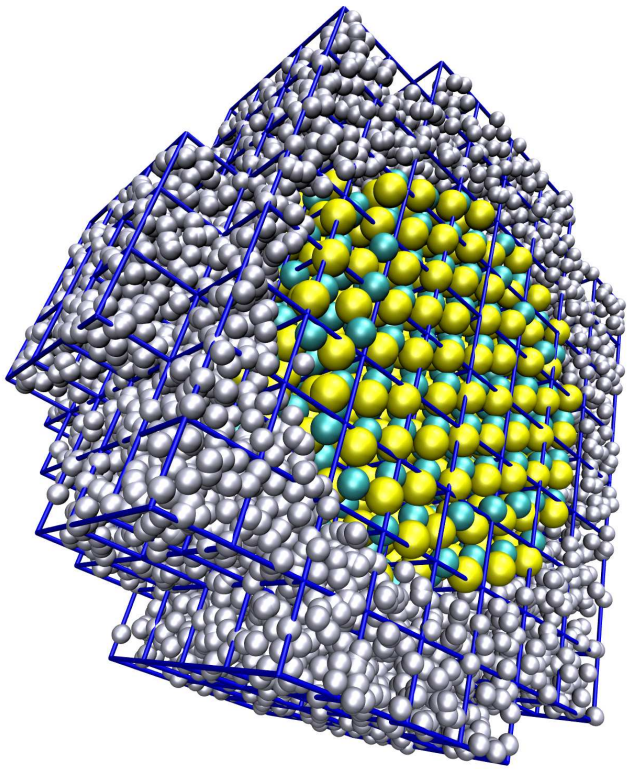

Figure 5.1: Cross section of a CdSe nanocrystal (cyan and yellow) immersed in the pressure bath of ideal gas particles (gray). The gas occupies only a thin layer, or atmosphere, around the nanocrystal. Built from rectangular cells (blue grid), this atmosphere adapts to the shape of the crystal and thus always provides the amount of gas particles necessary to exert the desired pressure.

2. Check if the boundaries of the atmosphere must change due to movement of crystal atoms. Remove cells and gas particles therein that are no longer needed.

3. Propagate the positions of the gas atoms for one time step (first step of the Velocity Verlet algorithm). Note that for this step forces obtained with the crystal atoms in their original position, i.e., before step 1, are used. Remove all gas particles that have left the atmosphere.

   Place n_{in} gas particles on the surface of the atmosphere, with positions drawn from a uniform distribution. n_{in} ("in" for "inject") is a number drawn from a Poisson distribution with an average of

   $$\overline{n}_{\mathrm{in}} = (2\pi m k_{\mathrm{B}} T)^{-1/2} A P \Delta t \ , \tag{5.2}$$

   where m is the mass of the gas particles, k_{B} is the Boltzmann constant, T and P are the desired temperature and pressure, A is the total surface area of the atmosphere, and Δt is the time step. Let $\mathbf{n}_x$ denote the local unit normal to the surface pointing into the system. The velocity distribution in that direction is given by

   $$p(v_x) = \frac{m}{k_{\mathrm{B}} T} v_x \exp\left(-\frac{m v_x^2}{2 k_{\mathrm{B}} T}\right) \ . \tag{5.3}$$

   The velocity distribution in directions perpendicular to $\mathbf{n}_x$ is Maxwell-Boltzmann. To take into account that the gas particles will not enter the atmosphere all at the same time but instead at a variety of times equally distributed over the period of one time step, propagate them for a random time step between zero and Δt, drawn from a uniform distribution. Keep only those particles that stay within the boundary of the atmosphere.

4. If necessary, add new cells and fill them with N_{app} gas particles. N_{app} ("app" for "appear") is a number drawn from a Poisson distribution with an average of

   $$\overline{N}_{\mathrm{app}} = \frac{P V_{\mathrm{app}}}{k_{\mathrm{B}} T} \ , \tag{5.4}$$

   where V_{app} is the total volume of all added cells. The positions of the particles are uniformly distributed over the cell volume and their velocity distribution is Maxwell-Boltzmann.

5. Compute forces and propagate all velocities of crystal and gas atoms (second step of the Velocity Verlet algorithm).

This algorithm propagates the system, nanocrystal and pressure bath, for one time step.

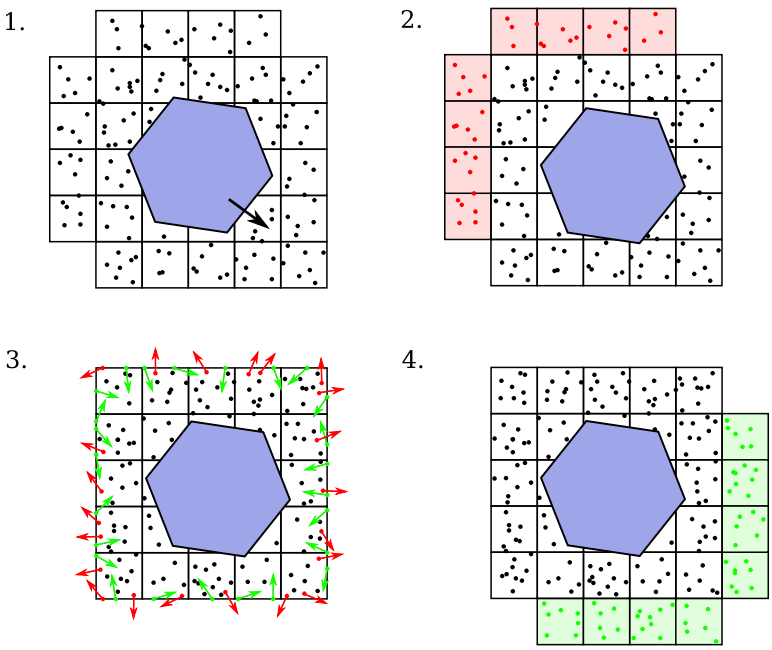

Figure 5.2: Illustration of the ideal gas algorithm. (1) Propagate crystal atoms (blue). (2) If necessary, remove cells and gas particles therein that are no longer needed (red). (3) Propagate gas atoms and remove those (red) that leave the atmosphere. Put new gas particles (green) on the surface and propagate them into the atmosphere; remove those (red) that end up outside. (4) If necessary, add new cells and fill them with gas particles (green).

5.3 Detailed balance for the ideal gas barostat

In the following we show that the algorithm presented in section 5.2 conserves the correct equilibrium distribution of a nanocrystal immersed in an ideal gas at constant pressure and temperature by satisfying a detailed balance condition.

The stationary distribution that the algorithm should conserve is that of a system at constant chemical potential μ, pressure P, and temperature T, where μ is given by an equation of state $\mu(P,T)$. Fixing only these intensive parameters, no constraint is put on extensive parameters like the number of particles N, the volume V, and the energy U. However, the volume of the system is determined by the external condition that the atmosphere of ideal gas particles around the nanoparticle be as small as possible. Thus, a constant μPT-ensemble is realized, and the desired stationary distribution is given by

$$\rho(x) \propto \frac{1}{h^{3N}} \, \mathrm{e}^{-\beta(U+PV-\mu N)} \, , \tag{5.5}$$

where x denotes the microscopic physical state of the system consisting of the positions and momenta of all particles, h is Planck's constant, and $\beta = 1/k_{\mathrm{B}}T$.

A sufficient condition for the conservation of $\rho(x)$ is the detailed balance condition [103]

$$\frac{p(x \to x')}{p(\bar{x}' \to \bar{x})} = \frac{\rho(x')}{\rho(x)} \, , \tag{5.6}$$

where $p(x \to x')$ is the probability to go from physical state x to state x' by the rules of the underlying dynamics and states $\bar{x}$ and $\bar{x}'$ are obtained by inverting the momenta in states x and x', respectively. In the following we explicitly calculate $p(x \to x')$ for the ideal gas barostat and show that equation (5.6) holds.

5.3.1 Combinatorics

As a first step, we cast equation (5.6) into a form that takes into account the fact that a single physical state x can have different representations in a computer simulation.

Let x denote the physical state of a classical system of N indistinguishable particles, which of course makes no reference to which particle is which. In a computer simulation we cannot avoid keeping a list of particle positions in some order. As a result, one physical state is represented by $N!$ "computer" states x_P, each corresponding to a distinct permutation of the stored list of positions, i.e.,

$$x \equiv \{x_1, x_2, \ldots, x_{N!}\} \, . \tag{5.7}$$

Taking this degeneracy of the physical state x into account, the stationary distribution of computer states is given by

$$\rho(x_P) = \frac{\rho(x)}{N!} \; . \tag{5.8}$$

To determine the total rate of transitions $W(x \to x')$ from one physical state x to another x', we must consider all moves $x_P \to x'_{P'}$ that are allowed by a given sampling algorithm:

$$W(x \to x') = \sum_P \sum_{P'} w(x_P \to x'_{P'}) \; . \tag{5.9}$$

Here $w(x_P \to x'_{P'})$ is the transition rate for a specific pair of computer states, depending on the time step Δt, the distribution of computer states $\rho(x_P)$, and the conditional probability $p(x_P \to x'_{P'})$ that a simulation at x_P will move to $x'_{P'}$ in a single step:

$$w(x_P \to x'_{P'}) = \frac{1}{\Delta t} \rho(x_P) p(x_P \to x'_{P'}) \; . \tag{5.10}$$

It seems that in any reasonable algorithm the way a physical state x is treated should not depend on the order in which particle positions are stored, so $\rho(x_P)$ and $\sum_{P'} w(x_P \to x'_{P'})$ do not depend on P. We can therefore write

$$\begin{aligned} W(x \to x') &= \frac{1}{\Delta t} N! \, \rho(x_1) \sum_{P'} p(x_1 \to x'_{P'}) \\ &= \frac{1}{\Delta t} \, \rho(x) \sum_{P'} p(x_1 \to x'_{P'}) \; , \end{aligned} \tag{5.11}$$

where we have used equation (5.8) in the second step. The sum in this equation accounts for the fact that the number of permutations P' accessible from any one permutation P may be different than in the reverse direction. Let's define this number as

$$\mathcal{N}(x \to x') \equiv \sum_{P'} H(x_P, x'_{P'}) \; , \tag{5.12}$$

where

$$H(x_P, x'_{P'}) \equiv \begin{cases} 1, & \text{if } x'_{P'} \text{ is accessible from } x_P \\ 0, & \text{otherwise.} \end{cases} \tag{5.13}$$

Since $H(x_P, x'_{P'})$ need not be a symmetric function, $\mathcal{N}(x \to x')$ may differ from $\mathcal{N}(x' \to x)$. We then have

$$W(x \to x') = \frac{1}{\Delta t} \mathcal{N}(x \to x') \rho(x) p(x_1 \to x'_1) \; , \tag{5.14}$$

if permutations are assigned such that x'_1 is accessible from x_1 in a single step. Here we have assumed that $p(x_1 \to x'_{P'})$ is the same for all accessible $x'_{P'}$, which is true for any algorithm that handles particle lists in a deterministic way.

To illustrate this point, consider a Monte Carlo move $x \to x'$ that attempts to insert a single particle, $N' = N + 1$. If a particle is added to x_P, we need some convention for labeling the new particle position. For concreteness, we choose to place the new position at the end of the original list of positions, but all that is really required is that there is some deterministic scheme for introducing a new label. Thus there is a one-to-one relationship between P and P', and $\mathcal{N}(x \to x') = 1$. If, in the reverse move $x' \to x$, a particle is removed from $x'_{P'}$, and the remaining particles retain their original label ordering, then the label ordering in the resulting state x_P is again completely determined, and $\mathcal{N}(x' \to x) = 1$.

If two particles are added in a single Monte Carlo move, two different resulting computer states $x'_{P'}$ are possible, depending on the order in which the particles are added. Both states are reached with equal probability, and $\mathcal{N}(x \to x') = 2$. In the reverse move, however, with the deterministic removal scheme described above, the resulting state x_P is again completely determined, and $\mathcal{N}(x' \to x) = 1$.

In the ideal gas barostat, more than one particle may be inserted in a single step and particles may be added to the simulation in two distinct ways: They may be injected through the boundary of the atmosphere or they may appear in a lattice cell that has just been added to the simulation. Since multiple orderings of these particles are possible, one computer state x_P could generate many different permutations of the added particles in $x_{P'}$.

The number of distinct computer states $x_{P'}$ accessible from x_P is

$$\mathcal{N}(x \to x') = N_{\text{in}}!\, N_{\text{app}}!\,. \tag{5.15}$$

Similarly,

$$\mathcal{N}(\bar{x}' \to \bar{x}) = N_{\text{in}'}!\, N_{\text{app}'}!\,, \tag{5.16}$$

where "in′" and "app′" refer to particles added in the reverse step $\bar{x}' \to \bar{x}$. Particles that are removed from the simulation, either because they leave the atmosphere or because they are located in a cell that is no longer needed, need not be considered, as long as there is a deterministic removal scheme as discussed above.

Detailed balance requires that the net transition rates between two physical states are identical at equilibrium,

$$W(x \to x') = W(\bar{x}' \to \bar{x})\,. \tag{5.17}$$

Using equation (5.14), we can now write down the corresponding requirement for the permutation-specific transition probabilities:

$$\frac{p(x_1 \to x'_1)}{p(\bar{x}'_1 \to \bar{x}_1)} = \frac{\rho(x')\,\mathcal{N}(\bar{x}' \to \bar{x})}{\rho(x)\,\mathcal{N}(x \to x')}\,. \tag{5.18}$$

Here we have used the fact that the stationary distribution of states is symmetric with respect to inversion of momenta, $\rho(\bar{x}') = \rho(x')$.

Inserting equations (5.5), (5.15), and (5.16) into equation (5.18), we arrive at the final expression of the detailed balance condition,

$$\frac{p(x_1 \to x_1')}{p(\bar{x}_1' \to \bar{x}_1)} = \frac{1}{h^{3\Delta N}} \frac{N_{\mathrm{in}'}!\, N_{\mathrm{app}'}!}{N_{\mathrm{in}}!\, N_{\mathrm{app}}!} \mathrm{e}^{-\beta(\Delta U + P\Delta V - \mu\Delta N)} , \tag{5.19}$$

where $\Delta N = N' - N$, $\Delta U = U' - U$, and $\Delta V = V' - V$ are the differences in particle number, energy, and volume, respectively, between the states x and x'.

5.3.2 Transition probabilities

We proceed by explicitly calculating the transition probabilities for the ideal gas barostat.

The transition probability from computer state x_1 to computer state x_1' can be written

$$p(x_1 \to x_1') = \mathcal{P}_{\mathrm{in}}(N_{\mathrm{in}})\mathcal{P}_{\mathrm{app}}(N_{\mathrm{app}}) \prod_{i\in\mathrm{in}} f_{\mathrm{in}}(\mathbf{q}_i', \mathbf{p}_i') \prod_{j\in\mathrm{app}} f_{\mathrm{app}}(\mathbf{q}_j', \mathbf{p}_j') \, . \tag{5.20}$$

Here, $\mathcal{P}_{\mathrm{in}}(N_{\mathrm{in}})$ and $\mathcal{P}_{\mathrm{app}}(N_{\mathrm{app}})$ are Poisson distributions governing the number of added particles. The positions $\mathbf{q}'$ and momenta $\mathbf{p}'$ of added particles are distributed according to $f_{\mathrm{in}}(\mathbf{q}', \mathbf{p}')$ for injected particles and $f_{\mathrm{app}}(\mathbf{q}', \mathbf{p}')$ for appearing particles. Note that we omit delta-function singularities, associated with deterministic time evolution, that appear symmetrically in $p(x_1 \to x_1')$ and $p(x_1' \to x_1)$. In particular, this concerns the deterministic motion of nanocrystal atoms. A detailed sketch illustrating the algorithm's treatment of the ideal gas is shown in Fig. 5.3. The following discussion separately treats the cases of particle injection and particle appearance.

Particle injection

N_{in} denotes the number of particles that are added to the system through the injection algorithm, where n_{in} particles are put on the boundary of the ideal gas atmosphere and then propagated into the system. The average number of those particles $\overline{n}_{\mathrm{in}}$ is given by equation (5.2). However, in general $N_{\mathrm{in}} \neq n_{\mathrm{in}}$ because there is chance that when propagated a particle will fly through a corner and leave the system (see Fig. 5.3), in which case it is no longer considered. We denote the probability for a successful injection, averaged over all possible injection points on the boundary R, by Π_R.

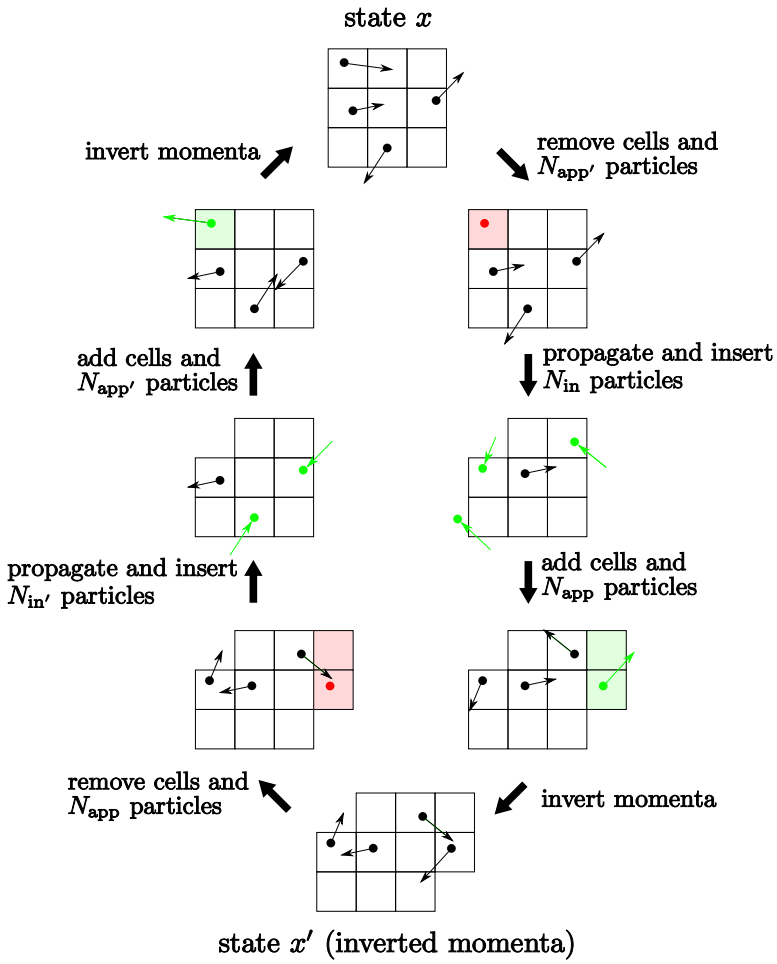

Figure 5.3: Illustration of a transition between two states x and x' of the ideal gas system, nanocrystal not shown. The system starts in state x, with thin arrows indicating particle velocities. In the first step, unnecessary cells and particles therein (red) are removed. In the next step, particles are propagated according to their velocities; if they leave the atmosphere, they are no longer shown. New particles (green) are injected through the surface, with trailing velocity arrows indicating their imaginary positions one time step in the past. In the final step, new cells and particles (green) are added to the atmosphere. In the reverse move the same sequence of steps is passed. Note that in both directions particles are injected through the same boundary, and not all injected particles remain inside.

