



**University of
Leicester**

Multiscale Methods for Nanoengineering

Kenny Jolley

Thesis submitted to the University of Leicester
in accordance with the requirements for the degree of
Doctor of Philosophy in the Faculty of Engineering

Supervisor: Dr. S. P. A. Gill

University of Leicester Engineering Department
Mechanics of Materials Research Group

Date of submission: April 2009

Multiscale Methods for Nanoengineering

Kenny Jolley

Mechanics of Materials Research Group,
Department of Engineering, University of Leicester,
Leicester, LE1 7RH, UK

Abstract

This thesis is presented in two sections. Two different multiscale models are developed in order to increase the computational speed of two well known atomistic algorithms, Molecular Dynamics (MD) and Kinetic Monte Carlo (KMC).

In Section I, the MD method is introduced. Following this, a multiscale method of linking an MD simulation of heat conduction to a finite element (FE) simulation is presented. The method is simple to implement into a conventional MD code and is independent of the atomistic model employed. This bridge between the FE and MD simulations works by ensuring that energy is conserved across the FE/MD boundary. The multiscale simulation allows for the investigation of large systems which are beyond the range of MD. The method is tested extensively in the steady state and transient regimes, and is shown to agree with well with large scale MD and FE simulations. Furthermore, the method removes the artificial boundary effects due to the thermostats and hence allows exact temperatures and temperature gradients to be imposed on to an MD simulation. This allows for better study of temperature gradients on crystal defects etc.

In Section II, the KMC method is introduced. A continuum model for the KMC method is presented and compared to the standard KMC model of surface diffusion. This method replaces the many discrete back and forth atom jumps performed by a standard KMC algorithm with a single flux that can evolve in time. Elastic strain is then incorporated into both algorithms and used to simulate atom deposition upon a substrate by Molecular Beam Epitaxy. Quantum dot formation due to a mismatch in the lattice spacing between a substrate and a deposited film is readily observed in both models. Furthermore, by depositing alternating layers of substrate and deposit, self-organised quantum dot super-lattices are observed in both models.

Keywords:

Molecular dynamics; Kinetic Monte Carlo; Finite Element Method; boundary conditions; heat transfer; multiscale modelling; atomistic/continuum coupling.

Acknowledgements

I would like to thank my supervisor Dr Simon Gill and my family for their continued support throughout my PhD.

Contents

ABSTRACT	I
ACKNOWLEDGEMENTS	II
1 INTRODUCTION.....	1
1.1 CONVENTIONAL MODELLING TECHNIQUES	9
(a) Ab initio	9
(b) Molecular Dynamics (MD)	10
(c) Statistical Mechanics	12
(d) Kinetic Monte Carlo (KMC)	13
(e) Finite Element Method (FEM).....	14
1.2 INTRODUCTION TO MULTISCALE MODELLING	15
1.3 PROBLEMS WITH CONCURRENT MODELS	16
 SECTION I	 19
2 INTRODUCTION TO MD SIMULATIONS	20
2.1 MATHEMATICS OF MD	20
(a) Newtonian dynamics.....	20
(b) Hamiltonian equations of motion.....	21
2.2 INTERATOMIC POTENTIALS.....	21
(a) Lennard-Jones	22
(b) Other potential energy functions	25
2.3 INTEGRATING THE EQUATIONS OF MOTION.....	27
2.4 DIMENSIONLESS UNITS	28
2.5 THE MEANING AND MEASUREMENT OF TEMPERATURE	29

2.6	TEMPERATURE CONTROL	30
(a)	Rescale the speeds	31
(b)	Velocity exchange.....	31
(c)	The Langevin thermostat	31
(d)	Nosé-Hoover thermostat	34
(e)	Berendsen thermostat.....	35
2.7	BOUNDARY CONDITIONS.....	35
2.8	LATTICE VIBRATIONS.....	37
(a)	Phonons.....	39
(b)	Thermal conduction by phonons	40
2.9	VALIDITY OF MD	42
2.10	METHODS FOR INCREASING COMPUTATION SPEED.....	43
2.11	SUMMARY	44
3	NON-EQUILIBRIUM MOLECULAR DYNAMICS (NEMD).....	45
3.1	INTRODUCTION.....	45
3.2	NEMD SIMULATIONS OF A 3D ROD.....	46
3.3	PROBLEMS OF PHONON REFLECTION, MISMATCH AND THE KAPITZA EFFECT.....	51
3.4	DIVERGENCE OF THERMAL CONDUCTIVITY	54
3.5	NEMD SIMULATIONS WITH EXTRA FEEDBACK CONTROL	55
4	A COUPLED ATOMISTIC/CONTINUUM MODEL FOR HEAT FLOW	60
4.1	THE CONTINUUM HEAT CONDUCTION MODEL	60
4.2	THERMAL BOUNDARY CONDITIONS FOR STEADY STATE COUPLED ATOMISTIC/CONTINUUM SIMULATION	63
4.3	THERMAL BOUNDARY CONDITIONS FOR TRANSIENT COUPLED ATOMISTIC/CONTINUUM SIMULATION	65
4.4	RESULTS	66
(a)	Rapid discrete boundary variations.....	67
(b)	Harmonically varying boundary	70
(c)	Simulation of a composite material	71
	DISCUSSION AND CONCLUSIONS FOR SECTION I.....	74

SECTION II	77
5 INTRODUCTION TO SURFACE DIFFUSION AND THE KMC METHOD	78
5.1 QUANTUM DOT FORMATION	79
5.2 THE KMC ALGORITHM	82
5.3 CONTINUUM KMC MODEL	87
6 KMC SIMULATIONS	89
6.1 SINGLE COMPONENT SIMULATIONS	89
6.2 TWO COMPONENT SIMULATIONS	96
7 SURFACE DIFFUSION WITH ELASTIC STRAIN	110
7.1 ALGORITHM	110
7.2 SURFACE ROUGHENING	112
7.3 KMC SIMULATIONS OF MOLECULAR BEAM EPITAXY (MBE)	119
CONCLUSION FOR SECTION II	125
OVERALL CONCLUSION	126
REFERENCES	127
APPENDIX	135
A DERIVATION OF HEAT FLUX IN 1D	135
B HAMILTONIAN FORMULATION	137
C ELLIPTICAL VOID IN A TWO COMPONENT MATERIAL	140

1

Introduction

Nano-engineering and multiscale modelling

Technological advances have often required the fabrication of smaller and smaller devices. For instance, since 1965, the number density of transistors in computer hardware has doubled every 2 years [1]. For this trend to continue, the feature size of each component will have to enter the nanometre scale. The latest Intel Core 2 Duo processors have architecture feature sizes of 45nm. At the nanometre scale, a great many new physical properties become more important. The materials can no longer be considered as continuous with constant bulk properties, as in many cases over 50% of the atoms are on a boundary. The behaviour of these small devices depends strongly on the atomic structure. Furthermore, quantum effects become more significant at this scale. One cannot simply scale down existing devices; they must be redesigned to function at the nanometre scale.

On the nanometre scale, one cannot consider a material as a uniform bulk with constant properties such as the Young's modulus etc. The surface atoms on a real crystal have different properties and are mobile and can reorder themselves on the surface to find a lower energy configuration. Surface reconstruction is such a process and is illustrated with the (111) surface of silicon in figure 1.1 below.




Fig. 1.1 High Resolution Electron Microscope (HREM) image of a (7x7) reconstruction on the Si(111) surface, [2].

Surface structures such as these are useful templates for the growth of 3D self-organised structures, because the deposited atoms are more likely to nucleate on points of minimal energy on the surface. An example of this is presented by Bansmann et al, [3], where the (111) surface of gold formed a ‘zigzag’ reconstruction of alternating Face Centre Cubic (FCC) and Hexagonal Close Packed (HCP) crystal structures as shown in figure 1.2. When cobalt is deposited onto the surface, it is found to form self assembled atomic clusters at the vertices of the zigzags as shown in figure 1.3.




Fig. 1.2 STM image of zigzag reconstruction on Au(111) surface showing commensurate FCC and HCP domains, [3].




Fig. 1.3 Image of self assembled cobalt clusters formed at the vertices of a zigzag surface reconstruction of Au(111), [3].

These reactive sites can also be manufactured. A stepped crystal surface can be created by cutting the crystal at a small angle relative to the surface. The edges of these steps become a lower energy site for adatoms to attach. Figures 1.4a and 1.4b illustrate this method.




Fig. 1.4 (a) Diagram of the fabrication of a stepped surface by cutting a crystal at a small angle, θ , to a crystal face. (b) Self assembled clusters form at the edges of the steps on a surface [4]. The sites at the edge of each step offer a lower energy than the terraces; therefore, as the experiment shows, these clusters preferentially form here.

The above examples of self assembly depend heavily on the underlying atomistic processes. An example of the significance of atomistic effects is the deposition of germanium atoms upon a silicon substrate by molecular beam epitaxy (MBE). One would assume a uniform layer of germanium would form upon the substrate; however as shown in figure 1.5 this is not the case.




Fig. 1.5 Germanium atoms deposited upon a silicon surface form self assembled atomic clusters. The colours represent the gradient of the surface, [5].

Germanium and silicon crystals have considerably different lattice spacing. This introduces an interfacial strain between the two materials. The surface will naturally evolve to minimise its energy. The large elastic strain energy is reduced by splitting the deposited germanium into discrete islands. This increases the surface energy, and hence the resulting size of the islands depends on the competition between strain energy and surface energy. This self-assembly process is known as Stranski-Krastanov growth. If the islands are small enough (below 20nm in diameter) they are known as quantum dots (QDs) due to their unique electronic properties. These are considered to be useful building block for quantum devices such as: single photon emitters for quantum computing and quantum cryptography, QD light emitting diodes (LEDs) for displays and more efficient photovoltaic devices.

The strain in a crystal lattice is also very important. The strain can affect bulk properties like thermal [6] and material [7] transport, which will undoubtedly affect the response of a strained crystal subject to thermal effects. In addition, atoms in a state of high strain are equivalently in a high energy state. A material will relax this strain in order to lower its internal energy if it is able to do so. Figure 1.7.a shows the result of deposition of multiple alternating layers of InP and InAs/GaAs by molecular beam epitaxy (MBE). The first layer of InAs/GaAs forms quantum dots as before. However when more InP is deposited on top, this capping layer is strained by the dots underneath. The next layer of InAs/GaAs then preferentially forms quantum dots in areas where the substrate is already strained as this is the lowest energy site. Therefore, the communication of the strain through the crystal causes the vertical stacking of QDs in this nano scale device, forming a so-called QD super-lattice. Another good example of the importance of material transport and mismatched strain is the growth of InAs nanowiskers on a GaAs substrate [8] by diffusion of In through a gold particle on the GaAs substrate.

The above pictures in figures 1.3 to 1.5 are examples of material transport. Thermal transport is also highly important for nano scale devices as the properties (such as electrical resistance) of these devices depend on the temperature of the system. Figure 1.6 illustrates the simulation parameters of a computer simulation of silicon in which nanopores of various sizes and densities are present. The conductivity of silicon is measured and the effect of nanopores on the thermal conductivity of silicon [9] is investigated. This is a good example of the value of computer simulations. Setting up and running these simulations is considerably simpler than trying to fabricate such samples in the lab. The authors find the introduction of

nanopores vastly reduces the thermal conductivity of silicon from a bulk value of $149 \text{ Wm}^{-1}\text{K}^{-1}$ to $7.1 - 0.6 \text{ Wm}^{-1}\text{K}^{-1}$ depending on pore size and density. They find that making the pores bigger with fixed spacing reduces the thermal conductivity, and that reducing the pore density (increasing spacing) will increase it. The reduction in thermal conductivity is thought to arise from the reduction in the channels for phonon transport and the increased phonon scattering at the pore surfaces.




Fig. 1.6 A schematic of the simulation of the effect of nanopores on the thermal conductivity of silicon. (a) Structure of porous silicon (blue balls) which is passivated with hydrogen atoms (green). (b) Structural variables of the nanopores with a pore diameter, d_p , and pore spacing, d_s [9].

Devices

Figure 1.7 shows two examples of nanoscale devices in which quantum dots are an integral part. Due to the different conduction band energy of different materials, a quantum dot appears as a large potential well to an electron. Quantum mechanics predicts that there are several bound states for the electron that depend on the shape and size of the quantum dots (which controls the shape and size of the potential well) as well as the depth of the potential well. The electronic properties are similar to large atoms in that there are several discrete energy levels and that bound electrons can absorb and emit photons at specific wavelengths. Quantum dots offer a way of tailoring these properties by altering the material composition, size and shape of the QD. Hence emitters and detectors for specific wavelengths can be created. Furthermore an isolated quantum dot will only absorb and emit single photons at a time. Therefore these are useful components in quantum computing and quantum cryptography applications. Even for use in simple detectors, these QDs have a higher

response time, higher temperature operation, higher light coupling to normal incidence light than conventional devices. Computer modelling allows for easy design and testing of the multiple parameters of the manufacturing process of nano devices which allows for easy tailoring of these devices to specific applications.





Fig. 1.7 Two examples of nanoscale devices are shown. (Left) A Single-photon generator/detector with InAs quantum dots [10]. (Right) A multi-stack InAs/GaAs quantum dot infrared photodetector is shown [11].

Computer modelling

At the nanometre scale, not only are interfaces, quantum effects and strain important, also thermal effects are significant. To illustrate this, an example of a computer model of a nano-indentation experiment is used to show the importance of temperature [12]. This also highlights some of the difficulties that arise in computer simulations.

(a) Strong (Nose-Hoover) Thermostat



(b) Weak (Berendsen) Thermostat




Fig. 1.8 Computer simulation of a nano-indentation experiment, [12]. The pictures show only the atoms that are displaced by the nano-indenter along with the top and bottom planes of the simulation. In both cases the thermostat is applied to all of the atoms in the simulation. (a) The atoms are thermostatted using the strong Nosé-Hoover thermostat. Here the dynamics of the atoms have been corrupted the leading to an incorrect result. (b) The weaker Berendsen thermostat controls the temperature without significantly corrupting the atomistic dynamics. The experimentally observed result is reproduced.

Figure 1.8 shows the results of a large scale MD simulation in which an indenter is pushed into a surface. Only the atoms that are displaced by the nano-indenter along with the top and bottom planes are shown. Two thermostating methods are compared, the strong Nosé-Hoover method and a weaker Berendsen method (see chapter 2.6). In each case, the thermostat is applied to all the atoms in the simulation. A thermostat is required to regulate the temperature of the simulation because the nano-indenter does mechanical work upon the system as it is pressed into the surface. The thermostats, then act on the atoms in order to remove this excess energy. If this was not done, the system temperature would increase significantly and reflections from the boundaries would adversely effect the simulation.

Figure 1.8 shows that one must be careful in applying thermostats to computer simulations. A thermostat alters the trajectories of the atoms which can have vast consequences on the system. The strong Nosé-Hoover thermostat has corrupted the dynamics of the atoms in the simulation, yielding an incorrect result.

In order to understand the physical properties of materials and devices at the nanometre scale, it is easier to simulate these small systems on computers. A computer can

track the trajectories of every atom if required, and, if storage space exists, any information about the system can be obtained. This is not possible with real systems; as such localised measurements are very difficult. Furthermore, identical systems can be subjected to different environmental conditions on a computer to investigate their effects, whereas in the real world nano-scale experiments are not precisely repeatable.

The major drawback with computer simulations is that it takes too much time to simulate enough atoms, due to the vast amount of degrees of freedom involved. A nanometre scale device can contain many billions of atoms. Even with the most powerful supercomputers, simulations are limited to a few million atoms over timescales of a few nanoseconds. However, as has been demonstrated above, one cannot ignore the atomistic effects at this scale. Therefore, in order to simulate the formation of self organised structures as seen in figures 1.4 and 1.5, a different approach is required.

A solution is to use multi-scale modelling. This is essentially a hybrid of existing discrete, stochastic and continuum numerical models with an artificial interface between them. Here the bulk area of a material (in which physics is well understood and predictable) is simulated with existing continuum methods. Surfaces, defects and boundaries are simulated with atomistic models. These separate models are then linked together in a way which does not adversely affect the results. This linkage is usually the major problem to be overcome when developing multiscale methods.

Two such multi-scale methods are developed in this thesis. In section I, heat conduction is modelled using the Molecular Dynamics (MD) method. This method is introduced in detail along with many methods of controlling MD simulations in chapter 2. Non-equilibrium molecular dynamics (NEMD) simulations are introduced in chapter 3 along with a steady state MD/continuum coupling method. In chapter 4, this concurrent MD/continuum thermal coupling model is developed in order to enable fully transient boundary conditions. A range of simulations are performed for extremely testing cases and results presented in 4.4.

In Section II multiscale computer simulation methods are introduced for use in surface diffusion problems. This section begins at chapter 5 with a detailed introduction to quantum dot formation. A detailed overview of the KMC method and the continuum KMC method are presented in 5.2 and 5.3 respectively. In chapter 6 both simulation methods are applied to

single and two component ellipses that are allowed to evolve by surface diffusion. This was done to validate the continuum method. Elastic strain is added to both algorithms in chapter 7. Simulations of surface roughening are then performed in order to identify optimal conditions for the atom deposition simulations. Lastly the formation of quantum dots is simulated. The deposition of alternating layers is simulated resulting in the formation of self organised vertically aligned quantum dots.

Before proceeding further, some standard numerical methods for simulating material at particular length and time scales are introduced as well as a short introduction to multiscale methods.

1.1 Conventional modelling techniques

There are many existing methods of computer simulation, each with very specific applications. For simulations of a single molecule, a first principles approach is best as the quantum effects at this scale are significant. Crystals of a few nanometres are best simulated with MD if computation time is an issue, as the quantum effects are minimal and the extra computational effort of ab initio methods is not required. In order to simulate large time scales at the nanometre scale, KMC is the best model. Here the motion of the individual atoms is not important, just the structure, which KMC simulates well. Finally, for systems with large length and time scales, material properties can be expressed by constitutive laws and the system is solved via the finite element (FE) method. Statistical mechanics also applies when there are many particles interacting for a large time. This section provides an overview of each method.

(a) Ab initio

The Ab initio (also known as ‘first principles’) method uses the rules of quantum mechanics to calculate energies and potentials of a group of atoms. The simulation proceeds by solving the many body Schrödinger equation. The only input is the atomic number of the elements involved. A numerical solution of the many-body wave function for more than a few particles is not viable, even on today’s largest supercomputers, due to the large storage requirements and vast amount of calculations required.

An accepted approximation is Density Functional Theory (DFT), in which the individual particles are replaced with a particle density, see [13] for review. The wave function of the DFT system now depends on, the electron density, and the locations of the atomic nuclei. This significantly reduces the complexity at the cost of accuracy. Even with DFT, ab initio calculations are limited to about 100 atoms.




Fig 1.9 An example of the results from a first principles calculation of a potential energy surface (PES) of a material [14]. The image shows the (2×1) view of the potential energy surface of a palladium atom on the (110) surface of dehydrated γ -alumina. The violet balls are oxygen and the yellow balls are aluminium.

Figure 1.9 shows the potential energy of a surface computed using Ab initio methods. These results can then be used to calculate the inter-atomic potentials and energy barriers to surface diffusion which can then be used in MD and KMC simulations.

(b) Molecular Dynamics (MD)

The MD method is presented in detail in chapter 2; however a brief overview is presented here.

A material is considered as a collection of classical particles (atoms), interacting via a pre-determined potential energy function, $V(x)$. The MD algorithm is quite simple; for each atom, the net force, $\underline{F}$, is calculated from the derivative of the potential energy, given the inter-atomic distances of the neighbouring atoms. Then each atom is moved according to Newton's second law of motion,

$$\underline{F} = m\underline{\ddot{x}} = -\frac{\partial V}{\partial \underline{x}} \quad (1.1)$$

where m is the mass of the particle, $\ddot{x}$ is its acceleration and $V(x)$ is the potential energy. The atomic structure and material properties are all contained within the potential energy function. This means that correct determination of this function and its constants is essential for correct simulation of materials.

It is possible to simulate several million atoms for a few nanoseconds using MD using a parallel code on a powerful supercomputer. A common example is the simulation of a nano-indentation experiment. A real indenter is typically of the order of 50 – 100nm, which creates an imprint on the substrate of several hundreds of nanometres. In general, the atoms that make up the indenter are not simulated. In order to reduce the computational requirements of the simulation, an additional potential term is added to the atoms of the simulation to mimic the indenter, as shown in figure 1.10.




Fig. 1.10 Example result of an MD simulation of a nano indentation experiment [15].

A major problem with these simulations is the boundaries. Fixing the boundaries will artificially stiffen the material and cause reflections of elastic waves. Freely moving boundaries cannot be used as the indenter would merely push the whole material downwards. Periodic boundaries are better, but the indenter does mechanical work on the substrate and hence raises its temperature. The heat generated must be conducted away from the indentation site; periodic boundaries will keep the substrate artificially hot. Section 1 of this thesis is particularly concerned with developing appropriate thermal boundary conditions for MD

simulations. A method of coupling MD and FE boundaries is derived to show how such issues can be addressed.

(c) Statistical Mechanics

For systems with a large number of particles, the laws of probability, quantum mechanics and thermodynamics are used to derive equations of state for the given system. This method has applications in many areas such as chemical reactions, photons in a cavity and population inversion in lasers and gases, to name a few. Statistical mechanics is useful for providing initial conditions for MD simulations as it gives the distribution of the speeds of atoms in a material. The distribution in atom velocities is not uniform, but rather fits a skewed profile known as the Maxwell Boltzmann distribution:

$$f(v) = 4\pi \left(\frac{m}{2\pi k_b T} \right)^{3/2} v^2 \exp \left[\frac{-mv^2}{2k_b T} \right], \quad (1.2)$$

where v , is the atomic velocity, m is the atomic mass, k_b is the Boltzmann constant and T is the absolute temperature. This is shown in figure 1.11.




Fig. 1.11 Plot of the Maxwell Boltzmann distribution of atom speeds for argon at 40K.

The peak of the curve is the most probable atom speed of:

$$v_p = \sqrt{\frac{2k_b T}{m}} \quad (1.3)$$

This is lower than the average speed, $\langle v \rangle$, due to the skewed nature of the curve. This average speed is given by:

$$\langle v \rangle = \sqrt{\frac{8k_b T}{m\pi}} \quad (1.4)$$

The RMS speed is given by:

$$v_{rms} = \sqrt{\frac{3k_b T}{m}} \quad (1.5)$$

It is found from MD simulations that the atoms fit this curve very well. All MD simulations at finite temperature discussed in this work are initialised with speeds picked from the Maxwell Boltzmann distribution.

(d) Kinetic Monte Carlo (KMC)

For long time scale processes like surface diffusion, the MD method is too slow. Typically one surface hop of an atom takes many thousands of atom vibrations, which is beyond the reach of MD.

In the KMC method, the vibrational motion of the atoms is not simulated, but merely represented by a temperature. A list of all the possible surface hopping events is generated along with the rates of these events. The rates can be calculated from the energy barrier between sites. One event is picked at random, depending on its rate, and executed. The time is then updated based on the rate of the event selected. This cycle is repeated many times to simulate the evolution of a surface. KMC simulations fall into two categories:

On-lattice

On-lattice KMC is only used for condensed phase simulations, mainly at zero strain. Each atom site and its neighbours are defined at the start of the simulation. The events are also defined at the start as hops to a vacant neighbour site, with the rates dependant on the initial state. The rate of each event can depend on the number and type neighbouring atoms.

Effects of elastic strain can be added by defining the starting lattice in a strained position. Then the rates of each event can be calculated from the ‘real’ initial energy of each atom.

One can also define the rates as dependant on the energy barrier between sites. Note that this energy barrier is overestimated in two-dimensions, where an atom cannot move around an atom to pass it.

Off-lattice

No lattice sites are defined here. Each atom is dragged from its current position to see if there is another stable minimum nearby, and the jump rate is dependant on the energy barrier. This method allows for dislocations, cracks and large scale deformations in strained crystals. In this thesis, these effects are not modelled in the KMC simulations and hence the simpler on-lattice model is used. The KMC method is discussed in more detail in section 2.

(e) Finite Element Method (FEM)

The finite element method is used to solve partial differential equations (PDE’s) approximately, by discretizing the functions on a mesh. The Galerkin method reduces the problem to the solution of a set of matrix equations.

To simulate a material, this method requires input of all material bulk properties, such as thermal conductivity, Young’s modulus and density. Also constitutive laws of all the physical processes which one wants to simulate need to be pre-defined in terms of differential equations. Then FEM is used to solve these PDE’s subject to user-given boundary conditions. In many situations (e.g. quantum dots on surfaces, figure 1.1) one cannot ignore the atomic structure, and hence finite element computations alone become over simplified.

An example of multi-physics FEM is shown in figure 1.12, where electrostatic forces control a pair of micro tweezers. The solution depends on solving a set of coupled PDEs for the electro-magnetic and elastic fields in the material.




Fig. 1.12 Finite element model of an electrostatically actuated comb drive used to open and close a pair of micro-tweezers. Electrostatic forces attract the combs to each other. The colours show the electric field, blank areas are the drive. [16].

FEM is capable of simulating objects from microns to kilometres. There is no maximum size limit, only a limit to the amount of detail. However, FEM cannot correctly simulate materials at the nanometre scale, where atomistic effects are significant.

1.2 Introduction to multiscale modelling

A solution to the problem of finite computational resources is multiscale modelling. It is not necessary to simulate the oscillations of every atom in a uniform bulk region because the results are predictable by using other methods such as finite elements. Therefore, for a given system, the atomistic simulation method should only be applied where it is needed (on the boundaries, or near defects such as grain boundaries, cracks, voids etc). Continuum approaches such as finite elements can then be used in the linear bulk regions.

The problem with this is that there needs to be a way of connecting these different simulation methods without adversely affecting the simulation results. Any interface between computer simulation methods will be an artificial one.

Figure 1.13 shows the typical length and time scales that are available for a range of simulation methods.




Fig. 1.13 Diagram comparing the length and time scales typically available for a range of computer simulation methods.

1.3 Problems with concurrent models

Dynamics

Finite temperature multi-scale modelling is not as simple as one might assume. The requirements of the continuum far field depend on the nature of the simulation, generally either sampling or dynamics. If the purpose of simulation is sampling of near equilibrium or steady state quantities, then typically only slowly-changing thermodynamical or statistical quantities are of interest and inertial effects are small. Rapid changes occur in truly dynamic situations such as fast fracture. Finite temperature simulations are complicated by the reflection of high frequency phonons from the interface between the atomistic and continuum regions. This leads to energy trapping and localized heating [15]. For the simulation of a single crack (figure 1.14a), for example, one simply needs to absorb all heat from the MD region.

Correct transmission of phonons across the MD/FE interface [15][17][18] is only necessary if the far boundaries can be seen during the simulation period (e.g. simulation of

Microelectromechanical systems, MEMS) or there are two atomistic regions which need to interact dynamically via the continuum medium (e.g. rapid growth of multiple cracks, figure 1.14b). It is assumed here that absorption of phonons at the interface using diffuse boundary conditions is a sufficient requirement. This type of approach [19] has allowed the elastic boundary conditions to be specified at a position remote from the atomistic region.

In recent years, concurrent multiscale methods have been developed for crystalline solids in which the complex response of the far field is represented by a coarse-grained continuum region constructed from finite elements [15][17]-[28]. These multiscale modelling methodologies have mainly focussed on the far-field representation of the elastic field at zero or constant temperature, although a few authors have looked at the thermal far-field [18][26][27]. Some models consider isothermal problems, but these are still in the minority. Even so, constraining simulations to constant temperature can be highly restrictive, especially in cases where work is being done on the system and heat is being generated. Liu et al. [29] have recently demonstrated that MD simulations of nanoindentation are very sensitive to restricted thermostatic control (see figure 1.8). Keeping the boundary temperature constant also restricts simulations to be near to thermal equilibrium, whereas non-equilibrium conditions (e.g. temperature gradients) may be of interest.




Fig. 1.14 Schematic diagrams of multiscale computer simulations of crack growth. (a) An MD simulation of a single crack coupled to a FE region. (b) A diagram of two MD simulations of crack growth that can interact with each other via a continuum region.