Dropping the subscript "in" for the moment, the probability to try an insertion for a number of n particles is given by the Poisson distribution

$$\hat{\mathcal{P}}_{\rm in}(n) = \frac{\overline{n}^n \mathrm{e}^{-\overline{n}}}{n!} \; . \tag{5.21}$$

To obtain the probability that a total of N particles survive the insertion process, we must sum over all possible realizations,

$$\mathcal{P}_{\rm in}(N) = \sum_{n=N}^{\infty} \hat{\mathcal{P}}_{\rm in}(n) \binom{n}{N} (\Pi_R)^N (1-\Pi_R)^{n-N} \; . \tag{5.22}$$

Resumming,

$$\begin{aligned} \mathcal{P}_{\rm in}(N) &= \frac{(\Pi_R)^N \mathrm{e}^{-\overline{n}}}{N!} \sum_{n=N}^{\infty} \frac{\overline{n}^n}{(n-N)!} (1-\Pi_R)^{n-N} \\ &= \frac{(\Pi_R)^N \mathrm{e}^{-\overline{n}}}{N!} \sum_{a=0}^{\infty} \frac{\overline{n}^{N+a}}{a!} (1-\Pi_R)^a \\ &= \frac{(\overline{n}\Pi_R)^N \mathrm{e}^{-\overline{n}}}{N!} \sum_{a=0}^{\infty} \frac{\overline{n}^a}{a!} (1-\Pi_R)^a \\ &= \frac{(\overline{n}\Pi_R)^N \mathrm{e}^{-\overline{n}}}{N!} \, \mathrm{e}^{\overline{n}(1-\Pi_R)} \; , \end{aligned} \tag{5.23}$$

we end up with another Poisson distribution

$$\mathcal{P}_{\rm in}(N_{\rm in}) = \frac{\overline{N}_{\rm in}^{N_{\rm in}} \mathrm{e}^{-\overline{N}_{\rm in}}}{N_{\rm in}!} \; , \tag{5.24}$$

with an average number of surviving particles

$$\overline{N}_{\rm in} = \overline{n}_{\rm in} \Pi_R \; . \tag{5.25}$$

For the evaluation of $f_{\rm in}(\mathbf{q},\mathbf{p})$ we adopt a coordinate system with respect to the local boundary of the ideal gas system, with $\mathbf{n}_x$ denoting the local unit normal to the boundary, pointing into the system. In our injection algorithm, a particle's position $\hat{\mathbf{q}}$ with components $(\hat{x},\hat{y},\hat{z})$ is assigned randomly, with $\hat{y}$ and $\hat{z}$ equidistributed on an area A with fixed coordinate $\hat{x} = x_0$. A momentum $\hat{\mathbf{p}}$ is drawn from a distribution $\varphi(\hat{\mathbf{p}})$, along with a timestep τ, equidistributed within $(0,\Delta t)$. Then the particle is propagated to a new position $\mathbf{q} = \hat{\mathbf{q}} + \tau\hat{\mathbf{p}}/m$. The distribution of the set of random variates $(\tau,\hat{y},\hat{z},\hat{\mathbf{p}})$ is given by

$$\hat{f}_{\rm in}(\tau,\hat{y},\hat{z},\hat{\mathbf{p}}) = \frac{\Theta(\tau)\Theta(\Delta t-\tau)}{\Delta t} \times \frac{h_R(\hat{y},\hat{z})}{A} \times \varphi(\hat{\mathbf{p}}) \; , \tag{5.26}$$

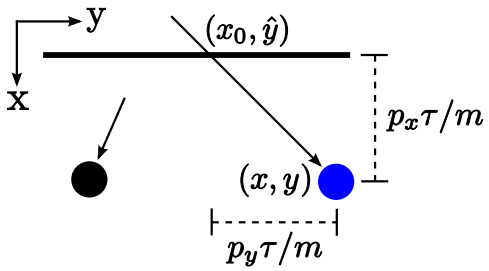

Figure 5.4: Illustration of the variable transformation (5.30), z-direction not shown. The arrows indicate particle velocities, with arrows starting at particle positions one time step in the past. The blue gas atom was injected from the point $(x_0, \hat{y})$ on the surface, and propagated for a time τ to its final position (x, y). The other particle could not have reached its position through an injection event.

where $\Theta(\tau)$ is the step function,

$$\Theta(\tau) = \begin{cases} 1, & \text{if } \tau \geq 0, \\ 0, & \text{if } \tau < 0, \end{cases} \tag{5.27}$$

$h_R(\hat{y}, \hat{z})$ is the characteristic function of the boundary R,

$$h_R(\hat{y}, \hat{z}) = \begin{cases} 1, & \text{if } (\hat{y}, \hat{z}) \text{ lies on } R, \\ 0, & \text{else,} \end{cases} \tag{5.28}$$

and $\varphi(\hat{\mathbf{p}})$ is the momentum distribution

$$\varphi(\hat{\mathbf{p}}) = \frac{\beta^2}{2\pi m^2}\, \Theta(\hat{p}_x)\, \hat{p}_x\, \mathrm{e}^{-\beta \hat{\mathbf{p}}^2/2m} \ . \tag{5.29}$$

We want to calculate the distribution of the position $\mathbf{q} = (x, y, z)$ and momentum $\mathbf{p} = (p_x, p_y, p_z)$ of an inserted particle. Thus we have to transform variables according to

$$\begin{aligned} \tau &= (x - x_0)m/p_x \ , \\ \hat{y} &= y - p_y\tau/m \ = \ y - (x - x_0)p_y/p_x \ , \\ \hat{z} &= z - (x - x_0)p_z/p_x \ , \\ \hat{p}_x &= p_x \ , \\ \hat{p}_y &= p_y \ , \\ \hat{p}_z &= p_z \ . \end{aligned} \tag{5.30}$$

An illustration of this transformation is shown in Fig. 5.4. With $\Delta x \equiv$

$x - x_0$, the determinant of the Jacobian for this transformation reads

$$\left|\begin{array}{cccccc} m/p_x & 0 & 0 & -\Delta x\, m/p_x^2 & 0 & 0 \\ -p_y/p_x & 1 & 0 & \Delta x\, p_y/p_x^2 & -\Delta x/p_x & 0 \\ -p_z/p_x & 0 & 1 & \Delta x\, p_y/p_x^2 & 0 & -\Delta x/p_x \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{array}\right| = m/p_x. \tag{5.31}$$

Transforming the distribution $\hat{f}_{\rm in}(\tau, \hat{y}, \hat{z}, \hat{\mathbf{p}})$ we arrive at

$$\begin{aligned} \tilde{f}_{\rm in}(\mathbf{q}, \mathbf{p}) &= \frac{\varphi(\mathbf{p})}{A\Delta t p_x/m}\, \psi_R(\mathbf{q}, \mathbf{p}) \\ &= \frac{\beta^2}{2\pi m A \Delta t} \mathrm{e}^{-\beta \mathbf{p}^2/2m}\, \psi_R(\mathbf{q}, \mathbf{p})\ , \end{aligned} \tag{5.32}$$

where

$$\psi_R(\mathbf{q}, \mathbf{p}) = \Theta\big[\tau(\mathbf{q}, \mathbf{p})\big]\Theta\big[\Delta t - \tau(\mathbf{q}, \mathbf{p})\big]\Theta(p_x) h_R\big[\hat{y}(\mathbf{q}, \mathbf{p}), \hat{z}(\mathbf{q}, \mathbf{p})\big] \tag{5.33}$$

decides, whether the set $(\mathbf{q}, \mathbf{p})$ could have been produced by the insertion algorithm.

We are interested in the distribution $f_{\rm in}(\mathbf{q}, \mathbf{p})$ of particles that survive the insertion process only, yet $\tilde{f}_{\rm in}(\mathbf{q}, \mathbf{p})$ is nonzero for particles that end up out of bounds after the insertion process. To correct for these particles, we write

$$f_{\rm in}(\mathbf{q}, \mathbf{p}) = \frac{1}{\Pi_R}\, \tilde{f}_{\rm in}(\mathbf{q}, \mathbf{p}) H_R(\mathbf{q})\ , \tag{5.34}$$

where

$$H_R(\mathbf{q}) = \begin{cases} 1, & \text{if } \mathbf{q} \text{ lies inside the boundary } R, \\ 0, & \text{else,} \end{cases} \tag{5.35}$$

is the characteristic function of the considered ideal gas volume. The normalizing constant Π_R, giving the probability that an insertion process will be successful, is obtained by integrating over phase space,

$$\Pi_R = \int \mathrm{d}\mathbf{q}\, \mathrm{d}\mathbf{p}\, \tilde{f}_{\rm in}(\mathbf{q}, \mathbf{p}) H_R(\mathbf{q})\ . \tag{5.36}$$

The final expression for the probability distribution of positions and momenta of inserted particles then reads

$$f_{\rm in}(\mathbf{q}, \mathbf{p}) = \frac{\beta^2}{2\pi m A \Delta t \Pi_R} \mathrm{e}^{-\beta \mathbf{p}^2/2m} \times \phi_R(\mathbf{q}, \mathbf{p})\ , \tag{5.37}$$

where

$$\phi_R(\mathbf{q}, \mathbf{p}) = \psi_R(\mathbf{q}, \mathbf{p}) H_R(\mathbf{q}) \tag{5.38}$$

decides, whether the set $(\mathbf{q}, \mathbf{p})$ could have been produced by a *successful* insertion move.

Volume change

The average number of particles $\overline{N}_{\mathrm{app}}$ added to newly considered ideal gas regions is given by equation (5.4). This average fully specifies the Poisson distribution

$$\mathcal{P}_{\mathrm{app}}(N_{\mathrm{app}}) = \frac{\overline{N}_{\mathrm{app}}^{N_{\mathrm{app}}} \mathrm{e}^{-\overline{N}_{\mathrm{app}}}}{N_{\mathrm{app}}!} \; . \tag{5.39}$$

The appearing particles are distributed uniformly throughout the added volume, and their momenta are Maxwell-Boltzmann distributed,

$$f_{\mathrm{app}}(\mathbf{q}, \mathbf{p}) = \frac{1}{V_{\mathrm{app}}} \left(\frac{2\pi m}{\beta} \right)^{-3/2} \mathrm{e}^{-\beta \mathbf{p}^2/2m} \; . \tag{5.40}$$

5.3.3 Detailed balance

Assembling equations (5.24), (5.37), (5.39), and (5.40), we can now write the transition probability (5.20) in detail,

$$\begin{aligned} p(x_1 \to x_1') = & \\ & \left(\frac{\beta^2}{2\pi m A \Delta t \Pi_R} \right)^{N_{\mathrm{in}}} \left[\frac{1}{V_{\mathrm{app}}} \left(\frac{2\pi m}{\beta} \right)^{-3/2} \right]^{N_{\mathrm{app}}} \mathrm{e}^{-\beta K_{\mathrm{add}}} \\ & \times \frac{\left(PA\Delta t \Pi_R \sqrt{\beta/2\pi m} \right)^{N_{\mathrm{in}}} \mathrm{e}^{-PA\Delta t \Pi_R \sqrt{\beta/2\pi m}}}{N_{\mathrm{in}}!} \\ & \times \frac{(\beta P V_{\mathrm{app}})^{N_{\mathrm{app}}} \mathrm{e}^{-\beta P V_{\mathrm{app}}}}{N_{\mathrm{app}}!} \prod_{j \in \mathrm{in}} \phi_R(\mathbf{q}_j', \mathbf{p}_j') \; , \end{aligned} \tag{5.41}$$

where K_{add} is the total kinetic energy of incoming particles. The last term simply decides, whether a transition between states x and x' is possible within our injection algorithm. Note that we can omit this factor in comparing transition probabilities because it is symmetric with respect to x and x' in the sense that only particles for which an injection is possible are carried back through the boundary in one time step, if their momenta are inverted. Simplifying,

$$\begin{aligned} p(x_1 \to x_1') = \mathrm{e}^{-\beta K_{\mathrm{add}}} & \left[\beta P \left(\frac{\beta}{2\pi m} \right)^{3/2} \right]^{N_{\mathrm{app}} + N_{\mathrm{in}}} \\ & \times \frac{\mathrm{e}^{-\beta P V_{\mathrm{app}} - PA\Delta t \Pi_R \sqrt{\beta/2\pi m}}}{N_{\mathrm{in}}! N_{\mathrm{app}}!} \; . \end{aligned} \tag{5.42}$$

Considering the fact that in the reverse move particles are injected through the same boundary R (see Fig. 5.3), the ratio of forward and backward rates

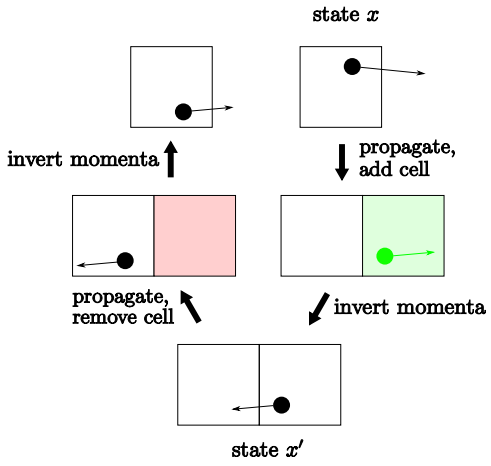

Figure 5.5: Violation of detailed balance by a naive implementation of the ideal gas barostat, as described in the text. Starting in state x, the particle is propagated over the boundary and removed. A new cell is added and filled with a new gas particle (green) to form state x'. The reverse move, in which the added cell (red) is removed again, does not lead back to the original state x.

then takes the form

$$\frac{p(x_1 \to x_1')}{p(\bar{x}_1' \to \bar{x}_1)} = \mathrm{e}^{-\beta\Delta U} \left[\beta P \left(\frac{\beta}{2\pi m}\right)^{3/2}\right]^{\Delta N} \mathrm{e}^{-\beta P \Delta V} \frac{N_{\mathrm{in}'}! N_{\mathrm{app}'}!}{N_{\mathrm{in}}! N_{\mathrm{app}}!} \,. \tag{5.43}$$

With the fugacity of the ideal gas,

$$\mathrm{e}^{\beta\mu} = h^3 \beta P \left(\frac{\beta}{2\pi m}\right)^{3/2} \,, \tag{5.44}$$

we arrive at

$$\frac{p(x_1 \to x_1')}{p(\bar{x}_1' \to \bar{x}_1)} = \frac{\mathrm{e}^{-\beta(\Delta U + P\Delta V - \mu\Delta N)}}{h^{3\Delta N}} \frac{N_{\mathrm{in}'}!\, N_{\mathrm{app}'}!}{N_{\mathrm{in}}!\, N_{\mathrm{app}}!} \,. \tag{5.45}$$

A comparison with equation (5.19) shows that our algorithm indeed satisfies detailed balance.

5.3.4 Other implementations of the barostat

We briefly discuss possible pitfalls that have to be avoided in the implementation of the ideal gas barostat.

As an essential and seemingly complicated feature of the algorithm presented in section 5.2, the boundary of the ideal gas atmosphere is updated

in two different steps. Consequently, particles leave the atmosphere and are injected through the same boundary in forward and backward moves. To illustrate how detailed balance can be violated if this is not the case, consider the following simple algorithm, which is similar to the one used in an earlier paper [101]:

1. Propagate all particles for one time step, remove particles that leave the atmosphere.
2. Inject particles through the surface.
3. Update the atmosphere: remove cells that are no longer needed and add new cells, filled with gas particles.

Clearly, in this algorithm particles are exchanged through a possibly different boundary in the backward move. Consider now a particle A that leaves the atmosphere and is removed from the simulation in step (1). If in step (3) a new cell is added adjacent to the place where A left in step (1), particle A must be restored in the new cell with its proper position and velocity, otherwise the reverse move is not possible and detailed balance is violated. For an illustration of this situation see Fig. 5.5.

As a possible solution, gas atoms leaving the atmosphere into cells that will be added during the same time step should not be removed from the simulation. Furthermore, particle insertions in newly added cells should only be accepted if inversion of momentum will not carry the particle back into the original atmosphere in one time step. However, the algorithm presented in section 5.2 can be implemented more efficiently than this alternative scheme.

5.4 Transition path sampling

In this section we show how the ideal gas barostat can be used within the framework of transition path sampling to study rare events like structural transformations in nanoparticles under pressure.

5.4.1 Transition path ensemble

Many interesting physical and chemical processes require the crossing of high free energy barriers. First order phase transitions near the thermodynamic transition point and chemical reactions are examples of such processes. In particular, structural transformations in nanocrystals display large hysteresis even on experimental time scales [10, 73]. On time scales accessible in a straightforward molecular dynamics simulation, significantly higher pressures and pressurization rates must be applied to observe the transformation. At elevated pressures, however, the transformation mechanism can be different and the formation of grain boundaries is observed [91, 97], in contradiction with experiments [10, 73].

One generally applicable approach to spanning disparate timescales is transition path sampling [17, 18]. Starting from an arbitrary trajectory displaying the transformation, a Monte Carlo procedure is used to generate an ensemble of reactive trajectories which can be analyzed to obtain the most probable transition path, the transition state, and even activation energies [54]. Because in a transition path simulation only reactive trajectories are considered, there is no need to artificially drive the system over the barrier and external parameters like temperature and pressure need not be adjusted from experimental values.

The central object of transition path sampling is the ensemble of reactive trajectories that lead from region A to region B in phase space, perhaps crossing a large free energy barrier. We consider trajectories X of fixed length t, consisting of a sequence of L points in phase space separated by a time step Δt,

$$X = \{x_1, x_2, x_3, \ldots, x_L\}\,, \quad t = (L-1)\Delta t\,. \tag{5.46}$$

Note that x_i in this context denotes different physical states of the system along a trajectory, not different computer representations of a single physical state. The probability to observe a particular transition path is given by

$$\mathcal{P}(X) \propto h_{\mathrm{A}}(x_1) h_{\mathrm{B}}(x_L) \rho(x_1) \prod_{i=1}^{L-1} p(x_i \to x_{i+1})\,. \tag{5.47}$$

Here, the characteristic functions $h_{\mathrm{A}}(x_1)$ and $h_{\mathrm{B}}(x_L)$ return unity only if the start and end points x_1 and x_L of the trajectory lie in regions A and B, respectively, and zero otherwise; $\rho(x_1)$ is the probability density to observe the starting point, and the product gives the probability to move from one point along the trajectory to the next according to the dynamics of the system. In the case of the ideal gas barostat, the functions ρ and p correspond to equations (5.5) and (5.42), respectively.

To sample the ensemble of trajectories defined by equation (5.47), a Monte Carlo procedure is used in close analogy to conventional Monte Carlo sampling in configuration space. Starting from an existing reactive trajectory a new trajectory is created, which is then accepted or rejected according to its probability relative to the old path. However, while trajectory space is sampled with a strong bias, the underlying dynamics of the system remain unaltered. As a result, trajectories obtained from a transition path sampling simulation are true dynamical trajectories, each evolving in time free of any bias.

5.4.2 Shooting and shifting

Two kinds of Monte Carlo moves are typically used to create new dynamical trajectories from old ones [18]. For a nanoparticle in the ideal gas pressure bath they take the following form.

In the first move, called *shooting*, a random point along a given trajectory is selected, dividing the trajectory in two parts. Then, applying the algorithm described in section 5.2, a new trajectory is obtained by propagating the shooting point forward and backward in time. Because of the stochastic nature of the pressure bath, new gas particles are introduced on the surface of the atmosphere every time step and the new trajectory will eventually diverge from the old one. To increase the probability of generating a new trajectory that still leads from A to B, it is advisable to regrow only one part of the old trajectory, either forward or backward in time.

In the second move, the *shifting move*, the trajectory is extended forward or backward for a random period of time, and a segment of equal length is removed from the other end of the trajectory.

In one of the review papers on transition path sampling [18] it has been shown that for stochastic dynamics, under fairly general conditions, the above procedures results in correct sampling of equation (5.47) with a very simple acceptance rule: new trajectories are accepted if they still connect regions A and B, and rejected otherwise. The corresponding acceptance probability takes the form

$$\mathrm{acc}(X \to X') = h_{\mathrm{A}}(x'_1) h_{\mathrm{B}}(x'_L) \; . \tag{5.48}$$

While systems like the ideal gas barostat, where the total number of particles is not constant, were not explicitly considered in Ref. [18], the above acceptance probability is correct provided the underlying dynamics of the system satisfy the condition of microscopic reversibility,

$$\frac{p(x \to x')}{p(\bar{x}' \to \bar{x})} = \frac{\rho(x')}{\rho(x)} \; . \tag{5.49}$$

Again, the state $\bar{x}$ is obtained by inverting the momenta of all particles in state x. As we have shown in section 5.3, the ideal gas barostat fullfills this condition. Therefore the described shooting and shifting moves together with the simple acceptance criterion (5.48) result in correct sampling of trajectories.

5.5 Simulation results

In this section we apply the transition path sampling algorithm presented above to identify the primary mechanism of the *h*-MgO to rocksalt transformation in faceted CdSe nanocrystals.

Bulk CdSe transforms from the four-coordinate wurtzite to the six-coordinate rocksalt structure at a pressure of about 2.5 GPa [84]. In a recent paper, Shimojo *et al.* [15] studied this transformation in molecular dynamics computer simulation of a bulk crystal and identified two main

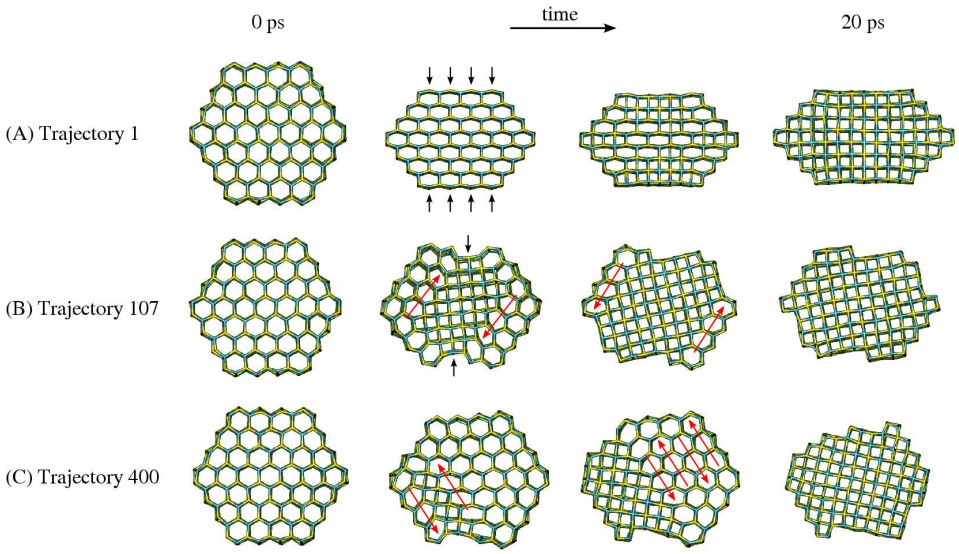

Figure 5.6: Evolution of the transformation mechanism during a transition path sampling simulation of a $Cd_{528}Se_{528}$ crystal. Rows (A), (B), and (C) show three selected trajectories, each evolving in time from left to right, showing the crystal along the hexagonal c-axis. In all three trajectories the crystal has the h-MgO structure at 0 ps; 20 ps later it is found in the rocksalt structure. (**A**) The simulation starts from an artificial trajectory, where the crystal is transformed through a uniform compression in the direction orthogonal to one set of parallel (100) surfaces, as indicated by the black arrows. Atoms at opposite vertices of hexagonal six-membered rings come together to form an additional bond. (**B**) After approximately 100 shooting moves, the mechanism changes: The first part of the transformation still follows a similar route, rearranging the mid-section of the crystal through the compression mechanism (black arrows). But then a different process begins, indicated by red arrows, which involves the sliding of parallel (100) planes. Note the different orientations of the final rocksalt lattice relative to the initial hexagonal lattice. (**C**) Yet a few hundred shooting moves later, the first mechanism has completely vanished and the crystal is transformed *via* the pure sliding planes mechanism. Starting on the surface, parallel (100) planes sequentially slide to transform the crystal, with adjacent planes preferably sliding in alternate directions.

mechanisms. In the first one the crystal is compressed in one of the equivalent [100] directions: the corresponding parallel (100) planes are flattened out and additional bonds are formed between atoms across six-membered hexagonal rings. This particular transformation route was first proposed by Tolbert and Alivisatos [73] for CdSe nanocrystals and will be called compression mechanism in the following. The second mechanism involves the sequential sliding of parallel (100) planes, where the 120° hexagonal bond angle along the planes is transformed into the 90° cubic angle. As a precursor to both mechanisms, the puckered (001) wurtzite layers are flattened out and the hexagonal five-coordinated *h*-MgO structure is formed as a metastable intermediate. Using transition path sampling, Zahn, Grin, and Leoni [83] recently showed that in the bulk crystal near the thermodynamic transition pressure of 2.5 GPa the sliding planes mechanism is highly favored over the compression mechanism and the transition occurs *via* nucleation and growth.