Kinetics

A great deal of research has been applied to increasing the practical length and time scales of MD and KMC simulations. Rare events such as atomic hops and configuration changes of molecules occur on timescales many orders of magnitude higher than the atomic vibrations simulated by MD. Given the initial and final atomic configurations of a rare event, action derived molecular dynamics (ADMD) [30][31] yields the minimum energy pathway with the constraint of energy conservation. The alternative bond-boost method [32], in which an additional boost potential is applied close to local minima, causes faster atomic jump rates while preserving the relative rates of rare events. This method requires no previous knowledge of the kinetics of rare events.

The KMC method for surface diffusion on metals considers only the rare atom hopping events and does not resolve any atomistic vibrations. This allows for a significant increase in both the length and time scale of computer simulations compared to MD. The length and time scales can be further increased by coarse-graining the atoms into clusters. This is only done away from boundaries and areas of high gradients where the loss in accuracy is not significant to the simulation result [33][34]. This can be further enhanced by adaptively refining and coarsening the mesh on the fly [35][36]. The time-scale of KMC simulations also can be increased by allowing surface adatoms that are far from step edges and defects to execute large steps over several atomic sites at correspondingly reduced rates [37]-[39].

Nurminen et al [40] simulate the strain field of embedded QDs by patterning the substrate such that the energy barriers to diffusion depend on the (x,y) position on the substrate. Hence, no computationally demanding calculation of the strain field is required. The authors find that QDs form within an optimal temperature range of 370K – 400K.

Consistency

A computer simulation of the same system using two different techniques must yield similar results. If it does not, then one of the methods is inaccurate. Therefore the MD and FE simulations must produce the same results, for the same set of conditions. The equations to be solved by the FE method must be derived from the MD inter-atomic potential, so that the results of both methods are the same.

Section I

This section is concerned with the development of a coupled FE/MD method for simulating heat flow in non-isothermal solids. Three issues need to be addressed. Firstly, before any multiscale coupling can be considered, it is necessary to be able to precisely stipulate the boundary conditions of a non-equilibrium molecular dynamics (NEMD) simulation. This is investigated in the context of steady state heat conduction in chapter 3.1. Secondly, for the coupled model, compatibility between the material parameters in the atomistic and continuum descriptions must be ensured. And thirdly, smooth transfer of information across the interface between the two models is required. These last two issues are discussed in chapter 4.

Existing elastically coupled atomistic/continuum multiscale methods differ in their approach to implementing the effects of (constant) finite temperature. Dupuy et al. [41] and Gill et al. [20] retain the dynamics of the atomistic system in the continuum region by allowing the finite element nodes to move under inertial forces. However, the dynamics of coarse-grained nodes are not physical and cannot model systems in which there is a temperature variation. This is because equipartition demands that a system will move towards a state in which each degree-of-freedom has the same thermal energy. If the number of degrees of freedom is reduced then the thermal energy it can store in the dynamics of those degrees-of-freedom is reduced. Therefore the kinetic energy per unit volume is not conserved through the body as the number of degrees-of-freedom per unit volume is not constant. Qu et al. [19] do not resolve the thermal vibrational motion of the atoms/nodes but link an isothermal MD simulation to a quasi-static elasto-plastic continuum. As Qu et al. [19] consider an isothermal problem; the missing vibrational energy does not need to be represented in the continuum. A similar philosophy is adopted in this section, except in the non-isothermal case the missing kinetic energy in the continuum does need to be explicitly conserved. This energy is represented by the usual continuum state variable, temperature, and is allowed to evolve according to the classical laws of heat conduction whilst interacting with the atomistic medium.

Chapter 2 below introduces MD simulations in detail.

2 Introduction to MD simulations

2.1 Mathematics of MD

This chapter introduces MD simulations and explains many of the concepts and implementation issues. Many researchers have written comprehensive reviews on the subject [15][42]-[44]. Here, the equations of motion are derived from the inter-atomic potential and solved efficiently with the Verlet time integration algorithm. Methods of controlling the state of MD simulations such as thermostats are explained in section 2.6. Also explained are some methods for greatly reducing the computer time required to simulate a given system.

(a) Newtonian dynamics




Fig. 2.1 Schematic of an MD system.

As explained in the introduction, an MD simulation proceeds by solving Newton's second law for a collection of classical particles (figure 2.1). Newton's second law for each particle is given by:

$$m_i \frac{d^2 \underline{x}_i}{dt^2} = - \frac{\partial V}{\partial \underline{x}_i} \quad (2.1)$$

where, m_i and $\underline{x}_i$ are the mass and position vector of the i^{th} particle and the total potential energy $V(\underline{x}_i) = \sum_i \sum_{j \neq i} \phi(r_{ij})$ is the sum of all the interatomic potentials, $\phi(r_{ij})$,

where $r_{ij} = |\underline{x}_i - \underline{x}_j|$ is the interatomic separation. In this work, only simple pair-wise atomic interactions are considered which is reasonable for most solid metals and fluids.

For two particles interacting via Newton's second law, the trajectories of these particles can be predicted analytically. In the harmonic limit for linear springs two atoms will simply oscillate with simple harmonic motion (SHM) and an equation such as $x(t) = a \cos(t) + b \sin(t)$ can be derived. Mathematicians have been trying for centuries to solve analytically the trajectories for three or more independent particles. This is known as the three body problem. Hence such complex systems need to be solved numerically.

(b) Hamiltonian equations of motion

Typically in the literature [15][44]-[46], the Hamiltonian formulation is used to describe equations of motion of particles in a simulation. Here the system is described in terms of generalised positions and momenta. An overview of the derivation of the Hamiltonian equations is presented in appendix B.

The Hamiltonian is the sum of the kinetic and potential energies of a system. In the case of a system of particles in a one-dimensional chain, the Hamiltonian becomes:

$$H = \sum_{i=1}^N \left[\frac{p_i^2}{2m_i} + V(r_{i,i+1}) \right], \quad (2.2)$$

where, p_i , is the momentum of particle i , m , is the particle mass, r_{ij} , is the inter-atomic separation and N is the total number of particles. The first term is the kinetic energy and the second term is the potential due to interaction with neighbouring particles and finally the sum is over all the atoms in the simulation.

2.2 Interatomic potentials

The interaction between atoms is described by the potential, V . This function represents the quantum mechanical properties of the atoms which one wishes to simulate. These quantum properties are responsible for the chemical properties of the atoms such as the

bond length, bond energy and the spatial arrangement of these bonds. In order to simulate different materials like diamond, silicon or water, the potential function used will depend on the separation and orientation of the many types of atoms. Therefore the potential energy function can be a very complicated expression. The general form of the potential is given by:

$$V(r_1, r_2, \dots, r_N) = \sum_i V_1(r_i) + \sum_{i,j>i} V_2(r_i, r_j) + \sum_{i,j>i,k>j} V_3(r_i, r_j, r_k) + \dots, \quad (2.3)$$

where r_n is the radius vector of the n th particle and the functions V_m are the m -body potentials. The first term is an external potential applied to each atom. This represents an external force field such as gravity, which the whole system is subjected to. This term is usually neglected. The second term is the pair-wise interaction between atoms and the third term contains the three body interactions, etc.

(a) Lennard-Jones

Atoms consist of a positively charged atomic nucleus with negatively charged electron clouds (orbitals) surrounding it. The net charge of an atom is therefore neutral. As two atoms approach each other, an attractive force between the electron clouds occurs (atomic bonding). However as the atoms get closer, the effective screening of the positively charged nucleus is decreased, and hence electrostatic repulsion of the nuclei becomes the dominant force and prevents the atoms from getting closer. The competition between these two forces results in a minimum energy at the equilibrium bond length.

In 1924 Jones proposed the Lennard-Jones (LJ) potential to describe pair wise interactions between atoms. This model has been used to simulate a variety of processes, but gives the best results for solid argon. The LJ potential depends on the separation between atoms only and is therefore described as axially symmetric.

A system of atoms using the LJ potential will organise themselves into a hexagonal close packed (HCP) crystal, as shown in figure 2.2, when solid. The model includes a long range attractive force between atoms, which comes from a van der Waals effect from the electron clouds, and also short range repulsion due to the positive nuclei. This model can also simulate non-linear effects such as thermal expansion.




Fig. 2.2 Diagram of the Hexagonal Close Packed (HCP) lattice.

The potential depends only on the distance between atoms, and hence is axially symmetric:

$$V(r_{ij}) = 4\epsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^6 \right], \quad (2.4)$$

where σ is the collision diameter (the point at which $V(r_{ij}) = 0$) and ϵ represents the bonding energy of the atoms, which is the minimum of $V(r_{ij})$. This minimum occurs at a separation distance of $2^{1/6}\sigma$. The corresponding force, derived from the negative gradient of the potential is given by:

$$F(r_{ij}) = -\frac{\partial V}{\partial r} = \frac{24\epsilon}{\sigma} \left[2 \left(\frac{\sigma}{r_{ij}} \right)^{13} - \left(\frac{\sigma}{r_{ij}} \right)^7 \right]. \quad (2.5)$$

The potential (2.4) and corresponding force (2.5) are plotted in figure 2.3. This shows clearly that the net force is zero at the equilibrium lattice spacing of this potential.




Fig. 2.3 Plot of the Lennard-Jones potential and corresponding interatomic force.

Occasionally, in the literature, the Lennard-Jones potential is written in terms of $a = 2^{1/6}\sigma$ (the equilibrium distance) to obtain:

$$V(r_{ij}) = \varepsilon \left[\left(\frac{a}{r_{ij}} \right)^{12} - 2 \left(\frac{a}{r_{ij}} \right)^6 \right]. \quad (2.6)$$

For potentials that decay within a few atomic distances like the Lennard-Jones potential presented here, it is not necessary to calculate pair wise interactions between every possible pair of atoms in the system. The solid phase MD simulations presented in this thesis were computed using nearest neighbour interactions only, this is acceptable as interactions from the next nearest neighbours are negligible. For simulations using nearest neighbour interactions only, the computation time scales linearly with the number of particles, $O(n)$.

For fluid MD and condensed phase simulations where physical processes such as crack growth allow atoms to change neighbours during the simulation, the nearest neighbours cannot be defined. Therefore the force on each atom depends on the sum of the pair-wise forces between all atoms within a cut-off radius, R . R is usually taken to be several equilibrium bond lengths. The truncated potential energy function, $V_T(r)$, is therefore now written as:

$$V_T(r) = \begin{cases} V(r) & r \leq R, \\ 0 & r > R. \end{cases} \quad (2.7)$$

The function $V_T(r)$ is not continuous, and hence an atom entering the cut-off region will suddenly experience a force which switches on and off as it crosses the point $r = R$. A “skin” region of thickness, D , in which a smooth step-like function reduces the value of the potential at R to zero at $R+D$ is introduced. The potential is then given by:

$$V_T(r) = \begin{cases} V(r) & r \leq R \\ \frac{V(R)}{2} \left[1 - \sin(\pi(r - R - \frac{D}{2})/D) \right] & R < r \leq R + D \\ 0 & r > R + D \end{cases} \quad (2.8)$$

Cycling through every atom in the simulation in order to compute the inter-atomic displacements in order to work out which atoms are interacting at each step takes an enormous amount of time. The simplest method is to maintain a list of local atoms, which will always be much smaller than the total number of atoms in the simulation. This local list need only be updated every thousand time-steps.

(b) Other potential energy functions

Morse potential

Morse proposed a function of exponentials to model the same effects as the LJ potential. Again there are attractive and repulsive terms representing bonding, and parameters to control bond energy and length:

$$V(r_{ij}) = \varepsilon \left(e^{2\beta(\rho - r_{ij})} - 2e^{\beta(\rho - r_{ij})} \right), \quad (2.9)$$

where ε represents the bonding energy, ρ is the equilibrium bond length, and β is another inverse length scale factor. See [47] [48] for examples of the use of the Morse potential.

Multi-body potentials

A simple pair-wise potential like LJ is only capable of simulating a HCP material. For materials like carbon or silicon, which have a diamond structure, and molecules like water and hydrocarbons, which can have complicated shapes depending on the interactions between many atom types, a higher order potential energy function is required. In these cases the atom

type and spatial arrangement of neighbours and bonding angles contribute towards the net force on each atom.

In order to describe all possible bonding geometries, 4, 5 or higher order terms would be required. This quickly becomes intractable and computationally inefficient to calculate. An alternative is to “wrap up” the local geometry into two- and three-body potentials, and employ a bond order function [15] to represent the spatial arrangements of bonds around the atom. These potentials are usually short ranged, involving only nearest or next nearest neighbours.

The diamond structure is modelled by the Tersoff potential [49]. This is used to simulate covalent materials with the diamond structure such as silicon, carbon and germanium [50]-[52].

Embedded atom method (EAM)

For metallic systems, atoms are positive ions embedded in a “sea of electrons” that permeate the whole crystal [15]. There are energies associated with pushing an atom into this “electron sea”, and electrostatic pair-wise interactions from the ions. In the embedded atom method (EAM), the total potential energy function, U , takes the form of:

$$U = \sum_i G_i \left(\sum_{j \neq i} \rho_j^a(r_{ij}) \right) + \sum_{i,j>i} V_{ij}(r_{ij}), \quad (2.10)$$

where the first term represents the embedding energy of atom i in the electron gas created by its neighbouring atoms. G_i is the embedding energy of atom i , ρ_j^a is the averaged electron density of neighbouring atoms, j , and V_{ij} is the pair-wise interaction between atoms. For a review see [53].

2.3 Integrating the equations of motion

Since there is no analytical solution for N-body problems, the equations of motion need to be solved numerically. The Verlet algorithm is a simple and efficient method of doing this. Given the force (and hence the acceleration) and a set of initial conditions (positions and velocities), the subsequent motion of the particles can be calculated. The general form is that once the forces are known for a given configuration, the positions at some time later Δt are calculated. New forces are then calculated, and the process repeated.

There are several methods for solving the trajectories numerically. The Verlet method is used in the work presented here and is the most widely used method in MD. It is accurate to third order in Δt and is simple to derive and use [54].