In CdSe nanocrystals, the analogous transformation is strongly influenced by surface structure and size of the crystals [16, 97, 98, 104]. Recent molecular dynamics simulations revealed that while spherical crystals transform much as observed in the bulk, the five-coordinated *h*-MgO structure can be stabilized in faceted crystals due to favorable surface energetics [97, 101]. However, in both faceted and spherical crystals the transformation from the hexagonal to the cubic structure can proceed through either the compression or the sliding planes mechanism, or a mixture of both [16, 97, 98].

To clarify the role of the two different mechanisms near the thermodynamic transition pressure, we apply the transition path sampling algorithm presented in the preceding section to the *h*-MgO to rocksalt transformation in faceted CdSe nanocrystals. While this transformation is observed at pressures of about 5.5 GPa in conventional molecular dynamics simulations, we here set the pressure to the bulk value of 2.5 GPa. We use the empirical pair potential for CdSe developed by Rabani [85]. Faceted $Cd_{216}Se_{216}$, $Cd_{528}Se_{528}$, and $Cd_{1050}Se_{1050}$ crystals with aspect ratios of about 1.0 are prepared by cleaving a bulk *h*-MgO lattice along equivalent (100) planes, with (001) and $(00\bar{1})$ planes terminating the crystals along the hexagonal [001] c-axis. We use trajectories of 20 ps length and a time step of 2 fs. At a temperature of 300 K, 44000 gas particles are needed on average to apply a pressure of 2.5 GPa on a $Cd_{528}Se_{528}$ crystal.

The transformation is monitored by calculating the percentage of crystal atoms with six nearest neighbors. Atoms are defined to be nearest neighbors if they are closer than 3.3 Å, the location of the first minimum of the radial distribution function. We accept trajectories if the crystal has no six-coordinated atoms initially and 20 ps later has more than 28%, 45%, and 54% six-coordinated atoms for $Cd_{216}Se_{216}$, $Cd_{528}Se_{528}$, and $Cd_{1050}Se_{1050}$ crystals, respectively. These values were determined from straightforward simulations in the rocksalt structure.

A first trajectory, necessary for starting a transition path sampling simulation, is obtained in the following way: Positions $\mathbf{q}_0$ of crystal atoms are initialized according to $\mathbf{q}_0 = (\mathbf{q}_{\mathrm{mgo}} + \mathbf{q}_{\mathrm{rs}})/2$, where $\mathbf{q}_{\mathrm{mgo}}$ represents the perfect h-MgO crystal and $\mathbf{q}_{\mathrm{rs}}$ are the positions of atoms in a rocksalt crystal obtained from a transformation via the compression mechanism. Particle velocities $\mathbf{p}_0$ are set to $\mathbf{p}_0 = (\mathbf{q}_{\mathrm{rs}} - \mathbf{q}_{\mathrm{mgo}})/\Delta t$ and then scaled to a temperature of 300 K. Keeping crystal atoms fixed, the atmosphere of ideal gas particles around the crystal is allowed to equilibrate for 10 ps. Starting from this point, a first trajectory is obtained by propagating the system forward and backward in time, resulting in a transformation of the crystal *via* the compression mechanism.

Crystals of all three sizes qualitatively show the same behavior. A typical transition path sampling run for a $Cd_{528}Se_{528}$ is illustrated in Fig. 5.6. Starting with the pure compression mechanism involving a uniform motion of all crystal atoms, within a few hundred trajectories the transformation evolves towards the sliding planes mechanism, nucleating on the surface of the crystal and growing along parallel (100) planes. In the thousands of subsequent trajectories we harvested this mechanism persists: While the relative directions of sliding planes can vary, resulting in different shapes of the final rocksalt crystal, the mechanism does not change qualitatively. This result strongly suggests that near the thermodynamic transition pressure the sliding planes mechanism is the most favorable transition route.

Chapter 6

Transition state analysis of solid-solid transformations in nanocrystals

"It is nice to know that the computer understands the problem. But I would like to understand it too." Eugene Wigner

A systematic simulation methodology is introduced which allows the accurate determination of experimentally measurable quantities characterizing solid-solid phase transformations under pressure. The atomistic mechanisms of nucleation and growth in a structural transformation of pressurized CdSe nanocrystals are identified using transition path sampling computer simulation. A committor-based transition state analysis is applied to extract activation enthalpies and volumes from transformation pathways at experimental conditions. The qualitative dependence of activation enthalpies on nanocrystal size is in good agreement with experimental data and supports the observed nucleation mechanism, which is characterized by a critical nucleus of remarkable shape located on the crystal surface. Based on committor distributions along typical transformation pathways, the coordination number is identified as a suitable reaction coordinate for the process.

6.1 Introduction

Identifying the atomistic mechanisms of structural phase transformations of solids on the nanoscale is an important step both for a deeper understanding of the physical process itself, as well as for the synthesis of new materials with tailored physical and chemical properties. Size-dependent behavior, typical for the nanoscale, can often be successfully described by simple thermodynamic considerations. For instance, the strong size dependence of the transition pressure that characterizes pressure-induced transfor-

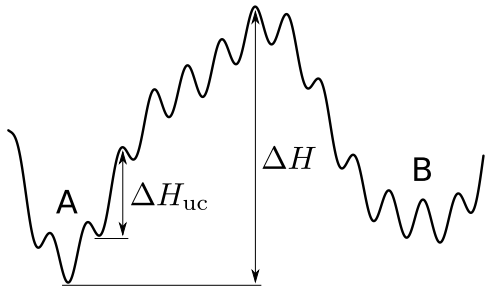

Figure 6.1: A rugged energy barrier separating two metastable conformations A and B of a solid. From the knowledge of the barrier ΔH_{uc} on the level of the unit cell alone it is not straightforward to reconstruct the activation energy ΔH of the entire transformation.

mations can be rationalized by models based on surface free energies [1, 8]. However, the specific surface structure and shape of a transformed particle and thus its surface free energy strongly depend on available atomistic transition routes [5]. In materials design, in turn, the facilitation or inhibition of such routes by precise control of surface configuration and shape of the nanoparticle can be a practicable way of stabilizing structures that are unstable otherwise [12, 105–107].

The thermodynamic and kinetic aspects of nanocrystal transformations have been studied in detail in pressure-experiments [2, 11, 12, 75, 88, 108]. Most significantly, Alivisatos and coworkers have demonstrated that it is possible to measure activation enthalpies and activation volumes for such processes [9, 10]. Based on this experimental evidence it is possible to speculate about the underlying transformation routes; for an unambiguous identification of the mechanism on the atomistic level, however, experimental time and space resolution currently remains insufficient.

The necessary atomistic perspective is conveniently provided by molecular dynamics computer simulation [13–16, 80, 83, 90–92, 98, 104, 109–114]. Nevertheless, the significance of a transformation mechanism observed with a computer simulation often remains unclear for two reasons. First, at experimental conditions of temperature and pressure, solid-solid transformations typically involve high kinetic barriers associated with nucleation events. The resulting waiting times often exceed seconds or even hours—timescales vastly out of range for atomistic computer simulations. The usual solution to this problem is to simulate at elevated pressure and temperature, whereupon transformations will readily occur on the accessible picosecond timescale. Under such conditions, however, observed mechanisms will in general not be comparable to mechanisms observed experimentally.

Second, a direct comparison between quantities that are observable both

in experiments and simulation is not easily achieved. Activation enthalpies and volumes, as measured in experiments [9, 10], in principle provide an excellent indicator of mechanism because the height of the kinetic barrier is a direct consequence of the specific atomistic pathways. Calculation of such quantities, however, require the identification of the central part of the transformation, the transition state containing the critical nucleus of the high pressure phase. Moreover, a full modeling of the experimental situation is necessary. In particular, it is not sufficient to identify the transformation pattern with the lowest energy barrier on the level of the unit cell, as illustrated in Fig. 6.1, because the activation enthalpy for the entire transformation sensitively depends on the precise sequence of such steps. Moreover, in a nanocrystal, certain transformation routes might not be accessible at all because they lead to particularly unfavorable surface configurations.

Here, we present a systematic methodology that overcomes these problems and establishes direct contact with experiments through the calculation of activation enthalpies and volumes at conditions close to experiments. We demonstrate its usefulness by identifying the nucleation mechanism of a structural transformation in CdSe nanocrystals. The timescale problem is avoided by means of transition path sampling [17,18,32,65,115], a simulation method for the study of rare events in complex systems. Using the committor function, a statistical measure for the progress of the transformation, we identify the critical nuclei of the high pressure phase and calculate the corresponding activation enthalpies and volumes. A comparison with experimental values yields good qualitative agreement and supports observed mechanism observed in the simulation, which involves a nucleation event on the surface of the crystal and subsequent growth by sliding of parallel crystal planes. We also show how committor calculations can be used to assess the quality of putative reaction coordinates, which capture the essentials of the transformations.

Briefly, the application of our methodology to CdSe nanocrystals consists of the following steps:

1. We perform transition path sampling simulations of the h-MgO to rocksalt transformation at pressures of 2.0, 2.5, 3.0, 3.5, and 4.0 GPa for crystals with diameters of 2, 3, 4, and 5 nm at a temperature of 300 K. For every crystal size and value of the pressure, at least 10 000 trajectories are collected (Fig. 6.2).

2. For 5 different trajectories chosen randomly from the ensemble of reactive trajectories for given pressure and crystal size, we calculate the committor as a function of time (Fig. 6.3) and identify all transition states along the trajectory (Fig. 6.4).

3. Along the same trajectories, using static Monte Carlo simulations, we calculate the enthalpy and volume of the crystal as a function of

time (Fig. 6.10). The average enthalpy and volume of the crystal in the h-MgO structure at the beginning of the trajectory, as well as in the final rocksalt configuration at the end are also calculated in long Monte Carlo runs. This allows us to estimate the dependence of the thermodynamic transition pressure on crystal size (Fig. 6.9).

4. For every transition state along a given trajectory, the activation enthalpy and volume are determined by comparing enthalpy and volume at the transition state with the respective values of the h-MgO structure (Fig. 6.10). These values are then averaged over all transition states along the trajectory and over all trajectories for given crystal size and pressure. The size and pressure dependence of the activation enthalpy and volume are compared to experimental values in Fig. 6.11.

5. An analysis of committor distributions along typical transformation pathways allows us to identify a suitable reaction coordinate (Fig. 6.12).

The individual steps outlined above are discussed in depth in the following sections.

6.2 Transition path sampling

In a recent paper [97], we have used molecular dynamics simulations to study the atomistic mechanisms of the wurtzite to rocksalt transformation in CdSe nanocrystals, modeled with the empirical pair potential by Rabani [85, 89]. As pressure-transmitting medium, we introduced a barostat of ideal gas particles that occupy a stochastically updated atmosphere around the nanocrystal [101]. Depending on crystal size and surface structure, qualitatively different transformation mechanisms are possible. In perfectly spherical crystals with unordered surfaces, the transformation proceeds mainly through shuffling of (100) planes and a simultaneous unbuckling of (001) layers. In crystals that have the shape of hexagonal prisms with well-defined and stable low-energy facets, the mechanism decomposes into two steps: The unbuckling part happens at lower pressures, taking the crystal to the 5-coordinated h-MgO structure which in the second step is transformed to rocksalt at higher pressures through the shuffling of (100) planes. Apart from this main route, a third mechanism was observed in some cases which was previously proposed by Alivisatos and coworkers [8] and involves the flattening-out of (100) planes and corresponds to a pathway

The pressures needed to induce the transformations to rocksalt discussed above lie between 5.5 and 10 GPa, depending on crystal size and shape. While these pressures are only moderately higher than those used experimentally, the timescales differ by orders of magnitude. To observe the transformation in simulation, the pressure has to be increased close to the point

where the low-pressure phase becomes unstable. The resulting transformations proceed rather violently, with simultaneous nucleation events from different sites on the crystal and the formation of grain boundaries [90,91,97].

Simulations under experimental conditions can be conducted with transition path sampling, a method that uses a Monte Carlo procedure to sample the ensemble of trajectories that connect two well-defined regions A and B in phase space [18]. This restriction of the space of observable pathways is particularly useful when A and B are separated by a barrier that is large compared to typical thermal fluctuations and the simulation of barrier-crossing trajectories becomes increasingly time-consuming with conventional methods. In a transition path sampling simulation, a Markov chain of paths is constructed by means of a trajectory displacement move called *shooting* [18], and individual trajectories are accepted or rejected as members of the chain according to a detailed balance condition that ensures correct sampling. While a strong bias is thus applied in trajectory space, the underlying dynamics along individual trajectories remain untouched and true dynamical pathways are observed. In particular, the frequency of occurrence of different mechanisms reflects the one observable in a hypothetical, extremely long conventional simulation run.

Recently, we have combined the ideal gas barostat used in our regular molecular dynamics simulations with transition path sampling to study the *h*-MgO to rocksalt transformation in faceted CdSe nanocrystals at a pressure of 2.5 GPa [65]. Sampling was started from artificially created first trajectories featuring a transformation through the mechanism proposed by Alivisatos and coworkers [8]. Within a few hundred iterations of the algorithm, the trajectories converged to the sliding-planes mechanism mentioned above, clearly identifying this route as the most favored one at lower pressures. Here, as the first step towards the identification of transition states in CdSe nanocrystals, we use the same simulation setup to study the transformations of crystals with different sizes (2, 3, 4, and 5 nm diameter—120, 432, 1056, and 2100 atoms) at pressures ranging from 2 to 4 GPa. The length of trajectories varies between 20 and 60 ps, depending on crystal size and pressure. Other simulation details, in particular regarding the implementation of the ideal gas barostat, can be found in Ref. [65].

6.2.1 Models of CdSe nanocrystals

The shape of CdSe nanocrystals used in pressure experiments has been determined with electron microscopy, revealing that the crystals have well-defined facets dominated by (001) and (100) surfaces [87] and an aspect ratio of about 1.2 in the wurtzite structure. While the hexagonal prismatic crystals used in our simulations are built according to these specifications, in real crystals this perfect geometry will be degraded by steps and other surface defects. Also, the role of faults in the perfect wurtzite stacking of

(001) planes observed experimentally is neglected in our model. Indeed, if the nanocrystal has surface defects or stacking faults, the *h*-MgO structure is not observed as a stable intermediate [97]. Nevertheless, we use the *h*-MgO to rocksalt transformation in these crystals as a model for the full wurtzite to rocksalt transformation. We do this because in the presence of surface defects transformations can get stuck in metastable intermediate states, resulting in significantly longer trajectories and therewith increased computational effort. Still, the simplified model used here is justified because of two reasons. First, the *h*-MgO to rocksalt transformation is the rate-limiting step, clearly occurring at higher pressures than the step from wurtzite to *h*-MgO. Second, also in simulations using other models [16, 98], a significant compression along the c-axis of the hexagonal wurtzite lattice was observed with increasing pressure; at the time of the transformation, the crystal is essentially in the *h*-MgO structure. An *a posteriori* justification of our approach is provided by the good agreement with experimental data, as discussed in section 6.5.

6.2.2 Defining the initial and final states

In a transition path sampling simulation, the initial and final states A and B, here the *h*-MgO and rocksalt structures, need to be well-defined regions in configuration space (or, more generally, phase space). These definitions should be chosen with care, as they must provide both a clear distinction between A and B and incorporate most of the equilibrium fluctuations within the states as well as all possible isomorphs. Often, this can be achieved with a suitable order parameter. In our case, a natural choice for such a parameter is the number of atoms with six nearest neighbors. Atoms are assumed to be nearest neighbors if their distance is smaller than 3.4 Å, a value roughly equal to the location of the first minimum of the radial distribution function. For crystals with diameters of 2, 3, 4, and 5 nm we use increasing thresholds of 8%, 20%, 30%, and 40% six-coordinated atoms, reflecting the decreasing surface-to-volume ratio, to define the rocksalt structure. These values, as will be confirmed later by committor analysis, lie safely beyond the values at the transition state but still allow the formation of all relevant isomorphs in the rocksalt structure. Preliminary simulations showed, however, that the use of the same order parameter for the definition of the *h*-MgO structure is problematic: Values that are too low do not include all equilibrium fluctuations within A, while with larger values at higher pressures critical rocksalt nuclei are frequently assumed to lie in A—correct sampling of trajectories is corrupted. We therefore use a bond angle criterion to define A: The angles between nearest neighbors atoms that have a value of 120° in the perfect *h*-MgO structure are monitored and the crystal is assumed to be in state A only when none of these angles exceeds the ideal value by more than 35°. While this definition excludes all other possible isomorphs of the crystal in

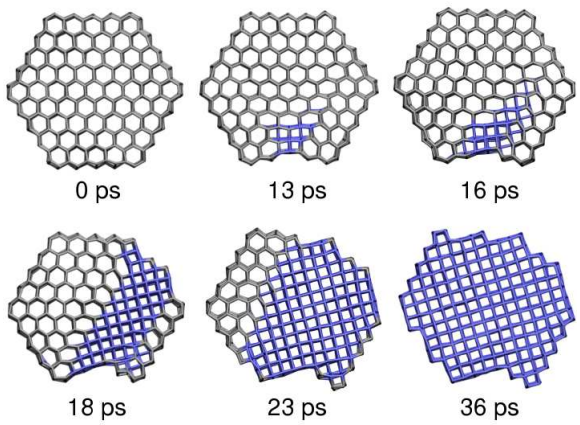

Figure 6.2: Snapshots along a typical transformation pathway of a 5 nm nanocrystal at a pressure of 3 GPa, viewed along the hexagonal c-axis. Gray atoms are in the hexagonal structure, blue atoms in the rocksalt structure. At 13 ps the crystal has reached the transition state (compare this configuration to the one in Fig. 6.4). Depending on the nucleation site and on the direction in which particular planes slide, different morphologies of the final rocksalt crystal are obtained.

the h-MgO structure, we confirmed that no trajectories are rejected because they lead to such incompatible morphologies. This indicates that the hexagonal prismatic shape, which is also observed experimentally [87], provides a strong basin of attraction for the sampled trajectories due to particularly favorable surface free energies.

6.3 Identifying the transition states

After convergence of the mechanism has been achieved in the transition path sampling simulations, we collect at least 10 000 trajectories; from these, we randomly select five for further analysis. A typical converged transformation pathway, displaying the mechanism of sliding (100) planes, is shown in Fig. 6.2. To identify the transition states, we calculate the committor p_{B} for all of the 1000 stored configurations along the paths. The committor is a statistical measure for the progress of the transformation and is defined as the probability of a particular configuration of the crystal to relax into state B, the rocksalt structure, rather than into the hexagonal structure. Thus, as the transformation proceeds, the committor changes its value from 0 to 1 and a configuration for which $p_{\mathrm{B}} = 1/2$ is identified as a transition state [18].

We estimate p_{B} using the computational scheme described in detail in

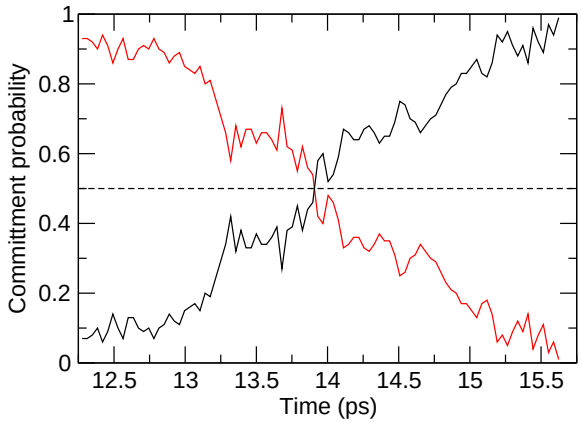

Figure 6.3: Committor p_B (black) and $p_A = 1 - p_B$ (red) along a transition path of a 5 nm nanocrystal, calculated with the scheme described in the text. The point around 13.9 ps where $p_B = p_A = 0.5$ marks the transition state.

Ref. [18], which ensures a consistent level of statistical error. From every configuration of the crystal, short molecular dynamics trajectories are started, each with particle velocities drawn from the appropriate Maxwell–Boltzmann distribution and a freshly created ideal gas atmosphere. When a trajectory reaches either stable state, it is terminated; the fraction of those that reach the rocksalt structure is an estimate for p_B with an Gaussian uncertainty of

$$\sigma = \sqrt{p_B(1 - p_B)/M}\ , \tag{6.1}$$

where M is the number of initiated trajectories. To ensure good statistics, a minimum of 10 trajectories are performed. Further trajectories are initiated until the configuration at hand can be excluded from the ensemble of transition states with an certainty of 95%, i. e. , when the estimate for p_B falls short of the interval $[1/2-2\sigma; 1/2+2\sigma]$. Configurations that live to see a maximum of 100 trajectories are assumed to be transition states. Figure 6.3 shows the time evolution of the committor for a barrier-crossing event in a typical trajectory. As the crystal crosses the barrier top, repeated recrossings of the $p_B = 1/2$ line are frequently observed and up to 25 transition states are identified along a single trajectory.

A compilation of transition states, identified through the committor criterion, are shown in Fig. 6.4. The shape of the critical rocksalt nuclei is surprising: Regardless of pressure or crystal size, the nuclei always appear on the surface, extending along the hexagonal c-axis over the full side length of the crystal. The transformations begin with the indentation of an (100) side facet, preferably in the facet's center. The resulting quasi-one-dimensional linear arrangement of atoms, each forming an additional bond with an op-

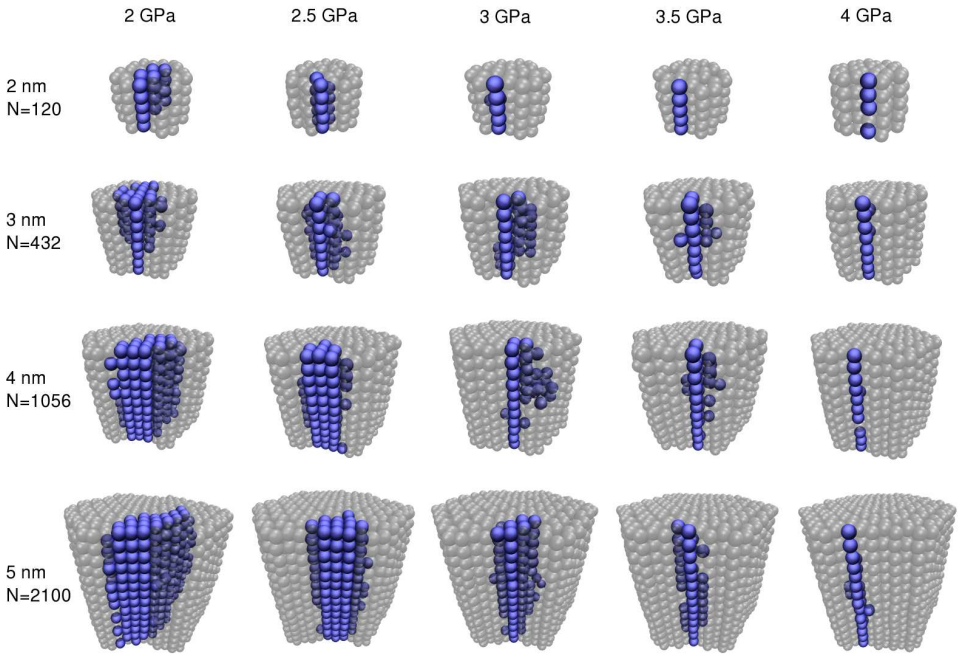

Figure 6.4: Critical rocksalt nuclei for a series of pressures and crystal sizes. Transparent gray atoms are in the hexagonal structure, blue ones in the rocksalt structure. Crystals have been rotated such that the nuclei appear on the same crystal facet. Mind that separate transition path sampling simulations were conducted for every crystal size and pressure; there is no dynamical connection of configurations at different pressures for fixed crystal size.

posite atom of the next (001) plane, forms the critical rocksalt nucleus at high pressures. At lower pressures, the transition state is reached after the rocksalt slab has grown in one of two equivalent [100] directions and along the lateral surface. This increased size of the critical nucleus at lower pressures is consistent with classical nucleation theory: Close to the thermodynamic transition pressure, the rocksalt structure is favored only slightly over the hexagonal structure. As a consequence, a larger amount of rocksalt is needed to pay the price of the interface than at higher pressures, at which the rocksalt structure is clearly more stable.

The differences of the observed nucleation mechanism to equivalent transformations in the bulk material are evident. In the bulk, nucleation will preferably happen at lattice defects like grain boundaries that locally facilitate the accommodation of an interface between the low and high pressure structures [9, 116]. In nanocrystals in the size range considered here lattice defects are easily annealed out during synthesis; nucleation from such sites thus is not a possibility. As a consequence, the surface serves as a low energy nucleation site. We speculate that the free energy penalty due to the large surface of the elongated nucleus is compensated by favorable strain energetics compared to a nucleus of more spherical shape located on the crystal's inside.

6.4 Measuring enthalpy and volume

To link our results to experiments, we quantify the observed rocksalt nuclei by calculating their activation enthalpy and activation volume, as discussed in Sects. 6.4.4 and 6.4.3. We do this using an efficient Monte Carlo version of the ideal gas pressure bath, which is discussed in Sec. 6.4.1. Experimentally, these quantities can be determined by evaluating the temperature and pressure dependence of the rate constant [9, 10]; in Sec. 6.4.2 we discuss thermodynamically meaningful definitions of enthalpy and volume that can be consistently compared to experiment.

6.4.1 A Monte Carlo algorithm for the ideal gas barostat

For the study of atomistic transformation mechanisms with computer simulation, molecular dynamics is the natural choice. In previous work we have shown how a barostat of ideal gas particles can be implemented efficiently within a molecular dynamics simulation such that it conserves the probability distribution of a system at constant pressure and temperature [65, 101]. However, for the calculation of static quantities like the average energy a Monte Carlo version of the barostat might be advantageous. In this section we present such a Monte Carlo algorithm and show that it obeys detailed balance.

The system we consider consists of a nanocrystal of N_c atoms immersed in a pressure bath that adapts to the shape and size of the crystal and consists of N non-interacting gas particles. These gas particles interact with crystal atoms via a purely repulsive potential that depends only on the particle distance r,

$$u(r) = \begin{cases} \epsilon \left[(r/\sigma)^{-12} - (r_\mathrm{cut}/\sigma)^{-12} \right] & \text{if } r < r_\mathrm{cut} \,, \\ 0 & \text{if } r \geq r_\mathrm{cut} \,. \end{cases} \tag{6.2}$$

In our simulations, we use the parameters $\epsilon = 1\,\mathrm{kJ/mol}$, $\sigma = 3.0\,\text{Å}$, and $r_\mathrm{cut} = 6.0\,\text{Å}$, choices that prevent gas particles from entering the crystal. The simulation box is divided into cubic cells with side length $l = r_\mathrm{cut} + 0.1\,\text{Å}$. Thus, all gas particles that could possibly interact with crystal atoms in a given cell will be found in the cell itself and its nearest neighbor cells. (The reason for not simply setting $l = r_\mathrm{cut}$ will become evident later in the discussion.) That given, we define the system volume to be a function of the coordinates of the crystal atoms: At any time, the accessible volume V of the system consists of all cells that contain crystal atoms and all cells adjacent to those. The probability distribution of this system is given by

$$\rho(x) \propto \mathrm{e}^{-\beta(U+PV-\mu N)} \,, \tag{6.3}$$

where x denotes the state of the system, specified by the positions of all particles, $U = U^{(\mathrm{cc})} + U^{(\mathrm{cg})}$ is the total energy of the system with contributions $U^{(\mathrm{cc})}$ and $U^{(\mathrm{cg})}$ from crystal-crystal and crystal-gas interactions, respectively, P is the pressure, μ is the chemical potential of the ideal gas, and $\beta = 1/k_\mathrm{B}T$ is the reciprocal temperature. To ensure that this probability distribution is conserved in our simulation, it is sufficient that the Monte Carlo moves used to evolve the system from an "old" state o to a "new" state n obey the detailed balance condition [65]

$$\frac{p(o \to n)}{p(n \to o)} = \frac{N_o!}{N_n!} \mathrm{e}^{-\beta[U_n - U_o + P(V_n - V_o) - \mu(N_n - N_o)]} \,. \tag{6.4}$$

Here, $p(o \to n)$ is the transition probability to go from state o to state n and the subscripts indicate respective quantities in these states. As discussed in Ref. [65], the prefactor accounts for all possible computer representations of a given physical state.

To sample the distribution given by Eq. (6.3), we consider two different Monte Carlo moves. The first one is a standard displacement move of a single crystal atom, with all other particles fixed. A crystal particle is chosen at random and displaced by a small vector, whose components are drawn from a symmetric distribution. For such a move, $p(o \to n)$ is the product of the probability $\pi(o \to n)$ to try the move and the probability $\alpha(o \to n)$ to accept it,

$$p(o \to n) = \pi(o \to n)\, \alpha(o \to n) \,. \tag{6.5}$$

As the system volume and number of gas particles are left unchanged, detailed balance is satisfied by the standard Metropolis acceptance probability [41], which takes into account the total change in energy,

$$\alpha(o \to n) = \min\left(1, \mathrm{e}^{-\beta(U_n - U_o)}\right) . \tag{6.6}$$

The second move affects the ideal gas atmosphere. Instead of displacing single gas particles, we want to resample the whole atmosphere in one step. Such a global move greatly enhances the efficiency of the calculation of static averages, because successive realizations of the ideal gas atmosphere are uncorrelated; in a molecular dynamics simulation, decorrelation will be typically reached after a few picoseconds. To resample the gas atmosphere, we proceed in the following way:

1. Remove all gas particles from the simulation.

2. Determine the system volume V_n, that is, the total volume of all cells that contain crystal atoms and their next-neighbor cells.

3. Fill the accessible volume with new gas particles. To this end, perform k trial insertion moves, where k is a number drawn from a Poisson distribution
$$\mathcal{P}(k) = \frac{\bar{k}^k \mathrm{e}^{-\bar{k}}}{k!} , \tag{6.7}$$
with an average of
$$\bar{k} = \beta P V_n . \tag{6.8}$$
One such insertion move consists of the following steps: Choose the position $\mathbf{r}$ of the gas particle uniformly distributed over the total available volume; compute the interaction energy $E(\mathbf{r})$ of the inserted gas particle with all crystal atoms; accept the insertion move with probability $\mathrm{e}^{-\beta E(\mathbf{r})}$.

This procedure is on average performed every N_c crystal particle displacement moves.

We now show that the above algorithm satisfies detailed balance. Because of a possibly large interaction energy with crystal atoms, in general not all of the k inserted particles will "survive" the acceptance criterion. As a result, there are many possible realizations of the move $o \to n$ leading to a final number of N_n accepted gas particles at positions corresponding to state n [65]. Thus, $p(o \to n)$ is the sum of the probabilities of all such realizations,

$$p(o \to n) = \sum_{k=N_n}^{\infty} \mathcal{P}(k) \binom{k}{N_n} \frac{\mathrm{e}^{-\beta U_n^{(\mathrm{cg})}} (1 - \Pi_n)^{k-N_n}}{V_n^{N_n}} , \tag{6.9}$$

where $1/V_n^{N_n}$ is the probability to put N_n particles in the right places and $\mathrm{e}^{-\beta U_n^{(\mathrm{cg})}}$ is the probability to accept them there. The factor Π_n denotes the probability that any one insertion move will be accepted, which depends on the positions of the crystal atoms and on the available volume V_n,

$$\Pi_n = \frac{1}{V_n} \int_{V_n} \mathrm{d}\mathbf{r}\, \mathrm{e}^{-\beta E(\mathbf{r})} \; . \tag{6.10}$$

The factor $(1-\Pi_n)^{k-N_n}$ therefore gives the probability that $k-N_n$ of the total k insertion moves will be rejected. Inserting Eqs. (6.7) and (6.8) into Eq. (6.9) we have, after some manipulation,

$$\begin{aligned} p(o \to n) &= \frac{(\beta P V_n)^{N_n}}{N_n! V_n^{N_n}} \mathrm{e}^{-\beta(U_n^{(\mathrm{cg})} + P V_n \Pi_n)} \\ &= \frac{1}{N_n!} \mathrm{e}^{-\beta(U_n^{(\mathrm{cg})} + P V_n \Pi_n - \mu N_n)} \; , \end{aligned} \tag{6.11}$$

where in the last step we have identified the fugacity of the ideal gas, $\mathrm{e}^{\beta\mu} = \beta P$. The ratio of transition probabilities of the forward and backward moves now becomes

$$\frac{p(o \to n)}{p(n \to o)} = \frac{N_o!}{N_n!}\, \mathrm{e}^{-\beta[U_n^{(\mathrm{cg})} - U_o^{(\mathrm{cg})} + P(V_n \Pi_n - V_o \Pi_o) - \mu(N_n - N_o)]} \; . \tag{6.12}$$

As the move affects only gas particles, $U_n^{(\mathrm{cg})} - U_o^{(\mathrm{cg})} = U_n - U_o$, and this expression is identical to the detailed balance condition of Eq. (6.4), except for the volume terms. However, if no new cells have been added to or removed from the simulation in step 2 of the algorithm, $V_n = V_o$ and $\Pi_n = \Pi_o$, and detailed balance is obeyed. In case the system volume has to be changed because since the last update of the gas atmosphere crystal atoms have been moved into or out of cells that do not hold any other crystal atoms, the respective volume terms do not cancel out. In this case, we can make use of the fact that gas particles in the cells that are removed or added do not interact with crystal atoms. (This is true only if the side length of the cells is larger than the cutoff distance of the gas-crystal interaction, and the number of crystal atom displacements between successive updates of the gas atmosphere as well as the maximum size of the particle displacement are not too large.) Then,

$$\begin{aligned} \Pi_n &= \frac{1}{V_n} \left(\int_{V_o} \mathrm{d}\mathbf{r}\, \mathrm{e}^{-\beta E(\mathbf{r})} - \int_{V_o \setminus V_n} \mathrm{d}\mathbf{r} \; + \int_{V_n \setminus V_o} \mathrm{d}\mathbf{r} \right) \\ &= \frac{1}{V_n} (V_o \Pi_o - V^- + V^+) \; , \end{aligned} \tag{6.13}$$

where we have introduced the symbols V^- and V^+ for volume that is removed and added, respectively. Using this relation, it follows that

$$V_n \Pi_n - V_o \Pi_o = V^+ - V^- = V_n - V_o \; , \tag{6.14}$$

Also in this case, detailed balance is thus obeyed. The last two equations reflect the fact that although the volume of the total system may change, the volume of the crystal is constant during the move. Thus, a possible change of the total volume is equaled by the change of the volume available to gas particles. In the next section we will indeed argue that $V(1-\Pi)$ and $V\Pi$ are correct measures for the crystal and gas volume, respectively.

6.4.2 Definitions of crystal volume and enthalpy

For the compound system consisting of pressure bath and nanocrystal, all thermodynamic properties are well defined. For instance, the total volume is determined by the positions of crystal atoms (see Section 6.4.1) and the total potential energy U is simply the sum of all pair-interaction energies. When we are interested in a specific property of the nanocrystal alone, for instance its volume V_c, energy U_c, or entropy S_c, a subdivision of the respective quantity of the total system between crystal and gas has to be made. To allow for comparison with experimentally determined quantities, it is essential that the introduced definitions are consistent with the quantities that enter the fundamental thermodynamic function, the Gibbs free energy of the nanocrystal

$$G_\mathrm{c} = U_\mathrm{c} + PV_\mathrm{c} - TS_\mathrm{c} \; . \tag{6.15}$$

In computer simulations of nanoparticles, several different methods are in use in which the volume is a function of the coordinates of the cluster atoms alone [100]; examples include volume definitions based on the volume of overlapping spheres [117] and convex hull constructions [118]. For a nanocrystal in explicit solvent, these definitions are inappropriate because they neglect possibly complicated effects of the solvent-crystal interaction. The correct thermodynamics of small systems in solution (and its connection to statistical mechanics) has been worked out long ago by Hill in the context of constant-pressure solution theory [119, 120]. The correct expression for the Gibbs free energy of a single small system in a solvent is

$$G_\mathrm{c} = -\beta^{-1} \ln \frac{\Delta(N_1, 1, P, T)}{N_1 \Delta(N_1, 0, P, T)} \; , \tag{6.16}$$

where

$$\Delta(N_1, N_2, P, T) = \sum_V Q(N_1, N_2, V, T) \mathrm{e}^{-\beta PV} \; . \tag{6.17}$$

Here, Δ and Q are the isobaric-isothermal and canonical partition functions, respectively, of an infinitely dilute system of N_1 solvent molecules and N_2 solute particles (nanocrystals, in our case). The thermodynamic properties of interest can be obtained by taking derivatives of Eq. (6.16). For the

volume we have

$$\begin{aligned} V_{\mathrm{c}} = \frac{\partial G_{\mathrm{c}}}{\partial P} &= \frac{\partial}{\partial P}\left\{-\beta^{-1}\ln\Delta(N_1,1,P,T)\right\} - \frac{\partial}{\partial P}\left\{-\beta^{-1}\ln\Delta(N_1,0,P,T)\right\} \\ &= \frac{\partial G(N_2=1)}{\partial P} - \frac{\partial G(N_2=0)}{\partial P} \\ &= V(N_2=1) - V(N_2=0)\ . \end{aligned} \tag{6.18}$$

Here, $G(N_2)$ and $V(N_2)$ denote the Gibbs free energy and volume, respectively, of a system with N_2 solute particles at constant N_1, P, and T. In the context of our simulations, $V(N_2 = 1)$ is simply the total volume of the compound system, and, given that the solvent species is an ideal gas, $V(N_2 = 0)$ is obtained from the ideal gas law,

$$V_{\mathrm{c}} = \langle V\rangle - \langle N\rangle(\beta P)^{-1}\ . \tag{6.19}$$

As a macroscopic analogy, consider a swimming pool completely filled with water. It follows from the above definition that the amount of water that is displaced by a dog jumping into the pool is a measure for the dog's volume. Similarly, we get for the enthalpy and energy of the crystal.

$$\begin{aligned} H_{\mathrm{c}} &= \frac{\partial(\beta G_{\mathrm{c}})}{\partial\beta} = U_{\mathrm{c}} + PV_{\mathrm{c}} \\ U_{\mathrm{c}} &= U^{(\mathrm{cc})} + U^{(\mathrm{cg})} = U\ , \end{aligned} \tag{6.20}$$

as might have been anticipated.

To calculate the average volume of the crystal in one of the stable structures, it is necessary to average V and N over many realizations of the pressure bath, as indicated by the angular brackets in Eq. (6.19). As we are also interested in the time evolution of the crystal volume during a transformation, we measure the instantaneous volume by averaging over pressure bath realizations with fixed positions of crystal atoms,

$$V_{\mathrm{c}}^{\mathrm{inst}} = V - \langle N\rangle_V(\beta P)^{-1}\ . \tag{6.21}$$

Here, V is the (fixed) total volume of the system and $\langle N\rangle_V$ is the average number of gas particles in the available volume. Using the algorithm discussed in Sec. 6.4.1, $\langle N\rangle_V$ is equal to the average number of trial insertions, Eq. (6.8), times the acceptance probability of a single insertion Π, $\langle N\rangle_V = \beta PV\Pi$. Inserting this relation into Eq. (6.21) yields $V_{\mathrm{c}}^{\mathrm{inst}} = V(1-\Pi)$, or alternatively for the volume of the gas, $V_{\mathrm{g}}^{\mathrm{inst}} := V - V_{\mathrm{c}}^{\mathrm{inst}} = V\Pi$, which provides the connection with the detailed balance condition of Eq. (6.12).

As a consequence of the definitions discussed above, thermodynamic properties like the volume and enthalpy of a nanocrystal explicitely depend on the details of the interaction with the pressure medium, a fact that can

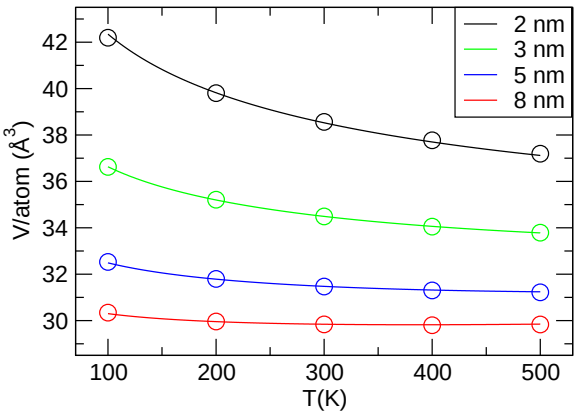

Figure 6.5: Volume per atom as a function of temperature for four different crystal sizes in the h-MgO structure at a pressure of 2 GPa. Lines represent fits to a model explaining the volume of the crystal in terms of bulk and surface contributions, as discussed in the text. With increasing temperature, the average distance between gas particles and crystal atoms on the surface decreases. For the nanocrystal sizes considered here, this surface effect dominates the bulk thermal expansion.

give rise to unexpected behavior that is not accounted for by volume definitions based on positions of crystal atoms alone. For instance, a convex hull approximation to the volume can be expected to underestimate the thermodynamic volume of the crystal because it neglects the "surface volume" associated with the space between the crystal surface and the first layer of gas particles. The thickness of this void depends on the crystal-gas interaction and will decrease with increasing temperature, as gas particle more closely approach the crystal surface. This effect can obscure the bulk behavior of thermal expansion and lead to an overall decrease of the volume with increasing temperature, as illustrated in Fig. 6.5.

As a simple model for the surface volume is illustrated in Fig. 6.6. Let us assume that, on average, gas particles can approach the crystal surface only till their potential energy is of order $k_{\mathrm{B}}T$. Using Eq. (6.2), the average distance to the surface is then approximately given by

$$d(T) = \sigma(k_{\mathrm{B}}T)^{-1/12} \,. \tag{6.22}$$

Assuming spherical shape of the nanoparticle, the volume can be written as

$$V_{\mathrm{c}} = \frac{4\pi}{3}\big(r_0\,\mathrm{e}^{\alpha T} + \gamma\, d(T)\big)^3 \,, \tag{6.23}$$

where r_0 is the radius of the crystal, α is the bulk coefficient of thermal expansion and the factor γ represents deviations from perfect spherical shape.

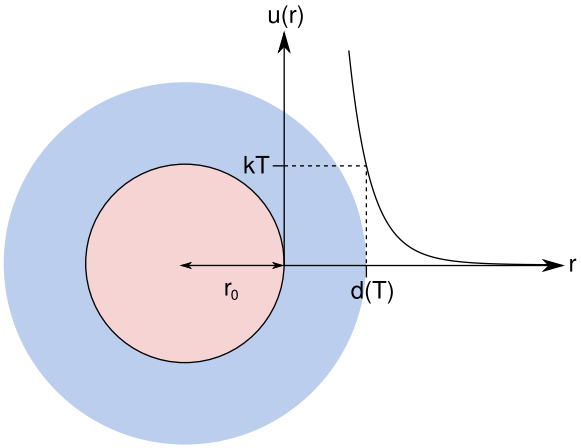

Figure 6.6: A simple model for the volume of a spherical nanocrystal in a pressure bath. Surface volume (blue) is added to the volume of the bulk core (red) as the repulsive interaction $u(r)$ between crystal and pressure bath particles produces a void shell of thickness $d(T)$.