Start with the Taylor expansion of the position variable $x(t)$:

$$x(t + \Delta t) = x(t) + (\Delta t)\dot{x}(t) + \frac{(\Delta t)^2}{2}\ddot{x}(t) + O(\Delta t)^3. \quad (2.11)$$

Now form the corresponding expansion for $x(t - \Delta t)$ by substituting $-\Delta t$ for Δt :

$$x(t - \Delta t) = x(t) - (\Delta t)\dot{x}(t) + \frac{(\Delta t)^2}{2}\ddot{x}(t) - O(\Delta t)^3. \quad (2.12)$$

Sum equations (2.11) and (2.12) and rearrange for $x(t + \Delta t)$ to obtain:

$$x(t + \Delta t) = 2x(t) - x(t - \Delta t) + (\Delta t)^2\ddot{x}(t) + O(\Delta t)^4. \quad (2.13)$$

The third order terms cancel, leading to an equation that only has fourth order errors. The force is only calculated once and the method only costs extra storage space for the previous position. Furthermore in order to obtain trajectories, there is no need to calculate the speeds. But if required, the speed at $x(t)$ can be found from:

$$\dot{x}(t) = \frac{(x(t + \Delta t) - x(t - \Delta t))}{2\Delta t} + O(\Delta t)^2. \quad (2.14)$$

The only drawback is that the calculated velocities, if needed, lag behind by one timestep. But this is usually of little consequence when dealing with averages over billions of calculations.

2.4 Dimensionless units

In order to implement the calculations on a computer, it is generally more convenient to use dimensionless units. This is done for three main reasons; firstly atoms exist at much smaller length and time scales than the macroscopic objects that the metric system was designed for. Hence the equations contain many small parameters. A computer can be more efficient when dealing with decimals that do not differ too much from unity, due to increased calculations required to handle exponents of very large or small numbers. Furthermore dimensionless units will prevent any problems caused by values becoming too small or large to be handled by the processor. Secondly the equations of motion are simplified because many of the parameters in the model can be absorbed into the units. Thirdly, solving the problem in dimensionless units allows one to apply the results to a whole class of similar problems. The most convenient set of units (for the LJ potential) are [55]:

$$\text{length:} \quad r_{Lab} = r_{MD} \sigma, \quad (2.15)$$

$$\text{energy:} \quad E_{Lab} = E_{MD} \varepsilon, \quad (2.16)$$

$$\text{time:} \quad t_{Lab} = t_{MD} \sqrt{\frac{m \sigma^2}{\varepsilon}}, \quad (2.17)$$

where subscript *lab*, refers to metric system (i.e. metres) and subscript *MD* refers to the dimensionless MD units of the program. Using these definitions, one can derive the following useful relations:

$$\text{velocity:} \quad v_{Lab} = v_{MD} \sqrt{\frac{\varepsilon}{m}}, \quad (2.18)$$

$$\text{force:} \quad F_{Lab} = F_{MD} \left(\frac{\varepsilon}{\sigma} \right), \quad (2.19)$$

$$\text{temperature:} \quad T_{Lab} = T_{MD} \left(\frac{\varepsilon}{k_b} \right), \quad (2.20)$$

$$\text{power:} \quad P_{Lab} = P_{MD} \sqrt{\frac{\varepsilon^3}{m \sigma^2}}. \quad (2.21)$$

2.5 The meaning and measurement of temperature

The thermal energy in a solid crystal is a manifestation of the atomic vibrations. A high temperature is simply equivalent to a large amplitude atomic vibration. The total amount of internal energy in a classical system is an even split between the potential energy and the kinetic energy. This is known as the equipartition theorem.

The thermal energy of a particle is $\frac{1}{2}k_bT$ per degree of freedom, where k_b is Boltzmann's constant. A degree of freedom is an independent axis of movement, for example a dimension, x , y , or z . A d -dimensional system with N particles has a total thermal energy E of:

$$E = \frac{dNk_bT}{2}. \quad (2.22)$$

This is equivalent to the total classical kinetic energy of the system of particles which is:

$$E = \frac{1}{2} \sum_{i=1}^N m_i |\dot{\underline{x}}_i - \langle \dot{\underline{x}}_i \rangle|^2, \quad (2.23)$$

where m_i is the mass of particle i , and the $\langle \dot{\underline{x}}_i \rangle$ term represents the drift velocity of the sample as a whole. In an MD simulation, the thermostats can be configured to remove any net momentum of the system. Furthermore, fixed boundaries prevent any drift in the sample entirely. Hence this term can be safely neglected. It is worth noting that in fluids the drift of matter is significant.

The temperature of a particle in terms of its velocity is found from the above expressions. For the i^{th} particle the temperature is:

$$T_i = \frac{m_i |\dot{\underline{x}}_i|^2}{dk_b}. \quad (2.24)$$

The absolute temperature at a definite point in time cannot be reliably obtained from (2.24) as the particle is oscillating and constantly changing velocity. Measurements are made in a similar way as an experimentalist would make measurements on a real experiment. The

average of a quantity, A , is the result of summing up successive measurements then dividing by the number of measurements made, M .

$$\langle A \rangle = \frac{1}{M} \sum_{n=1}^M A_n. \quad (2.25)$$

However, unlike a real experiment, one cannot assume that each measurement is independent. It is necessary to average over many oscillations in order to obtain an accurate value of the temperature. It is also useful to average over many atoms as well as many oscillations in order to obtain an accurate sample temperature. If one requires a precise measurement of the temperature over small spatial and temporal scales, then an average can be performed over multiple simulations of the same system, provided that the system is initialised randomly each time.

Statistical errors like the variance are sensitive to the correlations of the data. Therefore, one must average over many blocks of a size that is greater than the correlation, in order to get an accurate estimate of the variance [55]:

$$\sigma^2(\langle A \rangle_b) = \frac{1}{M_b - 1} \sum_{n=1}^{M_b} (A_n^2 - \langle A \rangle_b^2), \quad (2.26)$$

where M_b is the total number of blocks, A_n a typical block average, and $\langle A \rangle_b$ the overall average.

2.6 Temperature control

Simply integrating the equations of motion would cause the system to conserve energy over time. Eventually a system such as a 3D rod will attain a uniform temperature across the whole system. Hence, the equations of motion are said to be conservative. That is, the total energy E of all the atoms in a simulation will remain constant unless energy is deliberately added to the system. The total number of atoms N and the volume V are also constant. This is known as the microcanonical ensemble (constant NVE). A real material in equilibrium has a constant temperature and volume or pressure, which is represented by the canonical ensemble (constant NVT) and the isobaric-isothermal ensemble (NPT). In a simulation at

a constant energy, E , the temperature will fluctuate around a constant value. Any non-uniform initial conditions will smooth out over time to this equilibrium temperature.

Many methods have been proposed over the years to control the temperature and maintain realistic trajectories. Some of the most popular methods are presented here.

(a) Rescale the speeds

This simple idea just rescales the atom speeds in the reservoir depending upon the target temperature. ZhengXing [56] explain this and many other non-equilibrium molecular dynamics (NEMD) methods. Firstly, from a given target temperature, pre-compute what the average speed should be. Then calculate the average speed of the atoms in the simulation, and scale the speeds up or down depending on the desired average speeds. The scale factor used should depend on some function of the difference between measured and desired average speeds. This simple idea works, but there is no proof that the resulting trajectories explore the whole phase space available.

If this method is applied to a whole system with no net momentum initially, then the system will not gain or lose momentum. If, however, there is net momentum initially, or the thermostat is applied to two regions (hot and cold as in NEMD), then momentum is not conserved.

(b) Velocity exchange

This method is a simple yet effective way of imposing a temperature difference across a simulation. The speed of the slowest atom in the hot area is replaced with the speed of the fastest atom in the cold area and vice versa. This leads to an amount of energy ΔE added to the hot reservoir, and an equal amount subtracted from the cold reservoir. An important consideration here is that the method conserves momentum. But there is a computational speed penalty due to the extra calculations needed to sort the speeds. Furthermore the steady state temperature gradient is not specified directly by the user.

(c) The Langevin thermostat

The Langevin thermostat [57] is a stochastic thermostat which adds a random force to the particle motion along with an appropriate damping term such that:

$$m\ddot{\underline{x}}_i = -\frac{\partial V}{\partial \underline{x}_i} - \gamma m \dot{\underline{x}}_i + R \underline{f}_L, \quad (2.27)$$

where γ is a damping coefficient, $0 \leq R \leq 1$ is a uniformly distributed random variable. A new random number for each dimension is picked for each particle at every time-step.

$\underline{f}_L = \sqrt{\frac{6\gamma m_i T_c}{\Delta t}}$ is the magnitude of the stochastic force for a target temperature T_c and a time step Δt for each component ($L=x, y$ or z). The Langevin thermostat is a local thermostat as the target temperature is specified for each atom (i.e. a non-uniform temperature distribution can be specified). This is advantageous as it allows for a non-uniform temperature distribution to be specified at the boundaries. One drawback is that there is no feedback between the actual temperature and the target temperature for the Langevin thermostat. This is reasonable for equilibrium thermostating, for which it was designed, but far from equilibrium there is no guarantee that the target temperature will be achieved or maintained.

A local thermostat is useful as it allows for the creation of a diffusely thermostatted boundary which avoids the problems with phonon reflection associated with a sharp boundary. Figure 2.4 shows an example where the γ factor is changed linearly over a boundary such that at the ends the atoms have heavy damping whereas at the interface between thermostat and the atomistic region, the environment that the atoms are in changes very little. This is known as ‘stadium damping’, [19] and has been shown to provide absorption of pulses with minimal scattering. This is considered further in chapter 3.




Fig. 2.4 Variation of the γ value (red line) through a cross-section of an MD simulation with stadium boundaries is shown. The shaded outer regions of the plots in the top row in figure 2.5, shows the location of the boundaries where gamma is active ($\gamma > 0$). The dashed line in the upper right of figure 2.5 shows the location of this cross-section in the simulation region. The γ value describes the strength of the thermostat, where $\gamma = 0$ is equivalent to no thermostating at all and $\gamma = 1$ is the maximum thermostat strength.

Curtin et al [19], test the stadium boundaries by introducing a pulse of energy into the system. They do this by displacing the atoms at the centre of a simulation from their equilibrium positions. The displacements applied are given by a radially varying Gaussian function. The results of a small simulation with the stadium boundaries is compared to a larger system whose boundaries are far enough away as to not affect the central region of the simulation for the short time in which the pulse is observed. The authors find that the stadium boundaries are effective in absorbing the energy from the pulse such that the centre region behaves in a similar way to the larger system. Figure 2.5 shows their results.




Fig. 2.5 Plots show a snapshot of an MD simulation of a broad pulse at 100K with time increasing from left to right. The top row is a 150Å x 150Å region with stadium boundaries (lighter shaded region) and the bottom row is the centre of 400Å x 400Å simulation. The colours represent the displacement of atoms from their equilibrium positions with red corresponding to displacements greater than 0.25Å, [19].

(d) Nosé-Hoover thermostat

This thermostat proposed by Shūichi Nosé and Hoover [58]-[60] is a deterministic thermostat which maintains the average temperature of an atomic ensemble at a target value. This is widely used for constant temperature dynamical simulations due to its symplectic, volume conserving, time-reversible Hamiltonian structure [20]. In this case the motion of a thermostatted particle is described by

$$\begin{aligned}
 m_i \ddot{\mathbf{x}}_i &= -\frac{\partial V}{\partial \mathbf{x}_i} - \xi m_i \dot{\mathbf{x}}_i \\
 Q \dot{\xi} &= \frac{1}{M_T k_B T_c} \sum_{p=1}^{M_T} m_p \dot{\mathbf{x}}_p^2 - 1
 \end{aligned} \tag{2.28}$$

where Q is a thermal mass, ξ is a thermostating variable and the summation is over all the thermostatted particles, $p=1, \dots, M_T$. This is a global thermostat in that it enforces an ensemble of particles to maintain an average kinetic energy over time. It preserves the average temperature but it does not have any control over the distribution of the temperature within the thermostatted region. Any temperature distribution which satisfies this average is possible. This is acceptable for isothermal simulations, in which there is no driving force for the distribution to be non-uniform.

As the thermostat works on the average value of the temperature over a certain range of atoms, it is slow to react to changes. For example, an incoming pulse may drastically change the temperature in one area of the thermostat only, thus the average is not too different from the target, and thus the pulse can propagate through with minimal damping and reflect from any boundaries. This will lead to an unnatural confinement of heat in an atomistic region, and therefore must be avoided.

(e) Berendsen thermostat

The Berendsen thermostat [61] is a deterministic one similar to the Nosé-Hoover one above. In this case a damping/forcing term is added, whose sign and magnitude depends on the difference between a desired temperature, T_0 , and measured temperature T :

$$m_i \ddot{x}_i = -\frac{\partial V}{\partial x_i} - \gamma m_i \dot{x}_i \left(\frac{T_0}{T} - 1 \right). \quad (2.29)$$

This is equivalent to rescaling the velocities:

$$v_{NEW} = \lambda v_{OLD}, \quad (2.30)$$

where λ is given by:

$$\lambda = 1 + \gamma \Delta t \left(\frac{T_0}{T} - 1 \right), \quad (2.31)$$

And Δt is the time step.

2.7 Boundary Conditions

The boundary conditions for MD simulations in the condensed phase are a compromise between correct representation of the far field and minimization of the system size due to computational constraints. It is the aim of this section to further investigate the non-uniform thermal boundary conditions for MD, i.e. non-equilibrium MD (NEMD).

In order to better represent a bulk material, periodic boundary conditions can be used. Fixed or free edges tend to introduce artificial boundary effects, such as phonon reflections

and scattering etc. In the periodic case, the atoms at the edge interact with a copy of the atoms from the far side. In this way, all atoms behave as though they are a bulk atom, not a surface atom. The 3-dimensional rods introduced later chapter 3 have periodic sides but fixed ends.

An alternative is to use a space which folds around on itself. For example, in two-dimensional simulations one could use the surface of a sphere as the simulation domain. Here the atoms interact with each other over the great circles of the sphere. However this is only physically reasonable for large systems, and adds greater complexity to the code. For three-dimensional systems, a hyper sphere would be needed.

Periodic boundaries cannot be used for 1D NEMD simulations unless there are two copies of the simulation as illustrated in figure 2.6. The periodic boundaries must have the same temperature to prevent any unnatural heat flow. Figure 2.6 shows two ways in which periodic boundaries can be applied to a simulation of a 1D chain. This boundary can exist inside the cold (or hot) thermostat, or at any point along the chain.