In Fig. 6.5, this function is used for fitting to the volume data with common values of $\alpha = 2.2 \times 10^{-5}\,\mathrm{K}^{-1}$ and $\gamma = 1.35$ for all crystal sizes.

We would like to emphasize that the particular definitions of volume and enthalpy, Eqs. (6.19) and (6.20), are not subject to choice but need to be applied for the standard relations of thermodynamcis to stay valid [120]. To illustrate this point, we implement an alternative measure of the volume, based on hard sphere radii of crystal and gas particles. To determine this alternative volume $\hat{V}_\mathrm{c}$ of the crystal, only insertion moves of pas particles within distances larger than d_0 of any crystal atom are counted as successful. While this definition effects the volume and enthalpy of the crystal, the energetics of the pressure bath still correspond to Eq. (6.2). We therefore have easy access to a measure of the volume which does not depend on temperature and is purely geometric, comparable, for instance, to a convex hull approximation. Using both definitions, we calculate the equations of state $V_\mathrm{c}(P,T)$, $H_\mathrm{c}(P,T)$, $\hat{V}_\mathrm{c}(P,T)$, and $\hat{H}_\mathrm{c}(P,T) = U_\mathrm{c}(P,T) + P\hat{V}_\mathrm{c}(P,T)$ of a 2 nm crystal in the rocksalt structure for a broad range of pressures and temperatures; the particular value of $d_0 = 3.32$ Å is chosen to give agreement of both volume definitions at $T = 100$ K and $P = 2.5$ GPa. Polynomial fits to the resulting volume data are compared in Fig. 6.7.

In standard thermodynamics, useful interrelations between the derivatives of these equations can be derived [121]. For instance, using $\mathrm{d}H = T\mathrm{d}S + V\mathrm{d}P$ and the Maxwell relation $\left(\frac{\partial S}{\partial P}\right)_T = -\left(\frac{\partial V}{\partial T}\right)_P$, we see that

$$\left(\frac{\partial H}{\partial P}\right)_T = T\left(\frac{\partial S}{\partial P}\right)_T + V = -T\left(\frac{\partial V}{\partial T}\right)_P + V\,. \tag{6.24}$$

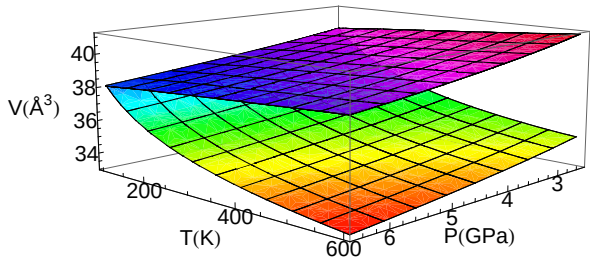

Figure 6.7: Volume per atom as a function of temperature and pressure for the thermodynamically correct definition of the volume V_c (lower surface) and the alternative $\hat{V}_c$ (upper surface), which neglects volume effects of the crystal-bath interaction. The surfaces are fits to data from Monte Carlo simulations of a 2 nm crystal in the rocksalt structure. While V_c decreases with increasing temperature (compare Fig. 6.6), $\hat{V}_c$ shows bulk-like behavior. Note that $\hat{V}_c$ is parametrized to give agreement with V_c at $T = 100$ K.

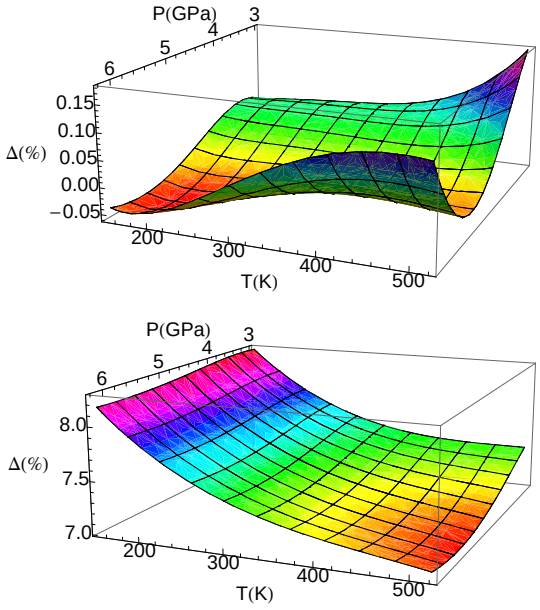

Figure 6.8: Relative difference Δ of the left and right hand sides of Eq. (6.24) for data calculated with the correct definition of the crystal volume V_c (top) and the alternative $\hat{V}_c$ (bottom). While deviations for the first definition are small and consistent with the error in the fitted functions and their derivatives, a large systematic error is observed for the latter definition.

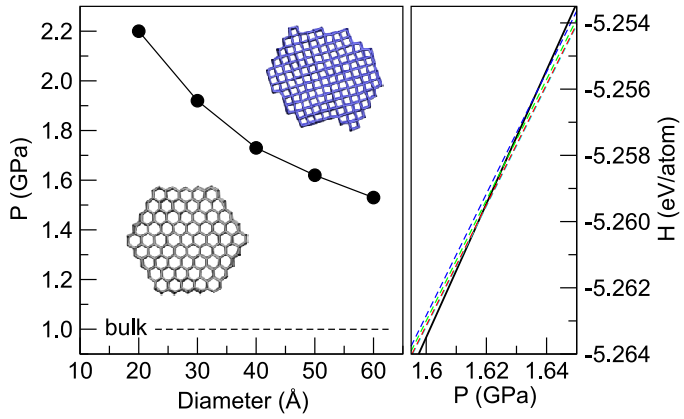

Figure 6.9: Left: Thermodynamic transition pressures of the h-MgO to rocksalt transformation as a function of crystal size, estimated from points of equal enthalpy. The dashed line indicates the bulk transition pressure of about 1.0 GPa [97]. Right: Enthalpy-pressure curves for a 5 nm crystal in the h-MgO structure (solid line) and in 4 different rocksalt configurations (dashed lines) obtained with transition path sampling.

This equation is equivalent to the well-known relation

$$c_V = c_P - TV\alpha^2/\kappa_T \tag{6.25}$$

between the heat capacities at constant volume and pressure c_V and c_P, respectively, the coefficient of thermal expansion α, and the isothermal compressibility κ_T. In Fig. 6.8, we check the validity of this relation for our nanocrystal system by comparing the right and left hand sides of Eq. (6.24) for both definitions of the crystal volume; good agreement is only achieved for V_c.

6.4.3 Size dependent transition pressure

A typical phenomenon in phase transformations of small systems is the size dependence of the thermodynamic transition point [1, 122]. For CdSe nanocrystals, a significant increase of transition pressure has been observed with decreasing crystal size [2, 8]. This behavior can be rationalized with simple thermodynamic models, in which different surface free energies are assigned to the two structures [8]. A higher surface free energy in the high-pressure phase, resulting from unfavorable configurations introduced during the transformation, stabilizes the low-pressure phase and pushes the transition pressure to higher values. In both experiment and simulation this behavior is not easily observed because transformations happen far from equilibrium and large kinetic effects obscure comparatively subtle changes in thermodynamics.

Here, we estimate the size dependence of the thermodynamic transition pressure by determining the pressures at which the two structures have equal enthalpy. Using the algorithm discussed above, we calculate the enthalpy-pressure curves of the crystals in the two stable structures in long Monte Carlo runs of 10^5 cycles each and determine their intersection. As the morphology of the rocksalt crystal is not unique, we calculate enthalpy-pressure curves for all different rocksalt crystals observed with transition path sampling and obtain different transition pressures for different isomorphs [1]. However, typical variations are small (≈ 0.05 GPa) and we define the average of the obtained values as the transition pressure for the given crystal size; the resulting data is shown in Fig. 6.9. Although our simulations reproduce the trend with crystal size observed in experiments, the experimental transition pressures are larger by a factor of about two. This difference might be partly due to deficiencies of the model used in our simulations. However, the experimental method of assigning the thermodynamic transition pressure to the midpoint of the hysteresis curve presumes that activation energies for the forward and backward transformations show a similar dependence on pressure, an assumption that is not necessarily true. Indeed, an extrapolation of the pressure dependence of the rate constants for both directions of the transformation yields a coexistence pressure of 2.3 GPa for a 2.5 nm nanocrystal [10].

6.4.4 Activation enthalpy and volume

Using the Monte Carlo algorithm discussed in Sec. 6.4.1, we calculate the enthalpy and volume as a function of time for all selected transition pathways: For every configuration of the crystal, we average over 1000 realizations of the pressure bath with crystal atoms held fixed. The resulting curves are exemplarily shown in Fig. 6.10 for the transformation of a 3 nm crystal at 2.5 GPa. To dispose of short-lived fluctuations, we perform the following smoothing operations. First, we average enthalpy and volume over time intervals of 0.5 ps. We then correct for fluctuations of the crystal kinetic temperature T_{kin} around the bath temperature $T = 300$ K. In a harmonic approximation, these fluctuations offset the potential energy $U^{(\mathrm{cc})}$ of the crystal by factors of $3/2 N_{\mathrm{c}} k_{\mathrm{B}} (T_{\mathrm{kin}} - T)$, which we subtract from $U^{(\mathrm{cc})}$. Recent work strongly indicates that such a harmonic approximation is valid even in regimes of high pressure and temperature [123]. We note, however, that the applied smoothing is not essential to the calculation and merely reduces the spread of activation enthalpies and volumes determined from different pathways.

[1] For 6 nm crystals no transition path sampling simulations were performed; for this size, the calculations were carried out using rocksalt configurations obtained in high-pressure molecular dynamics runs.

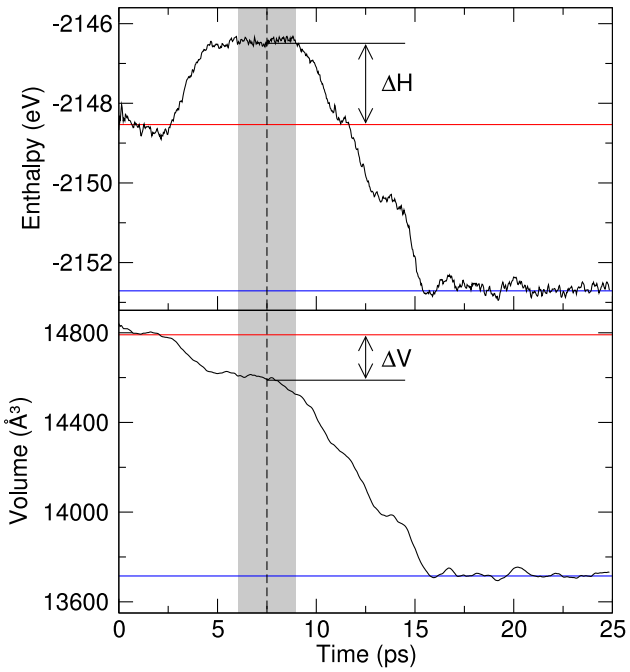

Figure 6.10: Enthalpy (top) and volume (bottom) as a function of time for a sample transformation pathway of a 3 nm crystal at 2.5 GPa. The gray field marks the interval in which the committor p_{B} changes from 0 to 1, the dotted line indicates the transition state at $p_{\mathrm{B}} = 0.5$, which coincides with the top of the enthalpic barrier. Red and blue lines indicate the values of enthalpy and volume in the metastable h-MgO and rocksalt structures, respectively. By comparing values of enthalpy and volume at the transition state with those in the hexagonal structure, the activation enthalpy ΔH and volume ΔV can be easily determined.

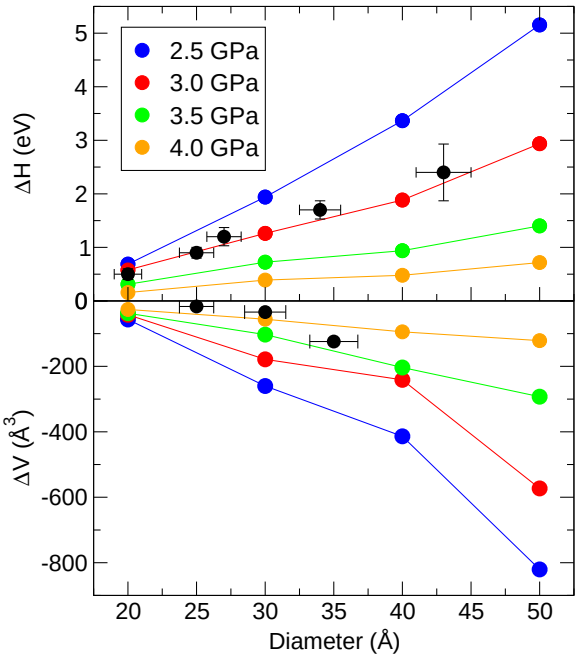

Figure 6.11: Size and pressure dependence of activation enthalpy (top) and volume (bottom). Black dots with errorbars indicate experimental data [10].

Finally, the activation enthalpy (and similarly, activation volume) is obtained as the difference of the enthalpy at the transition state and the average enthalpy in the *h*-MgO structure, as shown in Fig. 6.10. Values for different transition states along a given trajectory are averaged to give single-pathway activation enthalpies, which in turn are averaged over all 5 analyzed pathways to obtain the activation enthalpy for a given crystal size and pressure.

6.5 Comparison with experiment

In Fig. 6.11, activation enthalpies and volumes are compared with experiment. The size dependence of the experimental activation enthalpies, which were obtained from the temperature dependence of the rate constant at a pressure of 4.9 GPa [9,10], is in good agreement with data we obtained at 3 GPa. The discrepancy with regard to absolute values is not surprising and most likely due to deficiencies of the simple pair potential used in our simulations [85], which was designed to reproduce selected properties of bulk CdSe. On the other hand, the almost linear dependence of the enthalpic barrier on the crystal diameter is found in both experiment and simulation and is consistent with the particular shape of the critical nuclei shown in Fig. 6.4: With increasing crystal size, the critical rocksalt nuclei expand pri-

marily in one direction. Previously it was speculated that nanocrystals in the size range considered here would transform not through local nucleation and growth but in a single step, through a concerted motion of all crystal atoms [8, 9]. Such a mechanism, however, would result in an activation enthalpy that scales with the volume of the crystal, a behavior not observed in this work.

The comparison of activation volumes is complicated by the fact that, experimentally, activation volumes were determined at slightly different conditions than activation enthalpies, namely from changes in the rate constant as pressure was varied between 5 and 7 GPa, assuming a constant value of the activation volume in this pressure range [10]. Our simulations, however, indicate that the activation volume depends sensitively on the applied pressure. Thus, agreement between experiment and simulation might be expected to occur at higher pressures compared to activation enthalpies; indeed, we find fair agreement at 4 GPa.

Clearly, the observed trends with crystal size of activation enthalpy and volume can not continue to the bulk limit, where transformations are characterized by a critical nucleus of finite size that preferentially forms at defect sites of the crystal lattice. The transition to this different regime can be expected to occur at crystal sizes at which the concentration of such defects is no longer negligible.

6.6 Reaction coordinate

To substantiate the nucleation and growth scenario emerging from the visual inspection of the transition states, we identify the reaction coordinate, a dynamically relevant measure for the progress of the transformation. In principle, the committor function p_{B} is the ideal reaction coordinate in the sense that it tells what is likely to happen next. However, due to its unspecific nature, the committor does not provide direct insight into the transition mechanism nor can it be controlled in an experiment. Thus, it is desirable to express the committor in terms of variables with a transparent physical significance. As such a coordinate we propose to use the number of six-coordinated atoms, N_6, which can be viewed as a measure for the size of the rocksalt nucleus forming in the hexagonal matrix.

In contrast to the simple discrete criterion discussed in Sec. 6.2.2, N_6 is calculated using a continuous distance criterion. Every pair of atoms is assigned a "bond length" d_{ij} according to

$$d_{ij} = \begin{cases} 1 & \text{if } r_{ij} < a \\ \frac{b - r_{ij}}{b - a} & \text{if } a < r_{ij} < b \\ 0 & \text{if } r_{ij} > b\,. \end{cases} \tag{6.26}$$

Here, r_{ij} is the particle distance, $a = 3.0$ Å, and $b = 4.8$ Å. (These values

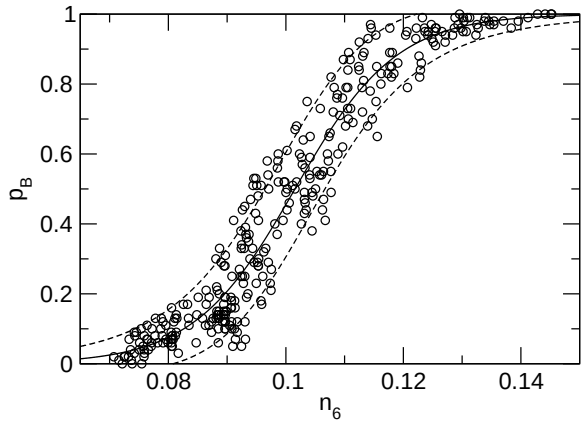

Figure 6.12: A good reaction coordinate: Committor p_{B} as a function of the fraction of 6-coordinated atoms, $n_6 = N_6/N_{\mathrm{c}}$, for configurations collected from transition pathways of a 3 nm nanocrystal at a pressure of 2.5 GPa. The solid curve is a fit to the data and the dashed curves indicate the 3σ-confidence interval around the fit. Only 14.8% of the data points fall outside of this interval, which is consistent with the statistical uncertainty of our committor calculation.

delimit the broad first minimum of the radial distribution function of unlike atom species.) The coordination number of particle i is defined by summation over all atoms of the other species, $c_i = \sum_j d_{ij}$. Correspondingly, we define the number of z-coordinated atoms as $N_z = \sum_i f_z(c_i)$, with the function

$$f_z(c) = \begin{cases} 1 - |c - z| & \text{if } \; z - 1 < c < z + 1 \\ 0 & \text{else}\,. \end{cases} \tag{6.27}$$

Figure 6.12 shows the committor vs. the fraction of six-coordinated atoms, $n_6 = N_6/N_{\mathrm{c}}$, for 330 configurations taken from four different transition trajectories of a 3 nm nanocrystal at a pressure of 2.5 GPa. As can be inferred from the figure, a particular value of N_6 specifies the corresponding value of the committor with good accuracy, indicating that N_6 is a satisfactory representation of the nucleation mechanism. The remaining spread of the data points is consistent with the statistical uncertainty of the committor calculation. To support this argument, the function $p(n_6) = \frac{1}{2}\{1 + \tanh[a(n_6 - b)]\}$ is used for fitting as a possible parameterization of the committor in Fig. 6.12. If this parameterization is perfectly valid, all deviations of the data points from $p(n_6)$ must be due to statistical errors because of the finite number of trajectories, $M = 100$ in this case, that are used to determine the committor of each configuration. Assuming Gaussian statistics, the standard deviation of p_{B} is given by Eq. (6.1); the corresponding 3σ confidence-interval is indicated in Fig. 6.12 and includes

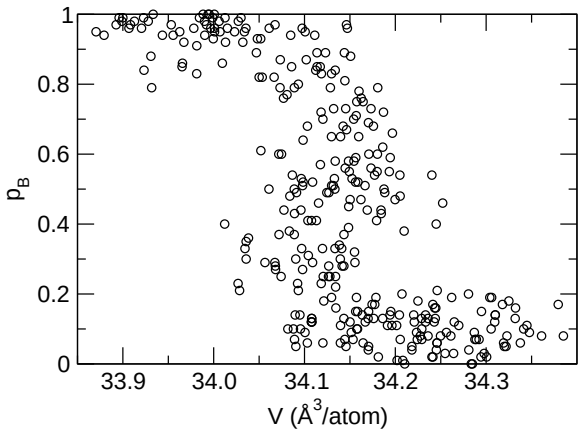

Figure 6.13: A poor reaction coordinate: Committor p_{B} as a function of the crystal volume for configurations collected from transition pathways of a 3 nm nanocrystal at a pressure of 2.5 GPa.

most of the data points.

As an example of a rather poor reaction coordinate, Fig. 6.13 displays an analogous analysis for the crystal volume. Although the volume changes considerably during the transformation, it is not a satisfactory indicator of the progress of the transition because of large fluctuations in configurational degrees of freedom that are not relevant for the nucleation event. Only slightly better results are achieved for the simple discrete coordination criterion discussed in Sec. 6.2.2, which is insensitive to small changes in bond angles and distances which play an important role during nucleation.

The success of N_6 as a reaction coordinate is a direct consequence of the small size of the crystals and the uniformity of the rocksalt nuclei. When crystal size is increased toward the bulk limit, nucleation from different sites becomes important and N_6 ceases to be a good reaction coordinate because many uncorrelated rocksalt nuclei of subcritical sizes may lead to N_6-values that are characteristic for the critical nucleus in the nano-scale crystal. Then, a cluster criterion identifying contiguous regions of 6-coordination is required for the definition of a good reaction coordinate.

6.7 Summary

In the study of phase transitions, a direct comparison between experiment and molecular dynamics computer simulation is often difficult to achieve due to long experimental time scales and a lack of properties that are easily determined with both methods. For the case of structural transformations, we introduce a systematic simulation methodology which allows to make di-

rect contact with experiments through the calculation of activation volumes and enthalpies. Simulations under experimental conditions are performed with transition path sampling, a technique based on a statistical description of pathways which is able to identify the most important mechanistic route of the transformation. Pathways generated with this technique are subject to a committor analysis which yields transition states containing the critical nuclei of the product phase and gives access to the desired quantitative properties.

As a case of special interest, we apply this methodology to the pressure-induced transformation in a model of CdSe nanocrystals suspended in an explicit pressure bath. Our results explain the size trend of activation enthalpies observed experimentally in terms of a critical nucleus of elongated shape located on the crystal surface and highlight the different nucleation mechanisms in bulk and nanocrystalline systems. Although applied here to the latter, the methodology is applicable also to the bulk case and also to other transformations of condensed matter systems frequently studied with computer simulations.

Appendix A

Dynamics of small phase space displacements

In Chapter 3 we discussed, in the context of the transition path sampling method, how extraordinarily small perturbations in phase space can be traced in time by exploiting their linear short-time dynamics and showed that the algorithm is stable for arbitrarily long times. To develop a better understanding of the errors associated with this method, in this chapter we investigate the time evolution of small phase space displacements beyond the linear regime. For a one dimensional mapping we show that under mild assumptions the terms of order n grow in time as $\mathrm{e}^{n\lambda t}$, where λ is the largest Lyapunov exponent of the system. We present evidence from computer simulations that the same behavior can also be expected in many particle systems. The interplay with numerical rounding errors which are ubiquitous in computer simulations are exemplified in detail for a liquid of soft spheres. The results are relevant in the context of numerical methods for the determination of a system's Lyapunov spectrum and in general for the stability of dynamical systems in the linear approximation.

A.1 Time evolution of small perturbations

We consider a one-dimensional system whose initial state is specified by the coordinate x_0 and which evolves in discrete time steps denumbered by n according to the propagator map Φ,

$$x_{n+1} = \Phi(x_n) \; . \tag{A.1}$$

Applying the propagator n times, we define a family of maps by

$$x_n = (\Phi \circ \Phi \circ \cdots \circ \Phi)(x_0) \equiv \Phi_n(x_0) \; , \tag{A.2}$$

where $\Phi_1 \equiv \Phi$. We now follow two trajectories x_n and y_n, which are initially separated by a small displacement

$$\delta_0 = y_0 - x_0 \,. \tag{A.3}$$

We expand δ_n in terms of δ_0,

$$\begin{aligned} \delta_n = y_n - x_n &= \Phi_n(x_0 + \delta_0) - \Phi_n(x_0) \\ &= \Phi_n'(x_0)\delta_0 + \frac{1}{2}\Phi_n''(x_0)\delta_0^2 + \frac{1}{6}\Phi_n'''(x_0)\delta_0^3 + \cdots \,, \end{aligned} \tag{A.4}$$

where the primes indicate derivatives. Dynamical instability of the system implies that, in the linear approximation,

$$|\delta_n| \sim |\delta_0| \mathrm{e}^{\lambda n} \,, \tag{A.5}$$

where λ is the largest Lyapunov exponent of the system [124]. Thus, for the one-dimensional map,

$$|\Phi_n'(x_0)| \sim \mathrm{e}^{\lambda n} \,. \tag{A.6}$$

Next, we want to estimate the time dependence of the second order term. Taking the derivative of equation (A.2) with respect to the initial condition x_0 gives

$$\Phi_n'(x_0) = \Phi'(x_{n-1})\Phi'(x_{n-2}) \cdots \Phi'(x_0) \tag{A.7}$$

and for the second derivative one obtains

$$\begin{array}{rcccccc} \Phi_n''(x_0) & = & \Phi''(x_{n-1}) & \Phi'^2(x_{n-2}) & \Phi'^2(x_{n-3}) & \cdots & \Phi'^2(x_0) \\ & + & \Phi'(x_{n-1}) & \Phi''(x_{n-2}) & \Phi'^2(x_{n-3}) & \cdots & \Phi'^2(x_0) \\ & & \vdots & & & & \\ & + & \Phi'(x_{n-1}) & \Phi'(x_{n-2}) & \cdots & \Phi''(x_1) & \Phi'^2(x_0) \\ & + & \Phi'(x_{n-1}) & \Phi'(x_{n-2}) & \cdots & \Phi'(x_1) & \Phi''(x_0) \,. \end{array} \tag{A.8}$$

Let us suppose that in this expression, as suggested by equation (A.6), the first derivatives give a contribution of e^{λ} each. Let us further assume that the second derivatives are all of the same order of magnitude, denoted by a. As Φ'' appears only once in every summand of equation (A.8), its contribution does not grow as we go to large times n. Thus,

$$\begin{aligned} |\Phi_n''(x_0)| &\sim a\big(\mathrm{e}^{\lambda(2n-2)} + \mathrm{e}^{\lambda(2n-4+1)} + \mathrm{e}^{\lambda(2n-6+2)} + \cdots + \mathrm{e}^{\lambda(n-1)}\big) \\ &= a\,\mathrm{e}^{2n\lambda}\big(\mathrm{e}^{-2\lambda} + \mathrm{e}^{-3\lambda} + \mathrm{e}^{-4\lambda} + \cdots + \mathrm{e}^{-\lambda(n+1)}\big) \\ &= a\,\mathrm{e}^{2n\lambda}\,\mathrm{e}^{-2\lambda}\big(1 + \mathrm{e}^{-\lambda} + \mathrm{e}^{-2\lambda} + \cdots + \mathrm{e}^{-\lambda(n-1)}\big) \\ &= a\,\mathrm{e}^{2n\lambda}\,\mathrm{e}^{-2\lambda}\frac{1-\mathrm{e}^{-n\lambda}}{1-\mathrm{e}^{-\lambda}} \quad \xrightarrow{n\to\infty} \quad a\,\mathrm{e}^{2n\lambda}\frac{\mathrm{e}^{-\lambda}}{\mathrm{e}^{\lambda}-1} \,. \end{aligned} \tag{A.9}$$

It follows that for large enough times n the logarithmic growth rate of the second order term is twice that of the linear term.

Of course, the same result can be obtained in more compact notation, which also allows an easy evaluation of higher order terms. We have

$$\frac{\partial x_n}{\partial x_0} \equiv \Phi'_n(x_0) = \prod_{i=0}^{n-1} \Phi'(x_i)\,, \tag{A.10}$$

and for the second derivative

$$\begin{aligned}\Phi''_n(x_0) &= \sum_{i=0}^{n-1} \Phi''(x_i) \prod_{j=0}^{i-1} \Phi'^2(x_j) \prod_{k=i+1}^{n-1} \Phi'(x_k) = \\ &= \sum_{i=0}^{n-1} \Phi''(x_i)\, \Phi_i'^2(x_0)\, \Phi'_{n-i-1}(x_{i+1})\,,\end{aligned} \tag{A.11}$$

where we have used equation (A.10) in the second step. Then,

$$\begin{aligned}|\Phi''_n(x_0)| &\sim a \sum_{i=0}^{n-1} \mathrm{e}^{2\lambda i} \mathrm{e}^{\lambda(n-i-1)} = a \sum_{i=0}^{n-1} \mathrm{e}^{\lambda(n+i-1)} \\ &= a \sum_{i=0}^{n-1} \mathrm{e}^{\lambda(2n+2-i)} = a\, \mathrm{e}^{2\lambda(n-1)} \sum_{i=0}^{n-1} \mathrm{e}^{-\lambda i} \\ &= a\, \mathrm{e}^{2\lambda(n-1)} \frac{1-\mathrm{e}^{-n\lambda}}{1-\mathrm{e}^{-\lambda}} \quad \xrightarrow{n\to\infty} \quad \mathrm{e}^{2n\lambda} \times \text{const.}\end{aligned} \tag{A.12}$$

Along the same lines we can also calculate the third order terms by taking the derivative of equation (A.11) with respect to x_0,

$$\begin{aligned}\Phi'''_n(x_0) = &\sum_{i=0}^{n-1} \Phi'''(x_i) \Phi'_i(x_0) \Phi_i'^2(x_0) \Phi'_{n-i-1}(x_{i+1}) \\ &+ \sum_{i=0}^{n-1} \Phi''(x_i) 2\Phi'_i(x_0) \Phi''_i(x_0) \Phi'_{n-i-1}(x_{i+1}) \\ &+ \sum_{i=0}^{n-1} \Phi''(x_i) \Phi_i'^2(x_0) \Phi''_{n-i-1}(x_{i+1}) \Phi'_{i+1}(x_0)\,.\end{aligned} \tag{A.13}$$

We again assume that the first two derivatives give contributions $|\Phi'_n(x_i)| \sim \mathrm{e}^{n\lambda}$ and $|\Phi''_n(x_i)| \sim \mathrm{e}^{2n\lambda}$, and that the third derivative Φ''' is roughly a

constant denoted by b. Then,

$$\begin{aligned} |\Phi_n'''(x_0)| &\sim \sum_{i=0}^{n-1} b\,\mathrm{e}^{\lambda(n+2i-1)} + 2\,\mathrm{e}^{\lambda(n+2i+1)} + \mathrm{e}^{\lambda(2n+i+1)} \\ &= \sum_{i=0}^{n-1} b\,\mathrm{e}^{\lambda(3n-2i-3)} + 2\,\mathrm{e}^{\lambda(3n-2i-1)} + \mathrm{e}^{\lambda(3n-i)} \\ &= \mathrm{e}^{3n\lambda} \sum_{i=0}^{n-1} b\,\mathrm{e}^{-\lambda(2i+3)} + 2\,\mathrm{e}^{-\lambda(2i+1)} + \mathrm{e}^{-\lambda i} \\ &\xrightarrow{n\to\infty} \quad \mathrm{e}^{3n\lambda} \times \text{const.} \end{aligned} \tag{A.14}$$

Similarly, yet higher derivatives can be evaluated and we expect the terms of order k to grow in time according to

$$|\Phi_n^{(k)}(x_0)| \sim \mathrm{e}^{kn\lambda}\,. \tag{A.15}$$

A.2 Unconstrained growth

In this section, we investigate the applicability of equation (A.15) to the case of higher dimensional systems, of which many-particle systems evolving according to Newton's equations of motion are of particular interest. In a simple model system of soft spheres the time evolution of two small phase space displacements is monitored which initially differ only in their magnitude. This particular trick lets us directly observe terms in the time evolution beyond the linear approximation. As the calculations are performed on a computer using floating point numbers in double precision representation, rounding errors may have significant influence on the time evolution of very small displacements. These effects are also accounted for in the following discussion. Except where explicitly mentioned, the model system consists of 108 WCA particles [69] in their liquid state. The total energy per particle is 1.0 and the density is 0.75, in reduced units. The equations of motion are integrated with the velocity Verlet algorithm [70] with a time step of $\Delta t = 0.002$.

We follow the time evolution of three trajectories $x_t^{(0)}$, $x_t^{(1)}$, and $x_t^{(2)}$, as well as the displacements $\delta_t^{(1)} = x_t^{(1)} - x_t^{(0)}$ and $\delta_t^{(2)} = x_t^{(2)} - x_t^{(0)}$, where the time t is continuous in contrast to the preceding section. Note that we are now dealing with vectors in phase space, consisting of the positions and velocities of all particles. For a system of N particles in three dimensional space, the dimension is $d = 6N$. We expand the displacements $\delta_t^{(1)}$ and $\delta_t^{(2)}$ up to second order,

$$\begin{aligned} \delta_t^{(1)} &= S\,\delta_0^{(1)} + \frac{1}{2}\delta_0^{(1)} A\,\delta_0^{(1)}\,, \\ \delta_t^{(2)} &= S\,\delta_0^{(2)} + \frac{1}{2}\delta_0^{(2)} A\,\delta_0^{(2)}\,. \end{aligned} \tag{A.16}$$

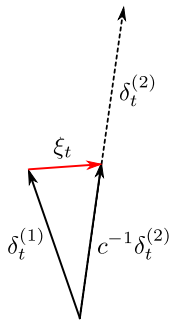

Figure A.1: The magnitude of ξ_t measures deviations from perfect alignment of the displacement vectors $\delta_t^{(1)}$ and $\delta_t^{(2)}$.

The terms involving S are first order, the terms involving A representations of second order terms in the Taylor expansion. For a system of more than one dimension, S and A are actually second and third order tensors, respectively. We now consider two initially parallel displacements, $\delta_0^{(2)} = c\,\delta_0^{(1)}$, with $|\delta_0^{(1)}| \equiv \sigma$. In the linear approximation, the initial proportionality holds for all times, $\delta_t^{(2)} = c\,\delta_t^{(1)}$. We want to evaluate deviations from this perfectly linear behavior by calculating the difference vector (see Fig. A.1)

$$\xi_t \equiv c^{-1}\delta_t^{(2)} - \delta_t^{(1)} = \frac{1}{2}(c-1)\,\delta_0^{(1)} A\,\delta_0^{(1)}\,. \tag{A.17}$$

Note that first order terms have canceled out. Now, let us assume that, in accordance with equation (A.15), the second order term, which involves second derivatives of the phase space point with respect to initial conditions, grows exponentially in time according to $|A| \sim \mathrm{e}^{2\lambda t}$, where a is a constant. Thus, for $c \gg 1$,

$$|\xi_t| \approx \left|\frac{1}{2}\,c\,\delta_0^{(1)} A\,\delta_0^{(1)}\right| \approx c\,\sigma^2 a\,\mathrm{e}^{2\lambda t}\,, \tag{A.18}$$

where a is a constant.

On a computer, rounding errors of the order of the machine precision have to be taken into account. We represent them by small phase space vectors of magnitude $\sqrt{d}\,\mathcal{F}$, where $\mathcal{F} = 10^{-15}$ is the size of the typical error of a double precision computer number of order 1. (We assume that this is the typical size of the components of phase space vectors.) These errors are added to the "correct" trajectories at every time step Δt and subsequently grow in time according to $\sqrt{d}\,\mathcal{F}\mathrm{e}^{\lambda t}$. Higher order terms are neglected here. We estimate the total accumulated error after a time interval t, which is

large compared to a single time step, by

$$\begin{aligned} \mathcal{F}_t^{\text{tot}} &= \sqrt{d}\,\mathcal{F}\left(\mathrm{e}^{\lambda t} + \mathrm{e}^{\lambda(t-\Delta t)} + \cdots + \mathrm{e}^{\lambda \Delta t}\right) \\ &\approx \frac{\sqrt{d}\,\mathcal{F}}{\Delta t}\int_0^t \mathrm{e}^{\lambda(t-\tau)}\,\mathrm{d}\tau = \frac{\sqrt{d}\,\mathcal{F}}{\lambda \Delta t}\left(\mathrm{e}^{\lambda t}-1\right) \equiv \hat{\mathcal{F}}\left(\mathrm{e}^{\lambda t}-1\right). \end{aligned} \tag{A.19}$$

To measure the relative deviation from linear behavior, we use the quantity $\epsilon(t)$ defined as $|\xi_t|$ divided by the magnitude of the smaller displacement,

$$\epsilon(t) = \frac{|\xi_t|}{|\delta_t^{(1)}|} \approx \frac{|\xi_t|}{\sigma \mathrm{e}^{\lambda t}}. \tag{A.20}$$

The magnitude of the actual difference vector $|\xi_t|$ observed in a simulation will be determined by competition of second order terms from Eqn. (A.18) and rounding errors from Eqn. (A.19). Depending on the magnitudes of the initial displacements, we will therefore see two different regimes as we follow the time evolution of two initially parallel displacements,

$$\epsilon(t) \approx \begin{cases} \hat{\mathcal{F}}/\sigma\left(1-\mathrm{e}^{-\lambda t}\right) & \text{if} \quad c\,\sigma^2 a\,\mathrm{e}^{2\lambda t} \ll \hat{\mathcal{F}}\left(\mathrm{e}^{\lambda t}-1\right), \\ c\,\sigma a\,\mathrm{e}^{\lambda t} & \text{if} \quad c\,\sigma^2 a\,\mathrm{e}^{2\lambda t} \gg \hat{\mathcal{F}}\left(\mathrm{e}^{\lambda t}-1\right). \end{cases} \tag{A.21}$$

For small values of σ and short times t, we can initially ignore second order terms. Thus, $\epsilon(t)$ will be dominated by rounding errors and show a plateau value of $\hat{\mathcal{F}}/\sigma$. For the model system we are considering ($d = 648$, $\lambda \approx 3.1$ and $\Delta t = 0.002$), $\hat{\mathcal{F}} = \sqrt{d}\,\mathcal{F}/\lambda\Delta t \approx 4\times 10^{-12}$, a number which can be considered as the "effective precision level". For larger values of σ or longer times t, deviations from linearity are dominated by second order terms and we expect exponential growth of $\epsilon(t)$. The time t_{x} at which the crossover between the two regimes occurs can be estimated by solving $c\,\sigma^2 a\,\mathrm{e}^{2\lambda t_{\text{x}}} = \hat{\mathcal{F}}\left(\mathrm{e}^{\lambda t_{\text{x}}}-1\right)$ for t_{x}, resulting in

$$t_{\text{x}} \approx \frac{1}{\lambda}\ln\frac{\hat{\mathcal{F}}}{c\,\sigma^2 a}, \tag{A.22}$$

where we have assumed $\mathrm{e}^{\lambda t_{\text{x}}} \gg 1$. In Figure A.2 this situation is examined in our model system. The crossover occurs at the expected point in time and the slope of $|\epsilon(t)|$ in the regime of dominating second order terms indicates that these indeed grow with an exponent equal to 2λ. Figures A.3 and A.4 further illustrate the significance of Eqn. (A.21).

A.3 Periodic rescaling

Inspired by the precision shooting algorithm discussed in Chapter 3, we now consider periodically rescaled displacements: After time intervals of length τ (at times $t = \tau, 2\tau, \ldots$) we reset the magnitudes of $\delta_t^{(1)}$ and $\delta_t^{(2)}$ to $|\delta_0^{(1)}|$ and $|\delta_0^{(2)}|$, respectively. Because of this rescaling, the dynamics is not continuous

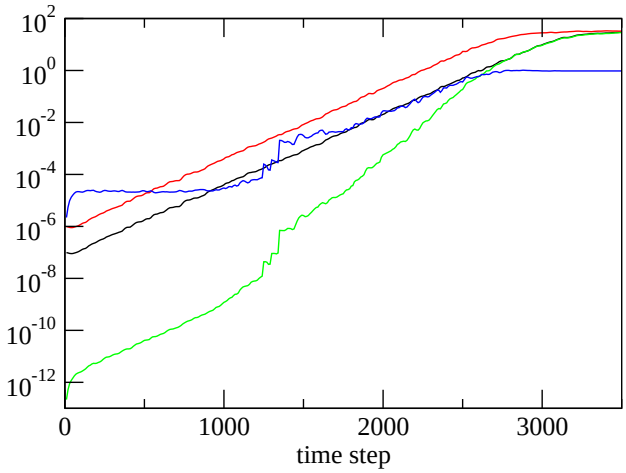

Figure A.2: Time evolution of the magnitude of two shooting displacements $|\delta_t^{(1)}|$ (black) and $|\delta_t^{(2)}|$ (red) with $\sigma = 10^{-7}$ and $c = 10$. The magnitude of the difference vector ξ_t is shown in green, the blue curve corresponds to $\epsilon(t)$. For $t \lesssim 1000\,\Delta t$, rounding errors dominate, $|\xi_t|$ grows with an exponent equal to λ, and $\epsilon(t)$ has a plateau value of about 2×10^{-5}. Multiplying this value by σ we get $\hat{\mathcal{F}} = \epsilon(t) \times \sigma \approx 2 \times 10^{-12}$ which compares well with the estimate given in the text. For $t > 1000\,\Delta t$, second order terms take over and $|\xi_t|$ grows with an exponent larger than λ, resulting in exponential growth of $\epsilon(t)$. The fact that $\epsilon(t)$ closely follows $|\delta_t^{(1)}| \approx \sigma \mathrm{e}^{\lambda t}$ in that regime indicates that, as we derived for the one-dimensional case in Section A.1, second order terms grow with an exponent equal to 2λ. Noting that $a \approx 0.1$ we can obtain a rough estimate for the crossover time t_x from Eqn. (A.22); this calculation indeed gives $t_\mathrm{x} \approx 1000\,\Delta t$ in good agreement with the simulation.

and we have to consider the time evolution of the displacement vectors separately for each time interval between individual rescaling points. Initially, the time evolution of the displacements follows the same rules discussed in Section A.2. But starting from $t = \tau$, the point of the first rescaling, things work differently. Because of deviations from perfect linearity (dominated either by rounding errors or second order terms), the two displacements are not aligned anymore. After rescaling,

$$\delta_\tau^{(2)} = c\left(\delta_\tau^{(1)} + \eta_\tau\right). \tag{A.23}$$

Here,

$$\eta_\tau \approx \frac{\hat{\xi}_\tau}{\mathrm{e}^{\lambda\tau}} \tag{A.24}$$

is the difference vector $\hat{\xi}_\tau$ (*before* rescaling) divided by the rescaling factor which is roughly $\mathrm{e}^{\lambda\tau}$. The time evolution of the displacements from the first

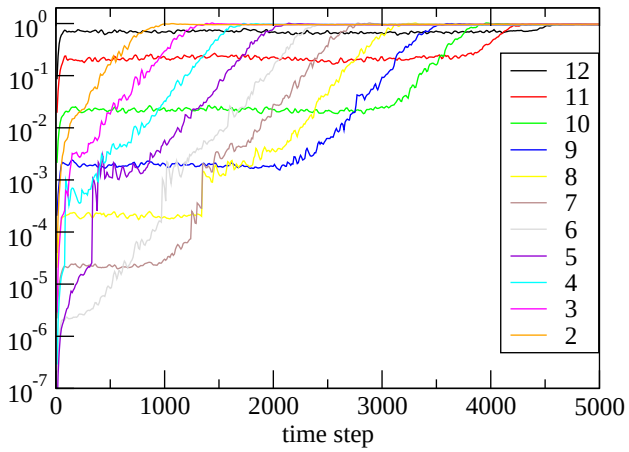

Figure A.3: Time evolution of $\epsilon(t)$ for 10 different pairs of displacements with $\sigma = 10^{-n}$ and $c = 10$. (Legend values indicate values of n.) Starting from $\sigma = 10^{-12}$, the plateau value of $\epsilon(t)$ decreases by a factor of 10 as σ is decreased by the same factor. The plateaus vanish for $\sigma \gtrsim 10^{-6}$, where second order terms are larger than rounding errors even for small times. Using the model developed in the text, an estimate for this critical size of σ can be obtained from $c\,\sigma^2 a = \hat{\mathcal{F}}$, which gives $\sigma \approx 2\times 10^{-6}$ in good agreement with the simulation.