Fig. 2.6 Two schemes for using periodic boundaries in NEMD simulations of a 1D chain. (a) Periodic boundary exists within the cold thermostat. (b) Periodic boundary exists outside the thermostats.

2.8 Lattice vibrations

Consider the simplest lattice, a 1D chain of atoms as shown in figure 2.7.




Fig. 2.7 Diagram of the classical model of a 1D chain of atoms. The atoms are modelled as point particles connected by identical springs.

Each atom is the same type with a mass M and inter-atomic spacing a . The length of the chain of N atoms is, $L = Na$. For small vibrations around the equilibrium spacing a the LJ potential is harmonic. The atomic vibrations are modelled by a wave equation,

$$u_n = A \exp[i(kx_n^0 - \omega t)], \quad (2.32)$$

where u_n is the displacement of atom n from its equilibrium position. A is the amplitude, x_n^0 is the undisplaced position of atom n and t is time. The vibration of atoms in a crystal lattice can be communicated throughout the crystal via the propagation of waves. The relationship between the wave frequency ω and the wave-number k is called the *dispersion relation*. In a one-dimensional chain the dispersion relation is (for derivation see [62]):

$$\omega = \sqrt{\frac{4K \sin^2(\frac{1}{2}ka)}{M}}, \quad (2.33)$$

where K is spring constant, a is lattice spacing, M is atom mass. There are N modes and hence the number of modes depends on the length of the chain. This 1D chain can support a maximum frequency of, $2\sqrt{K/M}$. This is known as the cut-off frequency. Modes with a wavelength shorter than the atomic spacing (i.e. a higher frequency) cannot exist. This system size dependence is not only important in determining the thermal properties of nanostructures but also in determining the accuracy of MD simulations (naturally of constrained size).

In the long wavelength limit $ka \ll 1$, the approximation $\sin(x) = x$ holds. This can be seen clearly in figure 2.8, therefore equation (2.33) becomes:

$$M\omega^2 = Kk^2a^2 \quad (2.34)$$

This gives a group velocity $d\omega/dk$ and phase velocity ω/k of: $a\sqrt{K/M}$.




Fig. 2.8 Plot of equation (2.33) which shows the normal mode frequencies of a 1d chain as a function of wave-number. The three points A, B and C all have the same frequency and translate to the same atomic positions. The gradient $d\omega/dk$ is the group velocity, where positive values are rightward moving waves and negative values correspond to waves moving left.

A one-dimensional crystal only has one longitudinal acoustic branch (fig 2.8) of lattice vibrations as the atoms are constrained to move along one axis only. In a real crystal, the atoms exist in 3 dimensions and can support both transverse and longitudinal waves. Furthermore, crystals with more than one atom in its primitive unit cell can support a set of optical branches (high frequencies) [62]. The total number of modes is always equal to 3 times the number of atoms in a 3-dimensional crystal.




Fig. 2.9 The picture shows a Sound wave travelling through crystal lattice (amplitude exaggerated), with the wavelength indicated [63].

(a) Phonons

The true quantum nature of atomic vibrations is not captured by classical models such as MD [64][65]. In a real crystal, the atoms are always oscillating about their equilibrium positions. Even at absolute zero, the Heisenberg uncertainty principle [66] requires a zero point motion due to the fact that it is impossible to know the exact position and momentum of a particle simultaneously such that:

$$\Delta p \cdot \Delta x \geq \frac{h}{4\pi} , \quad (2.35)$$

where h is the Planck constant (6.63×10^{-34} Js), Δp is the error in momentum and Δx is error in position.

A phonon is a quantized lattice vibration. Phonons, like all quantum objects, can be considered as a particle or a wave. Phonons behave like a simple harmonic oscillator in that they are restricted to discrete energy values and have a minimum zero point energy.

$$E_n = \left(n + \frac{1}{2}\right) \hbar \omega \quad (2.36)$$

where ω is the frequency of a lattice vibration mode and $\hbar$ is Planck's constant divided by 2π .

Like photons, phonons have an integer value of intrinsic angular momentum (spin) and are thus classified as bosons. Consequently, they do not obey the Pauli Exclusion Principle, nor do their numbers need to be conserved.

(b) Thermal conduction by phonons

Thermal energy in a solid is due to atomic vibrations. In classical mechanics heat conduction is via the propagation and interaction of normal modes of atomic vibration. In most real crystals there are two mode types: the *acoustic* branch (lower frequencies) and the *optical* branch (higher frequencies). In practice heat transport is predominantly due to the acoustic modes and the optical modes are often ignored. Each branch consists of longitudinal and transverse mode branches. The vibration is in the direction of wave propagation for longitudinal waves and perpendicular to it for transverse waves. For a d -dimensional system ($d=1,2$ or 3) there is one longitudinal branch and $(d-1)$ transverse branches. Hence in one-dimension there are no transverse waves but in three-dimensions there are two transverse branches.

In this work, insulating materials will only be considered. The conductivity of insulators is dominated by the flow of phonons. The derivation of the conductivity comes from the kinetic theory of gases [66] and is given by:

$$k = \frac{1}{3} \bar{v} l C \quad (2.37)$$

where $\bar{v}$ is the velocity of phonons (speed of sound), l is the mean free path and C is the specific heat per unit volume. Since the speed of sound remains fairly constant and the phonon number does not need to be conserved, this simple equation fits quite well.

The calculation of the heat capacity is a very involved method and is presented in [66]. The general form is:

$$C = \frac{1}{N} \int E(\omega) \frac{\partial f(\omega, T)}{\partial T} g(\omega) d\omega \quad (2.38)$$

where $E(\omega)$ is the energy of mode ω , $f(\omega, T)$ is the Bose-Einstein distribution which represents the occupation probability of mode ω and $g(\omega)$ is the density of states which is

the number of states with a frequency in the range $\omega \rightarrow \omega + d\omega$ and is derived from the dispersion relation.

The calculation of the density of states $g(\omega)$ from the potential $V(r)$ can be very difficult in 3D. Debye simplified this considerably by neglecting dispersion of the acoustic waves to obtain a heat capacity of [66]:

$$C = 9Nk_B \left(\frac{T}{\theta_D} \right)^3 \int_0^{\theta_D/T} \frac{x^4 e^x dx}{(e^x - 1)^2} \quad (2.39)$$

where $x = \hbar\omega/k_B T$ and the Debye temperature $\theta_D = \hbar\omega_D/k_B$. At high temperatures $T \gg \theta_D$ the heat capacity tends towards the classical value of $3Nk_B$, and at low temperatures $T \ll \theta_D$ the heat capacity varies with T^3 given by:

$$C = \frac{12Nk_B \pi^4}{5} \left(\frac{T}{\theta_D} \right)^3. \quad (2.40)$$

Equation 2.39 is plotted in figure 2.10 along with the classical Dulong-Petit law (dashed line).




Fig. 2.10 Debye prediction for the atomic heat capacity as a function of temperature. The Debye model is accurate at low temperatures. The classical limit is shown as a dashed line. It is generally accepted that classical MD is valid above the Debye temperature $(T / \theta_D) > 1$.

2.9 Validity of MD

At the Debye temperature $T = \theta_D$ the quantum heat capacity is $C_{Debye} = 0.95C_{Classical}$, as seen from fig 2.10. For this reason, it is commonly accepted that classical Newtonian dynamics are acceptable for modelling heat conduction in insulators above the Debye temperature. In this regime the Dulong-Petit law (dashed line in fig 2.10) for heat capacity is very reasonable for most materials. It is still valid for many materials at low temperatures (e.g. silicon) but fails for metals, where the contribution to the heat capacity from electrons becomes significant. Values for the Debye temperature for some common elements are shown in Table 2.1. In total, 30 of the elements have Debye temperatures above room temperature of which only three are not metals: carbon, silicon and germanium. The high value for carbon means that the heat capacity is 25% of the classical value at room temperature, i.e. $C_{Debye} = 0.25C_{Classical}$.

Element	Debye temperature, θ_D
Carbon	2230K
Beryllium	1440K
Silicon	645K
Chromium	630K
Iron	470K
Nickel	450K
Aluminium	428K
Germanium	374K
Silver	225K
Argon	93K

Table 2.1 Table of the Debye temperatures of some elements.

The appropriate thermal modelling methodology for insulators depends on the length scale and temperature of a system. A summary of this is presented in figure 2.11. At dimensions comparable to the phonon wavelength λ and temperatures much smaller than the Debye temperature θ_D (bottom left in figure 2.11), the quantum mechanical effects are strongly significant. At larger dimensions approaching the phonon mean free path and room temperatures, a semi-classical approach is more pragmatic such as perturbation theory (PT) or

the Boltzmann Transport Equation (BTE). At temperatures above θ_D the classical physics of MD is applicable.




Fig. 2.11 A summary of the appropriate thermal modelling techniques for insulators at different length scales and system temperatures. θ_D is the Debye temperature, a_0 is the lattice spacing and Λ is the phonon mean free path. The diffusion length $\sqrt{\alpha t}$ is given for the sake of completeness where α is the thermal diffusivity and t is the phonon relaxation time [67].

2.10 Methods for increasing computation speed

A costly part of each computation is evaluating the force given the inter-atomic distances. Therefore any time saved in this calculation will significantly increase the speed of the computer program. Depending upon the complexity of the potential function, sometimes it can be advantageous to pre-calculate some or all of the potential and assemble a ‘lookup table’ containing the value of the potential for a range of parameters. For the LJ potential, looking up the data in a table and interpolating can take just as long as calculating the potential each time. Therefore, the computer programs used here did not use lookup tables. For potentials that include trigonometric or exponential functions, tabulating the potential can reduce computation time. The exact benefit depends on computer hardware and the

programming languages' implementation of said functions. Therefore trial and error testing is required to find the optimal method.

Neighbour lists

In condensed phase simulations, the atoms exist in a fixed lattice site. The six nearest neighbours can be pre-calculated and saved in a table which contains all possible atom pairs which contribute to the net force on each atom. This removes the need to cycle through the whole simulation just to find which atoms are near each other. For solids this list remains unchanged throughout the simulation. For fluids and gases it needs to be continually updated.

2.11 Summary

In this chapter many of the methods of controlling and performing MD simulations have been described. Inter-atomic potentials were introduced in subsection 2.2. The equations of motion are solved numerically using the Verlet algorithm, which was described in detail in subsection 2.3. Various thermostating techniques for controlling the temperature of MD simulations, such as Langevin, Nosé-Hoover and Berensen were then described in subsection 2.6. A variety of methods of increasing the efficiency of computer simulation were also described in this chapter, such as dimensionless units, neighbour lists and lookup tables. The validity of MD was also reviewed in subsection 2.9. The optimisations and techniques described in this chapter are applied to the MD simulations to be performed in chapters 3 and 4.

3 Non-equilibrium Molecular Dynamics (NEMD)

In this chapter, MD simulations of a 3D rod are performed in which each end is thermostatted at a different temperature. This introduces a thermal gradient along the rod where thermal energy is constantly flowing from the hot end to the cold end. This is the steady state where the total energy of the system does not change, yet energy is flowing through it. This is known as non-equilibrium molecular dynamics (NEMD).

Here the effectiveness of some of the different thermostating methods introduced in section 2.6, Nosé-Hoover, Langevin and stadium damping are compared using a simulation of a 3D rod with periodic boundary conditions. An extra feedback control algorithm is then added to the stadium damping method in order to impose a specific temperature in the 3D rod at a point beyond the non-linear region of the thermostat. This is done to allow greater control of the temperature gradient within the simulation.

3.1 Introduction

Most molecular dynamics simulations are sampled from the micro-canonical (constant energy) or canonical (constant temperature) ensembles. However, there have been a few studies where a steady state temperature gradient has been imposed on such an atomistic simulation [6][43][56][68]-[74]. The technique is simply to use conventional thermostating techniques to enforce different temperatures on opposite ends of the sample. These NEMD simulations can then be used to determine the effective thermal conductivity of the medium, k , from Fourier's law for macroscopic heat flow:

$$q = -k\nabla T . \quad (3.1)$$

where q is the heat flux (averaged over time and space) in the unthermostatted region between the thermostats, and ∇T is the “measured” steady state temperature gradient. Importantly note that the measured temperature gradient is not the same as the temperature gradient imposed by the thermostats due to boundary effects. The situation described above is therefore not as straightforward as it may appear. This is discussed further in this section where details of the atomistic simulation are given below, along with a discussion of potential difficulties. If only determination of the thermal conductivity is of interest this can be

determined from equilibrium simulations using linear response theories such as the Green-Kubo formula in (3.4) [43].

The heat flux q in the unthermostatted region between the two thermostats must be measured in order to calculate k . This was achieved by tracking the heat energy that was added/ removed by the thermostats at each time step. The heat flux between each atomic slice is then the difference between these values minus the increase in internal energy U . Alternatively the local heat flux can be calculated [43] using (derivation in appendix A):

$$q_n = \frac{1}{2} a (\dot{x}_{n+1} + \dot{x}_n) F(x_{n+1} - x_n). \quad (3.2)$$

where, F is the force between atom n and $n+1$ and a is the lattice spacing.

3.2 NEMD simulations of a 3D rod

In all the MD simulations presented we restrict our interest to ballistic heat transport in insulators via phonon interaction (i.e. conduction by electrons is neglected). This is an inherently non-linear phenomenon as phonons do not interact in the harmonic limit. The temporal evolution of a particle of mass m_i at a position $\mathbf{x}_i$ in the main (unthermostatted) body of the system is described by Newton's second law (2.1) with forces between atoms being computed with the Lennard-Jones inter-atomic potential (2.4).

Only nearest neighbour interactions are considered. Physically, this potential is applicable to solid argon for temperatures below 60K. In accordance with this we take $\varepsilon = 120k_b$, $\sigma = 3.4\text{\AA}$ and $m = 6.68 \times 10^{-26} \text{ kg}$ [54][68][75] where k_b is Boltzmann's constant. Note that the exact form of the inter-atomic potential is not important to the general conclusions of this work. In fact, the principal philosophy behind this proposed modelling approach is to make it as simple and general as possible, so that it is not specific to particular details of the atomic model and can be implemented into a simulation code with only minor changes.

Consider a three-dimensional rod of Lennard-Jones atoms subjected to a temperature difference at each end. To establish a steady state temperature gradient in the rod, it is

necessary to inject kinetic energy into one end of the rod and to remove it from the other end. This is achieved by the use of thermostatting algorithms. The two well-known examples of very different thermostats will be considered here: the Langevin thermostat [15] [43] and the Nosé-Hoover thermostat [43] [46] [58]. These thermostats were described in detail in section 2.6. Like most thermostats, these have been designed to maintain a system at thermal equilibrium for constant temperature MD simulation.

In NEMD simulations, where temperature gradients exist, the temperature distribution in the thermostatted region can be highly non-uniform. In this case, the temperature imposed at the edge of the thermostatted region will not be the target temperature. The Nosé-Hoover thermostat therefore also offers less potential for the stipulation of a spatial variation in temperature over a boundary, although Li and E [23] have made some notable achievements in this regard by employing a number of Nosé-Hoover thermostats to control different regions of the simulation. One advantageous property, however, is that the temperature is controlled by feedback between the actual and target temperatures, so one can be confident that the desired temperature has been achieved even in non-equilibrium simulations, unlike the Langevin thermostat. For rapidly changing transient boundary conditions, local thermostats are more responsive than global ones, which only react to a change in the global average temperature.

The steady state thermostatic control of a rectangular rod of 100 atoms in length with a periodic square cross-section of 8x8 atoms is investigated. The atoms are in the minimum energy hexagonal close packed structure and oriented such that an 8x8 cross-sectional slice represents a (100) plane. From a continuum perspective this is effectively a one-dimensional heat conduction problem, as there is expected to be no net heat flow or temperature variation within a cross-sectional slice of atoms. Therefore, each group of 8x8 atoms within a slice are referred to by an incremental index j . It is intended that the ends of the rod be maintained at different, uniform temperatures in order to achieve a prescribed steady state temperature gradient within some region of the simulation, $j=0,...,M$, where the true dynamics of the system are preserved. The thermostatted regions (TR) are at the ends with the central true dynamics region (TDR) in the middle. This setup is illustrated in Figure 3.1 for two TR of M_T atomic slices. The separate regions are therefore defined by slice indices:

$$j = \begin{cases} -M_T, \dots, -1 & \text{for left TR} \\ 0, \dots, M & \text{for TDR} \\ M+1, \dots, M+M_T & \text{for right BR} \end{cases} \quad (3.3)$$




Fig. 3.1 Schematic of the three-dimensional molecular dynamics simulation. The Lennard-Jones rod 100 atoms in length has an 8x8 square cross-section. Thermostats are used to regulate the temperature at the ends. The only net temperature gradient will be along the length of the rod. Therefore the time ensemble averaged temperature in each cross-sectional slice of 64 atoms is expected to be uniform and is denoted $\langle T_j \rangle$ where j is the index of the slice, shown in detail on the horizontal axis. A total of M_T slices are thermostatted at each end (red TR region) in order to control the thermal boundary conditions of the inner true dynamics region of $M+1$ slices (yellow TDR region).

The Langevin damping coefficient, $\gamma = \frac{1}{2}\omega_D$, is taken to be half the Debye frequency [19] (above which there are no modes). Solid argon has a Debye temperature of 93K [62], which is equivalent to a Debye frequency of $\omega_D = 1.2 \times 10^{13} \text{s}^{-1}$. The time step used in the integration of the equations of motion was chosen such that approximately 50 samples were performed in each oscillation, which equates to a time step of $\tau = 2.15 \times 10^{-14} \text{s}$. Velocity Verlet time integration was used and micro-canonical simulations were performed to test energy conservation. The extremities of the rod are fixed ($\dot{x}_{-M_T} = \dot{x}_{M+M_T} = 0$) so that the atomic spacing is the zero Kelvin equilibrium spacing, although similar results are obtained for free end conditions.

The behaviour of three different thermostats introduced in section 2.6 is investigated: Nosé-Hoover, Langevin and stadium damping. Stadium damping is a variant of Langevin which has been shown to be an effective means of phonon absorption [19][76] and

importantly to produce the expected canonical ensemble [21]. In this case, the damping coefficient is a function of position, such that $\gamma = \gamma(\mathbf{x}_i)$. As shown in Figure 2.4, it is linearly ramped from a maximum value of $\gamma = \gamma_0$ at the rod ends down to zero at the edge of the thermostatted region. This forms a diffuse interface which allows phonons to move into the damping region and be slowly absorbed as they move through it. This avoids many of the problems associated with phonon reflection at a sharp interface [18]. More sophisticated methods for sharp interfaces based on memory kernels have been considered in the literature [15][22][24]-[26] but these usually have to be calibrated for a particular potential, are derived from the harmonic approximation (for which heat conduction is not observed), and generally are not so readily implemented.

The thermostats are applied to a rod of 100 atoms in length with $M_T=15$ slices of thermostatted atoms at each end and hence $M+1=70$ atomic slices in the TDR. The fixed target temperatures for the thermostats are T_L and T_R at the left (L) and right (R) ends respectively. A temperature difference is imposed such that $T_R = 0.6T_L$. The effect of temperature is investigated by considering $T_L=5,10,25$ and 50K. Simulations are allowed to reach a steady state over a time of $10^6 \tau$ and then the average temperature of the atoms in the j^{th} slice, $\langle T_j \rangle = \langle \frac{1}{2} m_j \dot{x}_j^2 \rangle$, is determined over a subsequent period of $10^6 \tau$, where $\langle . \rangle$ denotes the time ensemble average over the slice. The resulting steady state temperature distributions along the length of the rod are shown for the three thermostats in Figures 3.2a, 3.2b and 3.2c with the associated percentage error shown in figure 3.2d.




Fig. 3.2 Steady state temperature profile along the 8x8x100 atomic rod shown in Figure 3.1. The temperature difference is imposed by thermostating the two end atoms such that the left hand thermostat temperature is $T_L = 5, 10, 25$ and 50K and that at the right hand end is $T_R = 0.6T_L$. Results are shown for (a) a deterministic Nosé-Hoover thermostat, (b) a stochastic Langevin thermostat and (c) a variant of the Langevin thermostat, stadium damping.. The expected steady state temperature distribution for a constant thermal conductivity is shown as a dashed line. The deviations from the expected result arise from the discontinuity in the temperature profile at each end due to phonon scattering. (d) The maximum percentage error for each simulation shows that the error decreases as the temperature increases and is smallest for the Nosé-Hoover thermostat, although it is still greater than 5%.

3.3 Problems of Phonon reflection, mismatch and the Kapitza effect

From (3.1) it is expected that the temperature profile will vary linearly between the target temperatures at each end for a constant thermal conductivity (as illustrated by the dashed lines in Figure 3.2). As found in previous works [6] [43] [56] [68] [70] - [74], the simulation results do not conform to this expectation. There is a drop in the temperature at the edge of the thermostatted regions, such that the temperature gradient observed in the simulation is not the temperature gradient expected from the imposed temperature difference. The maximum deviation from the correct temperature profile (to be determined and shown in figure 3.7) is plotted in figure 3.2d. The percentage error is quite significant and clearly observable in all the temperature profiles, ranging from 5-20%. The largest errors are seen at low temperatures (where phonon interactions, due to sampling non-linearities in the potential, are reduced) and for stochastic thermostats (Langevin and stadium). This effect has been widely observed and is generally attributed to phonon mismatch at the interface between the thermostatted and unthermostatted regions [71]. The effect can be even more pronounced for other potentials, e.g. silicon [71]. For real physical interfaces this is known as the Kapitza effect, where it is observed that the thermal conductivity, like most physical properties, deviates from the bulk value near an interface. Even if the interface is artificial, as is the case here, it is difficult to avoid. To understand the origin of this boundary effect, and why its magnitude depends on the thermostating method, we refer to the Green-Kubo formula [71]. This linear response theory allows the thermal conductivity to be determined from equilibrium simulations, and is often used for this purpose rather than NEMD simulations. It states that the conductivity is proportional to the long-time average of the heat flux autocorrelation. Assuming local equilibrium, the conductivity between atoms i and $i+1$ is proportional to

$$k_i \propto \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \langle q_i(s) q_i(0) \rangle ds \quad (3.4)$$

where $q_i(t)$ is the instantaneous heat flux between atoms i and $i+1$. The net energy flux between particles is therefore due to long-term correlations between their motions. Any thermostat will always necessarily alter a particle's motion and corrupt this correlation and hence reduce the thermal conductivity at the edge of the thermostatted region. This leads to a Kapitza resistance across the boundaries due to phonon mismatch [77]. The Nosé-Hoover

thermostat in Figure 3.2a is seen to be better than the Langevin and stadium thermostat in Figures 3.2b and 3.2c as deterministic methods naturally exhibit longer correlation times than (uncorrelated) stochastic methods. In the steady state, the time-averaged heat flux at each point must be constant for conservation of energy. In this case (3.1) shows that the effective conductivity at a point must be proportional to the inverse of the temperature gradient such that $k \propto 1/\nabla T$. There is a larger-than-expected temperature gradient at each end for the Langevin and stadium thermostats in Figures 3.2b and 3.2c. This implies that the conductivity is small at the thermostat interfaces and that these stochastic thermostats strongly affect any temporal correlations between particle motions at that point. The case for the Nosé-Hoover thermostat in Figure 3.2a is improved, although the boundary conductivity is still reduced.

The Kapitza effect and adverse phonon scattering are far worse in a 1D system. Figure 3.3 shows the steady state temperature profile along a one-dimensional chain subjected to thermostating at each end. Here the stochastic Langevin thermostat struggles to even apply a temperature gradient, whereas the Nosé-Hoover thermostat does slightly better. An elementary solution to these adverse, non-linear boundary effects is proposed in section 3.5.




Fig. 3.3 Steady state temperature profile along a one-dimensional 100 atom Lennard-Jones chain with fixed ends. The temperature difference is imposed by thermostating the two end atoms to 50K and 40K respectively using either a stochastic Langevin thermostat or a deterministic Nosé-Hoover thermostat. The expected steady state temperature distribution for a constant thermal conductivity is shown as a dashed line. The large deviation from the expected result arise from the discontinuity in the temperature profile at each end due to phonon scattering.

(a) Phonon reflection

Phonons can reflect at the interface between the bulk and the thermostat regions [62][78]. Even in a constant temperature simulation, perturbations to the system can cause artificial heating in the central bulk region due to the reflection of phonons. Figure 3.4 shows the energy in the central region of a 3D rod during the introduction of a pulse at the centre. It can be clearly seen that using the Langevin thermostat, the central region remains hotter than it should be.




Fig. 3.4 Plot of the energy in the centre of a rod vs. time after a pulse is introduced. Several thermostating methods are compared along with the no damping case where the edges of the rod move freely. The results are similar to those presented in [19].

The Stadium boundaries used below are known to reduce phonon reflections [19]. In order to test this boundary a pulse was set off by increasing the velocity of the two centre atoms towards each other. This simulation was performed on a large (2000) atom system and a small system with stadium damping. Sufficient time passed for any pulses to reflect off the edges in the stadium system. In Figure 3.5 the total energy in the central region for both cases is compared. As can be seen, the energy decays at the same rate for both systems.




Fig. 3.5 Plot of the energy in the centre of a rod vs. time after a sharp pulse is introduced for both a large system (Black curve) and a small system with stadium boundaries (Red curve). The simulation was performed at (a) 1K and (b) 10K.

3.4 Divergence of thermal conductivity

One can measure the thermal conductivity using the NEMD experiment (Figure 3.2a) with the Nosé-Hoover thermostats. The temperature gradient is measured in the central linear part of the graph, and the heat flux is calculated in the simulation using equation (3.2). Then Fourier's law (3.1) can be used to calculate the thermal conductivity, k . Many authors have found that the measured conductivity depends on how long the atomistic chain is [43][72][79][80]. In investigations of the origins some have looked at on site potentials [81] and disorder [82]. Divergence of the conductivity is also found in 2D [83]. This effect exists in one and two dimensions, but vanishes in three dimensions. It is predicted that for momentum conserving potentials, (i.e. no on-site potential, and an interatomic potential that conserves energy such as LJ) a system of characteristic length N , will have thermal conductivity:

$$\begin{aligned}
 k &\propto N^{2/5} && \text{in 1D case, longitudinal heat flow only,} \\
 k &\propto \ln(N) && \text{in 2D case, longitudinal and transverse heat flow,} \\
 k &\propto \text{Const} && \text{in 3D case, longitudinal and transverse heat flow.}
 \end{aligned}$$

For coupling with continuum simulations, it is expected that an atomic system will obey Fourier's law (3.1) such that the thermal conductivity, k , will be independent of N in the macroscopic limit, i.e. although there are expected to be system size effects, the conductivity will converge to its bulk value as the system size increases. Therefore, although working with a one-dimensional chain is desirable from a model development point of view, it introduces problems which will not be an issue in three-dimensions. Therefore, only three-dimensional simulations will be considered for the remainder of this thesis.

3.5 NEMD simulations with extra feedback control

In these simulations, instead of scaling the thermostat to a desired value, the program scales the thermostat temperature in order to maintain an atom outside the thermostat at a desired temperature. This simple algorithm provides extra feedback between the actual temperature of the atoms and the thermostat's state.

The steady state thermostatic control of a rectangular rod is again investigated, but to avoid the problems due to boundary effects and global thermostating methods, the thermostatted regions (TR) at the ends are separated from the central true dynamics region (TDR) in the middle by two small buffer regions (BR). This setup is illustrated in Figure 3.6 for two BR of M_B atomic slices and two TR of M_T slices. The separate regions are therefore defined by slice indices

$$j = \begin{cases} -(M_B + M_T), \dots, -(M_B + 1) & \text{for left TR} \\ -M_B, \dots, -1 & \text{for left BR} \\ 0, \dots, M & \text{for TDR} \\ M + 1, \dots, M + M_B & \text{for right BR} \\ M + M_B + 1, \dots, M + M_B + M_T & \text{for right TR} \end{cases} \quad (3.5)$$




Fig. 3.6 Schematic of the three-dimensional molecular dynamics simulation for the thermal boundary conditions with the feedback control algorithm. In this case a buffer region of M_B slices (green BR region) is introduced between the TRs and the TDR to avoid any corrupting boundary effects at the edge of the TRs.