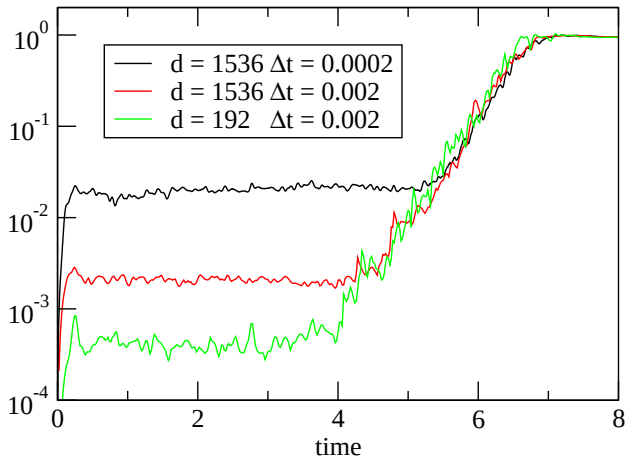

Figure A.4: Time evolution of $\epsilon(t)$ for three pairs of displacements with $\sigma = 10^{-9}$ and $c = 10$, integrated using two different values of the time step Δt and two different system sizes. The plateau value $\hat{\mathcal{F}}/\sigma = \sqrt{d}\,\mathcal{F}/\lambda\Delta t\sigma$ shows the expected dependence on Δt, the dependence on phase space dimension d is qualitatively correct.

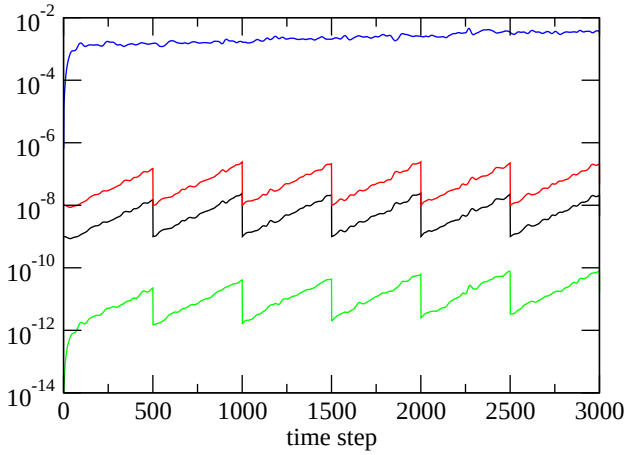

Figure A.5: Time evolution of $|\delta_t^{(1)}|$ (black), $|\delta_t^{(1)}|$ (red), $|\xi_t| = |c^{-1}\delta_t^{(2)} - \delta_t^{(1)}|$ (green), and $\epsilon(t)$ (blue) for $\sigma = 10^{-9}$, $c = 10$, and $\tau = 500\Delta t$. Note the slow increase of $|\xi_t|$ and $\epsilon(t)$.

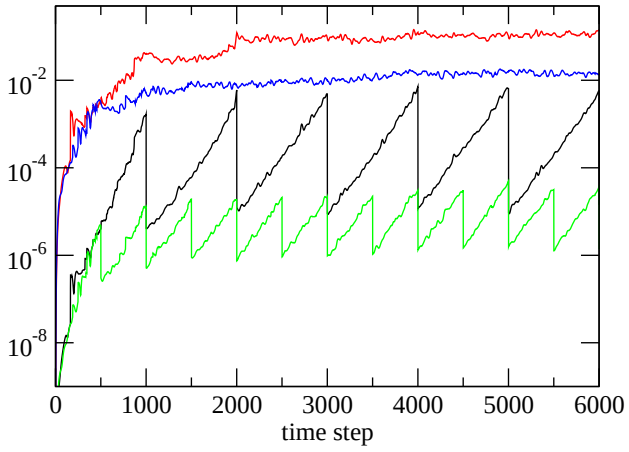

Figure A.6: Dependence of the plateau value on the rescaling interval τ. Two data sets are shown, both for pairs of displacements (not shown) with $\sigma = 10^{-4}$ and $c = 10$, but for different values of τ: $|\xi_t|$ (green) and $\epsilon(t)$ (blue) for $\tau = 500\Delta t$; $|\xi_t|$ (black) and $\epsilon(t)$ (red) for $\tau = 1000\Delta t$. In this regime, second order terms dominate and $\epsilon(t)$ grows proportional to $e^{\lambda t}$ up to the first rescaling point, where first order terms take over.

to the second rescaling point ($0 \leqslant \hat{t} \leqslant \tau$) then reads

$$\begin{aligned} \delta^{(1)}_{\tau+\hat{t}} &= S\,\delta^{(1)}_{\tau} + \frac{1}{2}\delta^{(1)}_{\tau} A\,\delta^{(1)}_{\tau}\,, \\ \delta^{(2)}_{\tau+\hat{t}} &= c\,S\big(\delta^{(1)}_{\tau}+\eta_\tau\big) + \frac{c^2}{2}\big(\delta^{(1)}_{\tau}+\eta_\tau\big)A\big(\delta^{(1)}_{\tau}+\eta_\tau\big)\,. \end{aligned} \tag{A.25}$$

Thus we get for the time evolution of the difference vector

$$\xi_{\tau+\hat{t}} \approx S\,\eta_\tau + \frac{1}{2}(c-1)\,\delta^{(1)}_{\tau} A\,\delta^{(1)}_{\tau}\,. \tag{A.26}$$

Here we assume $|\eta_\tau| \ll |\delta^{(1)}_{\tau}|$ and keep only the leading second order term. Along the lines of section A.2, we again introduce rounding errors. Depending on which regime we are in, the deviation from perfect alignment, after the first rescaling, will have a magnitude of

$$|\eta_\tau| \approx \begin{cases} \hat{\mathcal{F}}(\mathrm{e}^{\lambda\tau}-1)/\mathrm{e}^{\lambda\tau} = \hat{\mathcal{F}}(1-\mathrm{e}^{-\lambda\tau}) \approx \hat{\mathcal{F}} & \text{if } c\,\sigma^2 a\,\mathrm{e}^{2\lambda t} \ll \hat{\mathcal{F}}\left(\mathrm{e}^{\lambda t}-1\right), \\ c\,\sigma^2 a\,\mathrm{e}^{2\lambda\tau}/\mathrm{e}^{\lambda\tau} = c\,\sigma^2 a\,\mathrm{e}^{\lambda\tau} & \text{if } c\,\sigma^2 a\,\mathrm{e}^{2\lambda t} \gg \hat{\mathcal{F}}\left(\mathrm{e}^{\lambda t}-1\right). \end{cases} \tag{A.27}$$

η_τ grows exponentially in time,

$$|\eta_{\tau+\hat{t}}| \sim |\eta_\tau| \mathrm{e}^{\lambda\hat{t}}\,. \tag{A.28}$$

If we now add this term to the leading contribution in the respective regime, we obtain for the time evolution of $\epsilon(\tau+\hat{t})$ from the first to the second rescaling point

$$\epsilon(\tau+\hat{t}) \approx \begin{cases} \hat{\mathcal{F}}/\sigma + \hat{\mathcal{F}}/\sigma\big(1-\mathrm{e}^{-\lambda\hat{t}}\big) = \hat{\mathcal{F}}/\sigma\big(2-\mathrm{e}^{-\lambda\hat{t}}\big), \\ c\,\sigma a\,\mathrm{e}^{\lambda\tau} + c\,\sigma a\,\mathrm{e}^{\lambda\hat{t}} = c\,\sigma a\,\mathrm{e}^{\lambda\tau}\big(1+\mathrm{e}^{\lambda(\hat{t}-\tau)}\big). \end{cases} \tag{A.29}$$

As a consequence, at the second rescaling point

$$|\eta_{2\tau}| \approx \begin{cases} 2\,\hat{\mathcal{F}}, \\ 2\,c\,\sigma^2 a\,\mathrm{e}^{\lambda\tau}, \end{cases} \tag{A.30}$$

and for the time evolution from the second to the third rescaling point,

$$\epsilon(2\tau+\hat{t}) \approx \begin{cases} \hat{\mathcal{F}}/\sigma\big(3-\mathrm{e}^{-\lambda\hat{t}}\big), \\ c\,\sigma a\,\mathrm{e}^{\lambda\tau}\big(2+\mathrm{e}^{\lambda(\hat{t}-\tau)}\big). \end{cases} \tag{A.31}$$

Thus, starting from the first rescaling point, $\epsilon(t)$ has a plateau (albeit a slowly growing one) in both regimes. In the regime where second order terms are larger than rounding errors, the plateau value will depend on the time τ that passes before the first rescaling event; in the regime of rounding errors, this will not be the case. For fixed τ and c, in the regime $\sigma \lesssim 10^{-6}$

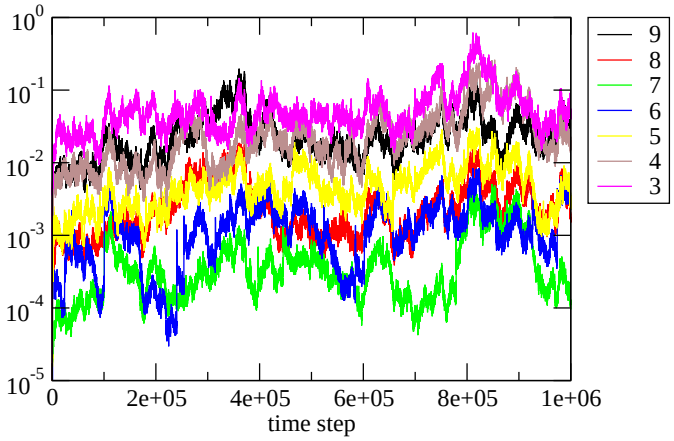

Figure A.7: Long-time behavior of $\epsilon(t)$ for different pairs of shooting displacements, with $\sigma = 10^{-n}$ and $c = 10$ (legend values indicate values of n). For $\sigma \geq 10^{-2}$ and $\sigma \leq 10^{-10}$ fluctuations, though arising for different reasons in the two regimes, eventually lead to complete separation of trajectories and $\epsilon(t) \approx 1$.

we expect the plateau value to decrease with increasing σ and then decrease again for $\sigma \gtrsim 10^{-6}$, as the system crosses over to a different regime.

As illustrated in Figures A.5 and A.6, the above argument correctly describes the behavior of our model system, at least for short times. However, for very long times the equations suggest a slow increase in $\epsilon(t)$ that will eventually lead to complete separation of trajectories. As discussed in Chapter 3, this does not happen, even for simulation times of order 10^7 time steps. We observe fluctuations in $\epsilon(t)$ over one or two orders of magnitude, but no long-time trend (see Fig. A.7). It is here that the slow alignment of phase space vectors in the direction of largest growth plays an important role [72], which is not captured in the simplified formalism used above, where we are only concerned with magnitudes of vectors and disregard their direction.

Appendix B

Nucleation and growth in structural transformations of nanocrystals

> This chapter contains a short manuscript based on Chapter 6 which has been recently accepted for publication in the journal *Nano Letters*.

Using transition path sampling computer simulations, we reveal the nucleation mechanism of a pressure-induced structural transformation in a model of CdSe nanocrystals. Consistent with experiments, the thermodynamic transition pressure of the transformation increases with decreasing crystal size. Through transition state analysis, we identify the critical nuclei and characterize them by calculating activation enthalpies and volumes. Our simulations reproduce the trends with crystal size observed in experiments. This result supports the observed transformation mechanism, which consists of nucleation on the crystal surface and growth by sliding of parallel crystal planes.

B.1 Introduction

In nanocrystals, mechanisms of structural phase transformations are often strongly influenced by the surface free energetics related to particular morphologies. For instance, in ZnS qualitatively different transformation pathways are observed for spherical and belt-like geometries [105, 106]. Control over such pathways promises new routes for creating materials by kinetically stabilizing structures that are unstable otherwise [12, 105, 106], but requires detailed understanding of the transformation mechanism on the atomistic level. While experiments provide detailed insights into the thermodynamics and kinetics of structural transitions, their time and space resolution is not

sufficient to identify transition mechanisms. Computer simulation methods can, in principle, yield such information but are complicated by long time scales. Here, we use transition path sampling (TPS) [17, 18] to overcome these difficulties and identify the nucleation pathway for a pressure induced structural transformation in CdSe nanocrystals. For the first time, close contact between simulation and experiment is established through computation of activation volumes and enthalpies from transition state analysis. Our simulations explain the trends with crystal size observed in experiments and shed light on the underlying nucleation mechanism, which is markedly different from mechanisms occurring in bulk materials.

The paradigmatic example of pressure-induced transformations in nanocrystals is the transition from the four-coordinated wurtzite to the six-coordinated rocksalt structure in CdSe nanocrystals. In an extensive series of experiments, Alivisatos and coworkers have studied the size-dependence of the hysteresis curve and of the transformation pressure, and have measured activation volumes and enthalpies [8–12]. While these experiments elucidate the thermodynamics and kinetics of the transformation, only limited information about the atomistic details of the transformation process can be extracted from them [10, 11].

Using molecular dynamics (MD) computer simulation, various transformation mechanisms have been identified for bulk [15,83] and nanocrystalline CdSe [16, 97], as well as for other semiconducting materials [14, 92, 98, 104, 110,113]. However, the significance of the observed mechanisms with respect to the experiments remains unclear for two reasons. First, a straightforward MD approach is plagued by the wide gap in timescales between experiment and simulation. To overcome the high free energy barrier, associated with the formation of a critical nucleus of the high-pressure phase, within the typical simulation timescale of a nanosecond, pressure must be increased close to the point where the low-pressure phase becomes mechanically unstable. Under these conditions, mechanisms may be quite different from those favored in a less violent regime [14, 65, 91, 97]. Second, for lack of a quantitative measure of the "right" mechanism, simulation studies so far had to be content with a mere description of observed mechanisms and no prediction of experimental observables could be made.

In this work, the time scale problem is avoided by using TPS, a computational technique for the simulation of rare events [17, 18]. In contrast to straightforward MD, TPS makes possible the study of phase transformations at experimental conditions. However, the significance of a transformation mechanism observed in simulations can only be assessed through the computation of experimentally observable quantities. To this end, we subject the trajectories harvested with TPS to a statistical analysis based on the committor [18], a statistical measure for the transition progress. This analysis yields transition states, which are familiar from chemical reaction dynamics and are instrumental for a detailed understanding of the transition

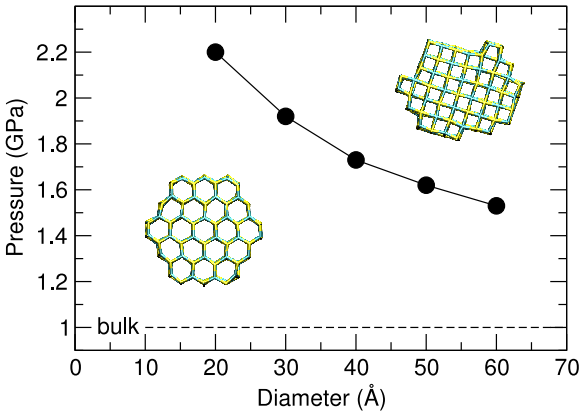

Figure B.1: Transition pressures as a function of crystal size for the h-MgO to rocksalt transformation (lines are guides to the eye). The dashed line indicates the bulk transition pressure [97]. A 30 Å crystal, viewed along the hexagonal c-axis, is shown in the hexagonal (lower left) and cubic (upper right) structure.

mechanism. From the transition states we compute activation volumes and activation enthalpies, making direct contact with experiments.

B.2 Simulation results

The basis of this work are extensive TPS simulations of faceted CdSe nanocrystals. The crystals have diameters of 20, 30, 40, and 50 Å (corresponding to $N = 120$, 432, 1056, and 2100 atoms) and are modeled with the empirical potential of Rabani [85]. The crystals, designed to resemble the particles studied experimentally, are hexagonal prisms, with (100) and (001) surfaces and an aspect ratio of about unity [87]. The wurtzite to rocksalt transformation for this crystal shape proceeds via the consecutive sliding of (100) planes [65] and passes through the five-coordinated intermediate structure h-MgO [16,98,125], which is stabilized by a favorable surface free energy [97]. The h-MgO structure, which is obtained from the wurtzite structure by a compression along the hexagonal c-axis, is not observed in experiments and disappears in the simulation of crystals with disordered surfaces or stacking faults [97]. Nevertheless, also these CdSe crystals show a compression along the wurtzite c-axis prior to the transformation, and, at the time of the transition, are essentially in the h-MgO structure. As indicated by straightforward MD simulations, the h-MgO to rocksalt transformation is the rate limiting step of the whole transition [97]. We therefore concentrate on this transformation in the following.

We perform TPS runs at a temperature of 300 K for pressures rang-

ing from 2 to 4 GPa, using a recently developed TPS algorithm for nanocrystals [65], which employs an ideal gas as pressure-medium [101]. (Simulation details can be found in Ref. [65].) Figure B.1 shows the size dependence of the thermodynamic transition pressure, which we identify with the pressure at which the hexagonal and rocksalt structures have equal enthalpy[1]. The transition pressure decreases with increasing crystal size approaching the bulk value of about 1 GPa [97]. This behavior is observed here for the first time in a computer simulation; it is consistent with experimental observations and can be rationalized in terms of the different surface free energies in the two structures [8]. In previous straightforward MD simulations with rapid pressurization, this size dependence was obscured by kinetic effects, preventing an accurate determination of the thermodynamic transition pressure [98].

To identify the transition mechanism, we determine transition states by calculating the committor p_{B}, a statistical measure for the progress of the transformation, along typical trajectories of the transition path ensemble [18]. The committor p_{B} is the probability of a particular configuration of the crystal to relax into the rocksalt structure rather than into the hexagonal structure. Thus, as the transformation proceeds, the committor changes its value from 0 to 1. A configuration of a nanocrystal is identified as a transition state if its committor equals one half.

Typical transition states of a 50 Å nanocrystal are shown in Fig. B.2a. It was speculated earlier [8, 9, 126] that in nanocrystals below a certain size a collective transformation mechanism, involving concerted motion of all crystal atoms, could be energetically favored over step-by-step growth of a solid-solid interface. Our simulations, however, support a directionally dependent nucleation and growth mechanism. Indeed, our TPS simulations start from a collective mechanism but rapidly converge toward the mechanism described in Ref. [65] and illustrated in Fig. B.2b. Transition states, defined through the committor criterion, are invariably characterized by a localized and connected region of rocksalt structure, a critical rocksalt nucleus. A conspicuous feature of the transformation process is that the rocksalt nuclei always appear on the surface of the crystal. At high pressures (3.5 GPa in Fig. B.2), the critical nucleus consists of a one-dimensional arrangement of atoms. At lower pressures (2 GPa in Fig. B.2), the transition state is reached after growth in one of two equivalent [100] directions. Consistent with classical nucleation theory, the critical nucleus shrinks as

[1]The enthalpy of a nanocrystal is defined as $H = U_{\mathrm{c}} + U_{\mathrm{g}} + PV$, where U_{c} is the interaction energy between crystal atoms, U_{g} is the interaction energy of crystal atoms with gas particles in the pressure bath, P is the pressure, and V is the volume of the crystal. The volume is defined as the difference of the total system volume V_{tot} and the volume occupied by the N_{g} gas particles according to the equation of state of the ideal gas, $V = V_{\mathrm{tot}} - N_{\mathrm{g}} k_{\mathrm{B}} T/P$. This definition is free of any geometrical parameters and is consistent with a thermodynamic theory for small systems in a solvent [120].

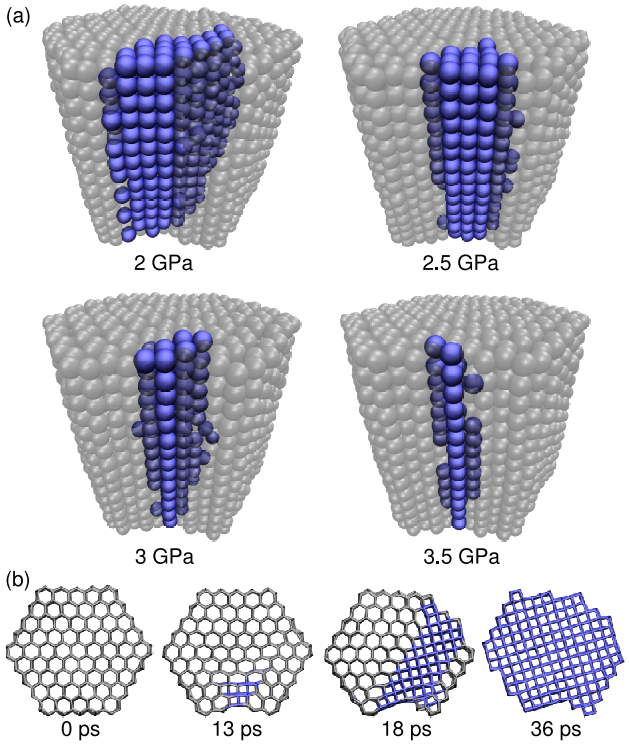

Figure B.2: Nucleation of the rocksalt structure. (a) The critical rocksalt nucleus in a 50 Å nanocrystal at four different pressures. Transparent gray atoms are in the hexagonal structure, blue atoms are in the rocksalt structure. Crystals have been rotated to facilitate comparison of the nuclei. As expected from classical nucleation theory, the size of the critical nucleus increases with decreasing pressure (for this size, the thermodynamic transition pressure is 1.6 GPa, see Fig. B.1). (b) Snapshots along a typical reactive trajectory at 3 GPa. The crystal is shown along the hexagonal c-axis to highlight the mechanism of sliding (100) planes. After 13 ps the crystal has reached the transition state.