It has been previously noted that there is not a straightforward relationship between the temperature that is imposed on the TRs and the temperature gradient which develops between them in the TDR. To ensure that there is precise control of the temperature within the TDR, the average temperature of the atomic slices at the edges of this region, $\langle T_0 \rangle$ and $\langle T_M \rangle$ are determined, where $\langle T_j \rangle$ is the time ensemble average of the temperature over all the atoms in slice j , as before. It is desired that these two atomistic temperatures attain the constant values T_0 and T_M respectively in the steady state. The target temperatures in the left and right TRs, T_L and T_R , must evolve to ensure that the actual temperatures at the periphery of the TDR ($\langle T_0 \rangle$ and $\langle T_M \rangle$) are maintained at the prescribed values (T_0 and T_M). A simple feedback loop is used to achieve this such that

$$\dot{T}_L = \frac{[T_0 - \langle T_0 \rangle]}{Q_T} \quad \dot{T}_R = \frac{[T_M - \langle T_M \rangle]}{Q_T} \quad (3.6)$$

where the constant Q_T determines the responsiveness of the thermostat. One benefit of this algorithm is that it provides feedback for the stochastic methods which was previously missing. Hence one can be certain that Langevin and stadium damping achieve the desired temperature.

Results are obtained for the same conditions as in Figure 3.2, with $T_0=5, 10, 25$ and 50K and $T_M = 0.6T_0$ with $M_T=15$, $M_B=10$ and a TDR of $M+1=70$. These temperature profiles are shown in Figure 3.7 for the three different thermostats under investigation, with $Q_T=100\tau$ and a thermal mass of $Q=10\varepsilon\tau$ for the Nosé-Hoover thermostat. It is found that all the thermostats correctly impose the prescribed temperatures at the boundaries of the TDR and that they all achieve a steady state in a similar time.




Fig. 3.7 Precise imposition of a steady state temperature gradient along a three-dimensional (8x8x100 atom rod) NEMD simulation by feedback control of the thermal boundary conditions using (3.6) for (a) Nosé-Hoover, (b) Langevin, (c) stadium damping thermostats. The time-averaged temperature profile along the rod is shown. The temperatures in the left and right thermostatted regions ($M_T=15$ atomic slices = 960 atoms each) are regulated at T_L and T_R by (2.9) such that the prescribed target temperatures in atomic slices $j=0$ and $j=69$ respectively are

maintained at $T_0=5, 10, 25$ and 50 K and $T_M = 0.6T_0$ for all four cases. A buffer region (BR) of $M_B=10$ atomic slices is introduced between the thermostatted regions (TR) and the true dynamics region (TDR) to avoid Kapitza effects at the TR/BR interfaces. Note that the temperature variation within the TR is non-uniform and unpredictable. The stochastic Langevin and stadium damping thermostat require a much larger temperature difference between the TR at either end to impose the same temperature gradient than the deterministic Nosé-Hoover thermostat. The Kapitza effect at the boundaries is smallest for the stadium damping case. The profiles in the TDR are not exactly linear due to changes in the thermal conductivity with temperature.

Inspection of the temperature variation over time at every slice within the TDR shows that it reproduces the Boltzmann distribution (that one would expect at the equivalent equilibrium temperature). The Boltzmann distribution curves at equal points along a rod for one case is shown below in figure 3.8.




Fig. 3.8 Speed distributions for atoms at each edge and the centre for the experiment with Nosé-Hoover boundaries.

The temperature profiles of the feedback simulations in figure 3.7 are not exactly linear. This is because the thermal conductivity changes with temperature. The resultant curve in the profile is therefore most apparent for the case with the biggest temperature difference. The three principal differences between the three algorithms are in the TRs and the BRs. Firstly, the temperature distribution in the TRs is highly non-uniform for the Nosé-Hoover and Langevin thermostats. This is because there is an abrupt change in the dynamics of the particles between the TRs and BRs and, in the case of the global Nosé-Hoover thermostat, this is controlling the average temperature not the temperature distribution. The temperature

distribution in the stadium damping case is more linear, which is in accordance with the linear variation in the damping parameter, $\gamma(x)$. Secondly, a major difference between the stochastic thermostats and the deterministic Nosé-Hoover thermostat in Figure 3.7 is that a much greater difference between the thermostat target temperatures, $T_L - T_R$, is required for the stochastic thermostats to achieve the prescribed temperature difference, $T_0 - T_M$. This is expected from Figure 3.2, where the stochastic Langevin thermostat is observed to corrupt the particle dynamics considerably leading to a larger drop in the conductivity at the interface with the thermostat region. This large temperature difference could be problematic if the lower target temperature in the thermostat region dropped below zero or the higher target temperature exceeded the melting point. However, this is only expected to be truly problematic in the presence of very high temperature gradients. The temperature gradient considered in this study is high for investigative purposes, but it is not unrealistic and can be easily attained when two surfaces at slightly different temperatures first come into contact for instance. The third difference between the thermostats is in the BRs. For the Nosé-Hoover thermostat the desired temperature gradient is only established within 5-6 atoms of the TR. For the Langevin thermostat there is a large drop in temperature over the first 2-3 atoms in the BR but then the correct temperature gradient is achieved. The best case is for the stadium damping thermostat, for which the boundary effect is small, and there is a smooth transition from the TR to the desired temperature gradient over 2-3 atoms. Hence it seems possible that the BR could be reduced in size for the stochastic algorithms. Note that energy is being injected by the thermostat into the left hand TR and removed from the right hand TR. Both ends are controlled by an identical algorithm which can handle either situation without modification. This is important when considering transient problems where the direction of heat flow across a boundary can reverse during the course of the simulation. The stadium damping algorithm produces the smoothest temperature profile across the TR and BR, is easy to implement, addresses each atom locally and provides a diffuse interface for phonon absorption. This method is therefore considered to be the best candidate for further implementation and is therefore the only thermostat considered for the coupled simulations in section 4.

4 A coupled atomistic/continuum model for heat flow

It is the aim of this work to develop algorithms which firstly allow the temperature of an MD simulation to be precisely controlled away from thermal equilibrium, and secondly allow MD simulations to be concurrently coupled with a continuum representation of the thermal far-field.

Therefore, in this chapter, a continuum representation of the NEMD simulation in chapter 3 is developed and concurrently coupled to an atomistic NEMD simulation to provide full control of the remote boundary conditions. A continuum model of the one-dimensional heat conduction problem is derived in section 4.1. In this case the finite difference method is used although a finite element method could be implemented just as readily. The model is initially developed in section 4.2 within the context of the steady state analysis of the previous section. This is then extended to the fully transient case in section 4.3, where the boundary conditions to the TDR are a function of time.

The coupled model is then applied to simulating the effects of time varying thermal boundaries on solid and composite 3D rods.

4.1 The continuum heat conduction model

A continuum finite difference model is employed in regions which overlap the TRs and the BRs, as shown in Figure 4.1a. The one-dimensional finite difference grid matches the initial regular positions of the atomistic slices in these regions. This is not a requirement of the model, for which the grid-spacing can be irregular and extend to any remote position. The nodal temperatures on the finite difference grid are denoted as $\tilde{T}_j$ and are at a fixed position $\tilde{x}_j$, where j denotes the slice number. These temperatures can evolve by the usual finite difference algorithm

$$c\tilde{T}_j = \tilde{q}_j - \tilde{q}_{j-1} \quad (4.1)$$

where,

$$\tilde{q}_j = k \left(\frac{\tilde{T}_j + \tilde{T}_{j+1}}{2} \right) (\tilde{T}_{j+1} - \tilde{T}_j) / (\tilde{x}_{j+1} - \tilde{x}_j) \quad (4.2)$$

is the continuum heat flux between nodal points j and $j+1$.




Fig. 4.1 A schematic diagram of the coupling between the NEMD simulation and the finite difference continuum representation for (a) steady state and (b) transient thermal boundary conditions. This is only shown for the left hand end of the rod, as a similar situation exists at the other end. The steady state boundary condition drives the atomistic and continuum temperatures at the CR/TDR interface to the same value using (3.6) and (4.6). The transient boundary condition is slightly more expensive to calculate as it conserves the heat fluxes across the TR/BR interface. This interface is chosen (rather than the BR/TDR interface) to minimize delays in the responsiveness of the thermostat. The transient boundary condition is also applicable to steady state conditions.

The thermal conductivity $k(T)$ is assumed to be a function of temperature T . The constant c is the heat capacity of an atom and is defined to be the amount of energy required to raise the temperature of an atom by 1K. In the simple case of a classical three-dimensional crystal with pairwise interactions this is readily determined to be $c = \frac{d(3k_B T)}{dT} = 3k_B$. The Debye model predicts the heat capacity for the more general case [62].

To implement a compatible continuum model it is necessary to have knowledge of the relevant macroscopic material parameters. In this case the only unknown parameter is the thermal conductivity. Close inspection of the steady state temperature profiles within the TDR in Figure 3.7 show that it is not precisely linear due to the variation in the thermal conductivity with temperature. To first order, it is proposed that the conductivity is a linear function of temperature such that $k(T) = k_0 + \nabla k T$. In the steady state the heat flux of (3.1) has a constant value:

$$q_{ss} = k(T) \frac{dT}{dx}. \quad (4.3)$$

This is readily integrated to give

$$q_{ss} x = k_0 (T - T_0) + \frac{1}{2} \nabla k (T^2 - T_0^2) \quad (4.4)$$

where the position is $x=0$ at the $j=0$ slice at which $T = T_0$. The observed temperature distribution closely fits the quadratic of (4.4) but knowledge of the steady state heat flux q_{ss} is required to determine an absolute thermal conductivity value.

The instantaneous (spatially averaged) atomistic heat flux between each slice, q_j , at time t can be determined from

$$q_j(t) = q_{j-1}(t) + \sum_{\text{slice } j} \mathbf{f}_T \cdot \dot{\mathbf{x}} + \frac{1}{2\tau} \sum_{\text{slice } j} m \left[|\dot{\mathbf{x}}(t)|^2 - |\dot{\mathbf{x}}(t-\tau)|^2 \right] \quad (4.5)$$

where the first term on the right hand side, q_{j-1} , is the heat flow out of the previous slice, the second term is the rate at which the thermostat adds energy to slice j (zero in BR and TDR) and the third term is the rate of change in the kinetic energy of the slice over the time step τ . The force on each atom due to the thermostat is $\mathbf{f}_T = R\mathbf{f} - \gamma m \dot{\mathbf{x}}$ from (2.27) for the stochastic

thermostats and $\mathbf{f}_T = -\xi m \dot{\mathbf{x}}$ from (2.28) for the Nosé-Hoover thermostat. The heat flux into the first slice (in the TR) at $j = -(M_B + M_T)$ is zero as these atoms are fixed in space. Hence the time averaged heat flux $\langle q_j \rangle$ can be calculated for every slice using (4.5). This was done for the steady state simulations of Figure 3.7. The heat flux in the TDR was found to be constant with $q_{ss} = \langle q_j \rangle$ for all $j=0, \dots, M$. The temperature distribution was then fitted to (4.4) and the thermal conductivity determined to be $k_0 = 1.016 \text{ W/(m.K)}$ and $\nabla k = -0.02 \text{ W/(m.K)}$.

Note that the long boundaries of the rod are periodic so the non-uniform thermal expansion of the rod cannot be accommodated. The compressive strain in the rod therefore increases with the temperature and varies along its length. It has recently been calculated that the thermal conductivity is a function of strain [6]. Ideally a continuum model will not require material parameters to be pre-determined, especially if they are a complex function of state, e.g. temperature, strain and crystallographic orientation. It is often difficult or too time consuming to completely characterise a parameter in terms of the many state variables. Therefore it is preferable to determine these parameters on-the-fly [23] or at least refine them during the course of a simulation. This is discussed within the context of the coupling methodology proposed in the following sub-section.

4.2 Thermal boundary conditions for steady state coupled atomistic/continuum simulation

Two finite difference continuum regions (CR) containing $N+1$ nodes and defined by (4.1) are now coupled to each of the two ends of the atomistic NEMD simulation, as shown in Figure 4.1a for the left hand region only. The nodal positions go from $\tilde{x}_{-N}$ to $\tilde{x}_0$ for the left hand CR and from $\tilde{x}_M$ to $\tilde{x}_{M+N}$ for the right hand CR. The nodes at $\tilde{x}_0$ and $\tilde{x}_M$ coincide with the edges of the TDR in the NEMD simulation at slices $j=0$ and $j=M$. The fixed boundary temperatures are now prescribed at the outermost boundaries of the finite difference grid and are denoted $\tilde{T}_{-N}$ and $\tilde{T}_{M+N}$. Unlike the analysis of section 3.5, the target temperatures at the boundaries of the TDR, T_0 and T_M , are now no longer fixed at a particular value. The NEMD simulation is now only of interest in the TDR. The TRs and BRs are only used to control the thermal boundary conditions to the TDR. These boundary conditions are determined by

matching conditions at the interface between the CR and TDR at $j=0$ and $j=M$. For steady state analysis it is sufficient to simply specify that the target temperatures for the edges of the TDR should be the same as those at the matching node in the finite difference model such that:

$$\begin{aligned} T_0 &= \tilde{T}_0 \\ T_M &= \tilde{T}_M \end{aligned} \quad (4.6)$$

The thermostat temperatures are determined by the feedback control equation (3.6) as before, as shown in Figure 4.1a.

Results are shown in Figure 4.2 for the stadium damping thermostat for $N=20$. The same temperature gradient prescribed in Figure 3.7 is maintained within the TDR, although this time the temperatures are defined at the remote boundaries of the finite difference simulation not at the edges of the TDR. The fixed end temperatures are taken to be $\tilde{T}_L = \tilde{T}_{-N} = 40\text{K}$ and $\tilde{T}_R = \tilde{T}_{M+N} = 20\text{K}$. The temperature profiles in the CRs at the left and right hand sides and the temperature in the TDR are shown. The unphysical temperatures in the TRs and BRs are not shown in these or subsequent simulation results. It is found that the temperature profile quickly settles down to the expected stable steady state. This is a considerable achievement as the atomistic region is highly dynamic and the temperature at the continuum/atomistic (CR/TDR) interface fluctuates rapidly. The stability of the method is not found to be very sensitive to the choice of the material parameters in the CR (k_0 , ∇k and c) which is important. A discontinuity in the gradient of the temperature curves at this interface is only observed if there is a significant thermal conductivity mismatch between the CR and TDR. As the heat flux and temperature gradient are known at all points in the simulation it is a simple matter to determine the thermal conductivity during the simulation. This value can be used to refine the measured value on-the-fly. This is especially useful if the conductivity is expected to change significantly due to a variation in the state variables during the simulation.