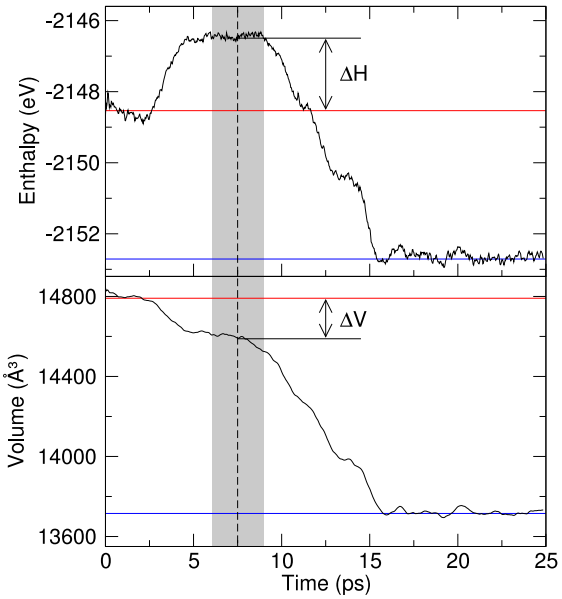

Figure B.3: Enthalpy and volume as a function of time along a typical transition path of a 30 Å nanocrystal, at a pressure of 2.5 GPa. The red and blue lines indicate the average of the enthalpy and volume, calculated from long Monte Carlo simulations, in the h-MgO and rocksalt structure, respectively. As the trajectory passes through the light gray field, the value of the committor changes from essentially zero to one; the dashed line around 7.5 ps indicates the point where the system is in the transition state ($p_{\mathrm{B}} = 0.5$). We obtain the activation enthalpy ΔH and volume ΔV from the difference of the respective value in the metastable h-MgO structure and the transition state, as indicated by the arrows.

pressure is increased, but even at a pressure of 4 GPa, close to the point where the transformation becomes observable with straightforward MD, the critical nucleus is strongly elongated and extends along the full length of the crystal. This unexpected feature holds for all crystal sizes considered here. Most likely, this particular morphology of the interface between the high pressure structure and the pressure bath is favored energetically over other solid-solid interfaces that cause large strain throughout large parts of the crystal. In the bulk, solid-solid transformations are believed to nucleate at lattice defects like grain boundaries for similar reasons [116]; in nanocrystals, the surface plays the analogous role of a low energy nucleation site.

To link our simulation results to experiments, we next calculate activation enthalpies ΔH and activation volumes ΔV for various nanocrystal sizes and pressures. Figure B.3 shows the time evolution of the enthalpy and volume of a 30 Å crystal during the transformation process. The en-

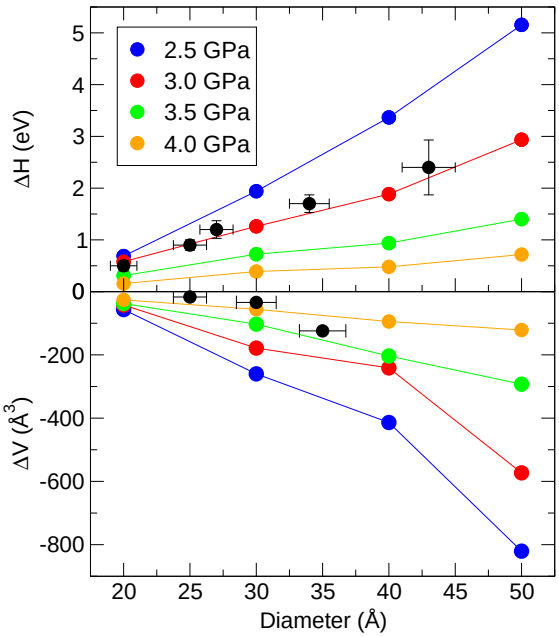

Figure B.4: Activation enthalpies and volumes as a function of crystal size, determined as averages over single trajectories like the one in Fig. B.3 (lines are guides to the eye). Black dots with errorbars mark experimental data: enthalpies were determined from the temperature dependence of the rate constant at 4.9 GPa, volumes from the change of the rate constant as pressure was varied between 5 and 7 GPa (experimental data are from the main text and from footnote (24) of Ref. [10]; no errors of measurement were provided for activation volumes).

thalpic barrier to the transformation is evident, and the transition state coincides with the barrier top, consistent with nucleation theory. We calculate ΔH and ΔV as the difference of the enthalpy and volume at the transition state and in the hexagonal structure. These quantities are then averaged over all transition states along a particular trajectory and over typically five randomly selected TPS trajectories (each having up to 25 transition states). The size- and pressure-dependence of ΔH and ΔV are shown in Fig. B.4, along with experimental data. Although our simulations do not entirely reproduce the true transition pressures, the agreement of the size-dependence of the enthalpic barrier at a pressure of 3 GPa with experimental data collected at 4.9 GPa [9, 10] is remarkable, giving evidence for the correctness of the observed sliding-planes mechanism. One might expect an analogous agreement of experimental activation volumes with our volume data at 3 GPa. However, a comparison of the size-dependence of activation volumes with experiment is complicated by the particular way the activation volumes have been determined from experimental data: assuming the activation volume is constant over a pressure range of a few GPa, Alivisatos and coworkers obtained activation volumes from the pressure-dependence of the rate constant between 5 and 7 GPa [10]. Our simulations clearly show that the activation volume decreases significantly as pressure is increased, consistent with nucleation theory. One should therefore consider the experimental volume data to be associated with higher values of pressure. The fair agreement of experimental activation volumes with our calculations at 4 GPa indeed supports this argument.

B.3 Concluding remarks

The size-dependence of the the activation enthalpy and volume sensitively depends on the particular transformation pathway and discriminates between alternative mechanisms. In particular, a mechanism involving a concerted rearrangement of all atoms in a single step would lead to an activation enthalpy that scales with the crystal volume. A transformation proceeding through planes sliding coherently with respect to each other would lead to a scaling with the surface area. In contrast, the linear behavior observed here and in experiments is a signature of a quasi-one dimensional critical nucleus that spans the entire crystal and grows only in one direction as the size of the nanocrystal is increased. This particular behavior of activation enthalpy and volume cannot extend to the bulk system in which the critical nucleus is finite and independent of system size. In the bulk, nucleation will occur preferentially at lattice defects rather than on the surface leading to a markedly lower kinetic barrier. Nevertheless, the computational approach proposed here is applicable also to this case, as well as to other structural and morphological transformations in nanocrystalline and bulk materials.

Bibliography

[1] A.N. Goldstein, C.M. Echer, A.P. Alivisatos, Science **256**, 1425 (1992)

[2] S.H. Tolbert, A.P. Alivisatos, Science **265**, 373 (1994)

[3] A.P. Alivisatos, Science **271**, 933 (1996)

[4] C.B. Murray, D.J. Norris, M.G. Bawendi, J. Am. Chem. Soc. **115**, 8706 (1993)

[5] X. Peng, L. Manna, W. Yang, J. Wickham, E. Scher, A. Kadavanich, A.P. Alivisatos, Nature **404**, 59 (2000)

[6] L. Manna, E.C. Scher, A.P. Alivisatos, J. Am. Chem. Soc. **122**, 12700 (2000)

[7] J.R. McBride, T.C. Kippeny, S.J. Pennycook, S.J. Rosenthal, Nano Letters **4**, 1279 (2004)

[8] S.H. Tolbert, A.P. Alivisatos, J. Chem. Phys. **102**, 4642 (1995)

[9] C.C. Chen, A.B. Herhold, C.S. Johnson, A.P. Alivisatos, Science **276**, 398 (1997)

[10] K. Jacobs, D. Zaziski, E.C. Scher, A.B. Herhold, A.P. Alivisatos, Science **293**, 1803 (2001)

[11] J.N. Wickham, A.B. Herhold, A.P. Alivisatos, Phys. Rev. Lett. **84**, 923 (2000)

[12] K. Jacobs, J. Wickham, A.P. Alivisatos, J. Phys. Chem. B **106**, 3759 (2002)

[13] R. Martoňák, C. Molteni, M. Parrinello, Phys. Rev. Lett. **84**, 682 (2000)

[14] B.J. Morgan, P.A. Madden, Nano Letters **4**, 1581 (2004)

[15] F. Shimojo, S. Kodiyalam, I. Ebbsjö, R.K. Kalia, A. Nakano, P. Vashishta, Phys. Rev. B **70**, 184111 (2004)

[16] N.J. Lee, R.K. Kalia, A. Nakano, P. Vashishta, Appl. Phys. Lett. **89**, 093101 (2006)

[17] C. Dellago, P.G. Bolhuis, F.S. Csajka, D. Chandler, J. Chem. Phys. **108**, 1964 (1998)

[18] C. Dellago, P.G. Bolhuis, P.L. Geissler, Adv. Chem. Phys. **123**, 1 (2002)

[19] G.M. Torrie, J.P. Valleau, J. Comp. Phys. **23**, 187 (1977)

[20] J.G. Kirkwood, J. Chem. Phys. **3**, 300 (1933)

[21] C.H. Bennett, in *ACS Symposium Series*, *ACS Symposium Series*, vol. 46 (ACS, Washington DC, 1977), p. 63

[22] D. Chandler, J. Chem. Phys. **68**, 2959 (1978)

[23] H. Eyring, J. Chem. Phys. **3**, 107 (1935)

[24] E. Wigner, Trans. Faraday Soc. **34**, 29 (1938)

[25] A. Laio, M. Parinelle, Proc Natl. Acad. Sci. **99**, 12562 (2002)

[26] A.F. Voter, F. Montalenti, T.C. Germann, Annu. Rev. Mater. Res. **32**, 321 (2002)

[27] G. Hummer, I.G. Kevrekidis, J. Chem. Phys. **118**, 10762 (2003)

[28] H. Jonssón, G. Mills, K.W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, ed. by B.J. Berne, G. Ciccotti, D.F. Coker (World Scientific, 1998), p. 385

[29] E. Weinan, W. Ren, E. Vanden-Eijnden, Phys. Rev. B **66**, 052301 (2002)

[30] E. Weinan, W. Ren, E. Vanden-Eijnden, Chem. Phys. Lett. **413**, 242 (2005)

[31] C. Dellago, P.G. Bolhuis, in *Advances in Polymer Science*, *Advances in Polymer Science*, vol. 221, ed. by C. Holm, K. K. (Springer, 2009), p. 167

[32] C. Dellago, P.G. Bolhuis, P.L. Geissler, in *Computer Simulations in Condensed Matter: From Materials to Chemical Biology. Volume 1*, ed. by M. Ferrario, G. Ciccotti, K. Binder (Springer Lecture Notes in Physics, 2006), p. 349

[33] P.G. Bolhuis, D. Chandler, C. Dellago, P.L. Geissler, Ann. Rev. Phys. Chem. **53**, 291 (2002)

[34] C. Dellago, D. Chandler, in *Molecular Simulation for the Next Decade*, ed. by P. Nielaba, M. Mareschal, G. Ciccotti (Springer, 2002), p. 321

[35] C. Dellago, in *Handbook of Materials Modeling*, ed. by S. Yip (Springer, 2005), p. 1585

[36] C. Dellago, in *Free energy calculations: Theory and applications in chemistry and biology*, ed. by A. Pohorille, C. Chipot (Springer, 2007)

[37] C. Dellago, P.G. Bolhuis, in *Topics in Current Chemistry*, vol. 268, ed. by M. Reiher (Springer, 2007), p. 291

[38] G.E. Crooks, D. Chandler, Phys. Rev. E **64**, 026109 (2001)

[39] P.L. Geissler, D. Chandler, J. Chem. Phys. **113**, 9759 (2000)

[40] P.G. Bolhuis, C. Dellago, D. Chandler, Faraday Discuss. **110**, 421 (1998)

[41] M. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.N. Teller, E. Teller, J. Chem. Phys. **21**, 1087 (1953)

[42] C. Dellago, P.G. Bolhuis, D. Chandler, J. Chem. Phys. **110**, 6617 (1999)

[43] P.L. Geissler, C. Dellago, D. Chandler, J. Phys. Chem. B **103**, 3706 (1999)

[44] P.G. Bolhuis, C. Dellago, D. Chandler, Proc. Natl. Acad. Sci. **97**, 5877 (2000)

[45] L. Onsager, Phys. Rev. **54**, 554 (1938)

[46] R. Du, V.S. Pande, A.Y. Grosberg, T. Tanaka, E.I. Shakhnovich, J. Chem. Phys. **108**, 334 (1998)

[47] D. Ryter, Physica A **142**, 103 (1987)

[48] A. Berezhkovskii, A. Szabo, J. Chem. Phys. **125**, 104902 (2006)

[49] M.M. Klosek, B.J. Matkowsky, Z. Schuss, Ber. Bunsenges. Phys. Chem. **95**, 331 (1991)

[50] E. Pollak, A.M. Berezhkovskii, Z. Schuss, J. Chem. Phys. **100**, 334 (1994)

[51] P. Talkner, Chem. Phys. **180**, 199 (1994)

[52] G. Hummer, J. Chem. Phys. **120**, 516 (2004)

[53] D. Moroni, P.R. ten Wolde, P.G. Bolhuis, Phys. Rev. Lett. **94**, 235703 (2005)

[54] C. Dellago, P.G. Bolhuis, Mol. Sim. **30**, 795 (2004)

[55] P.L. Geissler, C. Dellago, J. Phys. Chem. B **108**, 6667 (2004)

[56] T.S. van Erp, D. Moroni, P.G. Bolhuis, J. Chem. Phys. **118**, 7762 (2003)

[57] T.S. van Erp, P.G. Bolhuis, J. Comp. Phys. **205**, 157 (2005)

[58] P.G. Bolhuis, Proc. Natl. Acad. Sci. **100**, 12129 (2003)

[59] P.G. Bolhuis, Biophys. J. **88**, 50 (2005)

[60] D. Moroni, P.G. Bolhuis, T.S. van Erp, J. Chem. Phys. **120**, 4055 (2004)

[61] R.J. Allen, P.B. Warren, P.R. ten Wolde, Phys. Rev. Lett. **94**, 018104 (2005)

[62] R.J. Allen, D. Frenkel, P.R. ten Wolde, J. Chem. Phys. **124**, 024102 (2006)

[63] R.J. Allen, D. Frenkel, P.R. ten Wolde, J. Chem. Phys. **124**, 194111 (2006)

[64] P.L. Geissler, C. Dellago, D. Chandler, J. Hutter, M. Parrinello, Science **291**, 2121 (2001)

[65] M. Grünwald, P.L. Geissler, C. Dellago, J. Chem. Phys. **127**, 154718 (2007)

[66] P.G. Bolhuis, J. Juraszek, Proc. Nat. Acad. Sci. **103**, 15859 (2006)

[67] P.G. Bolhuis, J. Phys.: Condens. Matter **15**, S113 (2003)

[68] H.A. Posch, W.G. Hoover, Phys. Rev. A **39**, 2175 (1989)

[69] J.D. Weeks, D. Chandler, H.C. Andersen, J. Chem. Phys. **54**, 5237 (1971)

[70] D. Frenkel, B. Smit, *Understanding Molecular Simulation* (Academic Press, New York, 2002)

[71] C. Dellago, H.A. Posch, W.G. Hoover, Phys. Rev. B **53**, 1485 (1996)

[72] C. Dellago, W.G. Hoover, H.A. Posch, Phys. Rev. E **65**, 056216 (2002)

[73] S.H. Tolbert, A.P. Alivisatos, Annu. Rev. Phys. Chem. **46**, 595 (1995)

[74] L. Manna, E.C. Scher, A.P. Alivisatos, Journal of Cluster Science **13**, 521 (2002)

[75] D. Zaziski, S. Prilliman, E.C. Scher, M. Casula, J. Wickham, S.M. Clark, A.P. Alivisatos, Nano Letters **4**, 943 (2004)

[76] M.D. Knudson, Y.M. Gupta, A.B. Kunz, Phys. Rev. B **59**, 11704 (1999)

[77] S. Limpijumnong, W.R.L. Lambrecht, Phys. Rev. Lett. **86**, 91 (2001)

[78] S. Limpijumnong, W.R.L. Lambrecht, Phys. Rev. B **63**, 104103 (2001)

[79] S. Limpijumnong, S. Jungthawan, Phys. Rev. B **70**, 054104 (2004)

[80] M. Wilson, P.A. Madden, J. Phys.: Condens. Matter **14**, 4629 (2002)

[81] H. Sowa, Acta Cryst. A **57**, 176 (2001)

[82] H. Sowa, Acta Cryst. A **61**, 325 (2005)

[83] D. Zahn, Y. Grin, S. Leoni, Phys. Rev. B **72**, 064110 (2005)

[84] H. Sowa, Solid State Sciences **7**, 1384 (2005)

[85] E. Rabani, J. Chem. Phys. **116**, 258 (2002)

[86] M.P. Allen, D.J. Tildesley, *Computer Simulation of Liquids* (Clarendon Press, Oxford, 1991)

[87] J.J. Shiang, A.V. Kadavanich, R.K. Grubbs, A.P. Alivisatos, J. Phys. Chem. **99**, 17417 (1995)

[88] S.M. Clark, S.G. Prilliman, C.K. Erdonmez, A.P. Alivisatos, Nanotechnology **16**, 2813 (2005)

[89] E. Rabani, J. Chem. Phys. **115**, 1493 (2001)

[90] S. Kodiyalam, R.K. Kalia, H. Kikuchi, A. Nakano, F. Shimojo, P. Vashishta, Phys. Rev. Lett. **86**, 55 (2001)

[91] S. Kodiyalam, R.K. Kalia, A. Nakano, P. Vashishta, Phys. Rev. Lett. **93**, 203401 (2004)

[92] R. Martoňák, L. Colombo, C. Molteni, M. Parrinello, J. Chem. Phys. **117**, 11329 (2002)

[93] A.P. Alivisatos, J. Phys. Chem. **100**, 13226 (1996)

[94] J.Z. Jiang, J.S. Olsen, L. Gerward, D. Frost, D. Rubie, J. Peyronneau, Europhys. Lett. **50**, 48 (2000)

[95] J.E. Jørgensen, J.M. Jakobsen, J.Z. Jiang, L. Gerward, J.S. Olsen, J. Appl. Cryst. **36**, 920 (2003)

[96] L. Manna, L.W. Wang, R. Cingolani, A.P. Alivisatos, J. Phys. Chem. B **109**, 6183 (2005)

[97] M. Grünwald, E. Rabani, C. Dellago, Phys. Rev. Lett. **96**, 255701 (2006)

[98] B.J. Morgan, P.A. Madden, Phys. Chem. Chem. Phys. **8**, 3304 (2006)

[99] A.J. Kulkarni, M. Zhou, K. Sarasamak, S. Limpijumnong, Phys. Rev. Lett. **97**, 105502 (2006)

[100] S.E. Baltazar, A.H. Romero, J.K. Rodríguez-López, H. Terrones, R. Martoňák, Comput. Mater. Sci. **37**, 526 (2006)

[101] M. Grünwald, C. Dellago, Mol. Phys. **104**, 3709 (2006)

[102] A. Malevanets, R. Kapral, J. Chem. Phys. **112**, 7260 (2000)

[103] C.W. Gardiner, *Handbook of Stochastic Methods: for Physics, Chemistry and the Natural Sciences* (Springer, Berlin, 2004)

[104] B.J. Morgan, P.A. Madden, J. Phys. Chem. C **111**, 6724 (2007)

[105] Z.W. Wang, L.L. Daemen, Y.S. Zhao, C.S. Zha, R.T. Downs, X.T. Wang, Z.L. Wang, R.J. Hemley, Nature Materials **4**, 922 (2005)

[106] F. Huang, J.F. Banfield, J. Am. Chem. Soc. **127**, 4523 (2005)

[107] Y. Liebhold-Ribeiro, D. Fischer, M. Jansen, Angew. Chem. Int. Ed. **47**, 4428 (2008)

[108] S.H. Tolbert, A.B. Herhold, L.E. Brus, A.P. Alivisatos, Phys. Rev. Lett. **76**, 4384 (1996)

[109] B.J. Morgan, Phys. Rev. B **78**, 024110 (2008)

[110] S.E. Boulfelfel, D. Zahn, Y. Grin, S. Leoni, Phys. Rev. Lett. **99**, 125505 (2007)

[111] R. Martoňák, C. Molteni, M. Parrinello, Comput. Mater. Sci. **20**, 293 (2001)

[112] C. Molteni, R. Martoňák, M. Parrinello, J. Chem. Phys. **114**, 5358 (2001)

[113] R. Martoňák, D. Donadio, A.R. Oganov, M. Parrinello, Nature Materials **5**, 623 (2006)

[114] S. Leoni, R. Ramlau, K. Meier, M. Schmidt, U. Schwarz, Proc. Nat. Acad. Sci. **105**, 19612 (2008)

[115] M. Grünwald, C. Dellago, P.L. Geissler, J. Chem. Phys. **129**, 194101 (2008)

[116] J.M. Besson, J.P. Itié, A. Polian, G. Weill, Phys. Rev. B **44**, 4214 (1991)

[117] L.R. Dodd, D.N. Theodorou, Mol. Phys. **72**, 1313 (1991)

[118] F. Calvo, J.P.K. Doye, Phys. Rev. B **69**, 125414 (2004)

[119] T.L. Hill, J. Am. Chem. Soc. **79**, 4885 (1957)

[120] T.L. Hill, *Thermodynamics of Small Systems* (Dover, New York, 1994)

[121] H.B. Callen, *Thermodynamics and an Introduction to Thermostatistics*, 2nd edn. (John Wiley & Sons, New York, 1985)

[122] P. Buffat, J.P. Borel, Phys. Rev. A **13**, 2287 (1976)

[123] D. Antonangeli, M. Krisch, D.L. Farber, David G. Ruddle, G. Fiquet, Phys. Rev. Lett. **100**, 085501 (2008)

[124] E. Ott, *Chaos in dynamical systems* (Cambridge University Press, 1993)

[125] K. Sarasamak, A.J. Kulkarni, M. Zhou, S. Limpijumnong, Phys. Rev. B **77**, 024104 (2008)

[126] L.E. Brus, J.A.W. Harkless, F.H. Stillinger, J. Am. Chem. Soc. **118**, 4834 (1996)