Fig. 4.2 Steady state results for the coupled atomistic-continuum simulation. Target temperatures of 40K and 20K are now prescribed at the edges of the continuum regions (at $j=-20$ and $j=100$). The target temperatures at the boundaries of the TDR (at $j=0$ and $j=80$) are now variable and determined from the coupling between the continuum model and the NEMD simulation.

4.3 Thermal boundary conditions for transient coupled atomistic/continuum simulation

The previous work in this thesis has only considered the steady state response of NEMD simulations. However, one of the advantages of a continuum model is that full control of the remote boundary conditions is obtained. These boundary conditions are dependent on the temperature or the heat flux at the boundary and can therefore be functions of time. In this section the methodology is extended to consider such cases. This is achieved by discarding the temperature matching condition (4.6) at the atomistic/continuum interface and explicitly enforcing the conservation of thermal energy between the models instead. A first approach would be to simply equate the continuum and atomistic heat fluxes on either side of the CR/TDR interface. For the left-hand side shown in Figure 4.1b this would require $\tilde{q}_{-1} = q_0$.

However the instantaneous atomistic heat flux is rapidly changing so this is enforced on average over time by controlling the thermostat target temperatures at each end such that,

$$\dot{T}_L = \frac{1}{Q_Q} \int_0^t (\tilde{q}_{m-1} - q_m) dt \quad \dot{T}_R = \frac{1}{Q_Q} \int_0^t (q_{m-1} - \tilde{q}_m) dt, \quad (4.7)$$

where Q_Q determines the response rate of the system to disparities in the heat flux and the integer m is the index of the slice where this condition is to be applied. The integral ensures that no heat is lost over time. The temperatures in the atomistic and continuum regions must also be connected. This is achieved by defining the continuum temperature at the CR/TDR interface to be the time-averaged value at the same point in the atomistic simulation such that,

$$\begin{aligned}\tilde{T}_0 &= \langle T_0 \rangle \\ \tilde{T}_M &= \langle T_M \rangle.\end{aligned}\tag{4.8}$$

The remaining issue when dealing with thermal transients is the responsiveness of the system. There is an inherent time delay in the system as a change in the target temperature of the TR takes a small time to effect a change in the heat flux at another point in the simulation. This is unavoidable with this methodology so the objective is to minimize the effect of this delay. This can be achieved by optimizing the position, m , at which heat flux conservation between the continuum and atomistic models is enforced. Selecting this to be the BR/TDR interface ($m=0$ on left hand side) gives reasonable results although there is still a small noticeable delay between the two descriptions. It is found that the optimal position is at the TR/BR interface as shown in Figure 4.1b. This minimizes the distance between the thermostat and the system controller (4.7). The BR has been reduced in size to $M_B = 5$ based on the steady state observations for the stadium damping thermostat in Figure 3.7. This is because the system is more responsive if the buffer zone is smaller.

4.4 Results

Results for a variety of test cases are shown in figures 4.3 – 4.8. These demonstrate the ability of the model to respond to changes in the boundary conditions at the atomistic/continuum interface over time, including multiple reversals in the heat flux from heat entering to heat leaving the boundary. Even though the thermal change is large and rapid, there is no observable delay between the response of the continuum and atomistic regions. In each case the combined atomistic/continuum results agree very well with the results from a full finite-difference simulation (shown as dashed lines). All parameters remain the same as before with the additional parameter $Q_Q = 6\varepsilon\tau$. The choice of value for this parameter is not

critical and the algorithm is not particularly sensitive to it. All the graphs are plots of the average of multiple simulations at 10000 τ intervals to reduce thermal noise.

(a) Rapid discrete boundary variations

The first case shown in Figure 4.3a has the remote boundary temperatures fixed at $\tilde{T}_L=40\text{K}$ and $\tilde{T}_R=20\text{K}$. The simulation starts with one half held at 40K while the other half is held at 20K. The centre slice is initially fixed such that no heat flows between the two halves while the system is thermalised. After sufficient time the centre slice is allowed to interact as normal, and the system evolves towards a steady state. The second case is shown in Figure 4.3b. Here the system is allowed to reach a steady state with the outer continuum temperatures set at $\tilde{T}_L=40\text{K}$ and $\tilde{T}_R=30\text{K}$. The left outer edge $\tilde{T}_L$ is then instantaneously dropped to 20K, and the system allowed to evolve. The boundary conditions at the CR/TDR interface are truly transient as the temperature and heat flux evolves continuously over time at this point. Figure 4.3c illustrates the third case. A central part of the TDR is thermalised at 40K while all surrounding atoms and the CRs are thermalised at 20K. As in Figure 4.3a, this is achieved by fixing the atoms between these different temperature regions. When the interactions are turned on again, the 40K region rapidly cools down to the temperature of the surrounding 20K region, reaching equilibrium after a time of $10^5\tau$. The fourth case, shown in Figure 4.3d, is for a coupled system that is initially thermalised at 20K. The left hand temperature $\tilde{T}_L$ is then instantaneously increased to 40K. The system evolves as expected towards the steady state. As expected the transient boundary conditions produce the same result as the steady state boundary conditions once equilibrium has been achieved. However, the steady state boundary conditions in subsection 4.2 are still of value for purely steady state problems as they do not require the additional complication of the flux calculation of (4.5). Importantly, the continuum coupling has virtually no computational overhead, with the NEMD part of the simulation accounting for practically all the processor time.




Fig. 4.3 Snapshots of the temperature profile evolution for transient boundary conditions for an entirely continuum model (dashed lines) and a coupled atomistic/continuum model, where the TDR is $0 \leq j \leq 80$. Four cases are considered: (a) the simulation starts with one half held at 40K while the other half is held at 20K; (b) the system is allowed to reach a steady state with the outer edges set at 40K and 30K. Then the left outer edge is suddenly dropped to 20K and the system allowed to evolve towards its new steady state; (c) a central region within the TDR is thermalised at 40K while all surrounding atoms and the continuum are at 20K. The hot inner region cools over time until the system has cooled down entirely to a steady state temperature of 20K after a long time; (d) Initially at 20K, the left hand boundary is subject to an instantaneous temperature increase to 40K. Again, the system evolves as expected towards the steady state. In each case the coupled simulation result agrees very well with the result from the continuum simulation.

In the problems of Figure 4.3, the system always evolves towards a steady state. A situation which better demonstrates the potential of the coupled atomistic/continuum method proposed here is shown in Figure 4.4. This is a situation where the system is large (infinite)

and will not reach a steady state (within a finite time). Energy is constantly injected into the central slice of the TDR by adding a random force to the atoms there. The temperature of these atoms increases and heat flows out towards the boundaries. Figure 4.4 shows the temperature profile evolution for two approximations to the full atomistic simulation of this problem (where we imagine that it is not possible to obtain the full solution due to computational limitations). Case (A) is an atomistic approximation of a reduced system employing a 100 slice NEMD simulation with the temperature fixed at the local boundaries (using the algorithm in section 3.5); case (B) represents a similar NEMD simulation of 100 slices coupled to a large continuum region. This is expected to provide more realistic boundary conditions as the temperature at the atomistic/continuum boundaries can evolve over time. The inset in Figure 4.4 shows the temperature profile at a time $t=10^4\tau$ when heat is just beginning to cross the atomistic system boundaries. At this point the profiles are similar. The main figure illustrates the thermal profile at a much later time, $t=10^5\tau$ (the earlier case (B) $t=10^4\tau$ result is shown again for reference). The temperature profiles are now quite different. The constraint of the local fixed temperature boundary conditions in case (A) has had a major effect. A steady state temperature gradient has developed such that the heat flow from the centre balances the rate of energy input into the system. This is an artefact of the system size. In case (B) the atomistic simulation is embedded within a large continuum region so that it cannot see the remote boundaries within the time scale of the simulation. It is important to note that case (B) does not require significantly greater computational time than case (A), i.e. the continuum region and coupling algorithm has a very small computational overhead compared to the NEMD simulation. This class of problem is representative of a situation where work is being done on an atomistic simulation but the simulation is reduced in size (for computational efficiency) such that heat is crossing the simulation boundaries (without being properly accounted) for during the simulation time, e.g. differential thermal contact such as an AFM tip on a substrate [70], nanoindentation [41], wear [28] or crack growth [76].




Fig. 4.4 Transient analysis for two identical NEMD simulations except for their boundary conditions : case (A) has the temperature fixed at 20K at its boundaries (at $j=0$ and $j=100$) whilst the other, case(B), has its boundaries coupled to a much larger continuum simulation. Energy is injected into the centre of the NEMD simulation at $j=50$ causing the temperature at the centre to increase and heat to flow out towards the boundaries. The inset picture compares cases (A) and (B) at a time ($t=10^4 \tau$) when the heat flux is starting to be affected by the boundary conditions. At this point the temperature profiles are similar. As the system continues to evolve, the temperature gradient in case (A) balances the rate of energy input into the system and achieves a steady state. This is a direct consequence of the local boundary conditions. Meanwhile, case (B) continues to evolve as heat flows across the atomistic/continuum boundary into the large CR which provides more appropriate boundary conditions.

(b) Harmonically varying boundary

A fully transient boundary varies in time constantly. In order to thoroughly test the thermal coupling method, the edge of the continuum region is subject to a cyclic variation in its temperature between 20K and 40K. Each oscillation has a time period of 50000τ . Figure 4.5a is a plot of the temperature at three equally spaced points along the MD region as a function of simulation time. Also shown is the temperature at the edge of the continuum which is being controlled (black curve). On the same axis a full finite elements solution is plotted, and is found to be in good agreement with the coupled simulation. Figure 4.5b shows

the measured temperature profile (solid line) and finite elements prediction (dashed line) of the 3D rod at five points in time during an oscillation of the boundary. These plots show that the temperature profile along the rod is in good agreement with the full finite elements solution.




Fig 4.5 Results of the coupled MD/FD simulation with fully transient boundary conditions. The temperature of the far edge of the continuum is driven sinusoidally between 20K and 40K. (a) Plot of temperature (averaged over 1000 τ) vs. simulation time at equally spaced points along the 3D rod. (b) Plot of the temperature profile of the rod at equally spaced points in time during a single oscillation of the remote boundary. These results were presented at MMM2008 in Tallahassee, Florida [84].

(c) Simulation of a composite material

There is little point in performing coupled MD/FD simulations where a simple finite elements simulation yields the same results, other than for testing purposes. For this reason, a composite 3D rod was investigated. Here the atoms in the central 30 atomic slices of the rod have their mass increased by a factor of 10. Therefore, composite rod consists of three regions, two outer M regions and a central $10M$ region. This creates a real interface where phonon reflections and a Kapitza effect are expected to occur.

Steady state NEMD simulations of this heavy ($10M$) material were performed in order to measure the thermal conductivity which is required for the finite elements simulation. This is not required for the coupled simulations, and is only used in the full FE simulations for comparison. The $10M$ region has a temperature dependant conductivity given by

$K = K_0 + \Delta k T$ with $K_0 = 0.35$ and $\Delta k = -0.0075$. This is much lower than that obtained for the original rod, which was $K_0 = 1.037$ and $\Delta k = -0.02$.

Figure 4.6 shows the steady state temperature profile (solid line) of the composite rod with the remote boundaries fixed at $\tilde{T}_L = 40\text{K}$ and $\tilde{T}_R = 20\text{K}$. The profile is an average over a time of 80000τ long after any transient effects have vanished. A large temperature drop in the coupled MD/FD simulation occurs at the interface between the M and $10M$ regions. This is not seen in the standard continuum solution (dashed pink line).

The dashed red line is the same continuum solution except at the interface between regions a cell with a much lower conductivity is used. This conductivity is also temperature dependant as can be readily seen in figure 4.6 from the different size temperature drops each side of the $10M$ region. This conductivity is given by:

$$K = \frac{1}{2}[T_j + T_{j+1}]\Delta k + K_0 \quad (4.9)$$

where j is the index of the slice, $K_0 = 0.0001$ and $\Delta k = 0.001$. This was obtained from measurements of the MD simulation.





Fig 4.6 The steady state temperature profile for a coupled atomistic/continuum (MD/FD) model of a composite rod. The temperatures are fixed at the edges of the continuum regions (CRs) to be 40K and 20K respectively. The temperature in the atomistic TRs and BRs is not shown. A region in which atoms have a large mass ($m=10$) is located in the centre of the atomistic rod. There is a large drop in the temperature at the interface between the regions of different masses due to the Kapitza effect. Results for the pure FD model (pink dashed line) do not reflect this unless a reduced interface thermal conductivity k is used at the interface between the high and low mass regions.


Transient results

Discrete jumps in temperature of the far left continuum boundary are applied to the system to investigate its transient response. The simulation starts with all continuum nodes and atoms thermalised at 20K. After sufficient time ($10^4\tau$) the edge of the continuum is raised instantly to 40K. After a time of $5 \times 10^4\tau$ the edge of the continuum is set back to 20K again. Figure 4.7 shows successive temperature profiles averaged over $10^4\tau$ during the (a) heating and (c) cooling phases. Plots (a) and (b) are the same data compared with the different FD comparison models (similarly for (c) and (d)).

Without the artificially low conductivity cell at the interface, the FD model is vastly different from the coupled simulation result. This highlights the value of the full atomistic model where atomistic effects are significant. The continuum solution could not predict the temperature drop at the interface, and only matches up when a FD cell of reduced k is used at the interface. These results were presented at IUTAM 2008 [85].



(a)



(b)




Fig. 4.7 MD/FD simulation results (solid lines) demonstrating the control of transient boundary conditions during the heating and cooling phase of a composite rod. (a) Initially at a uniform temperature of 20K, the far left continuum boundary temperature is instantly changed to 40K at $t = 0$ to heat the system. The dashed line is the full FD solution and (b) FD solution with a reduced interfacial conductivity k given by (4.9). (c) The left continuum boundary temperature is now instantly changed back to 20K. This occurs directly after the heating phase. The dashed line is (c) a pure FD solution (d) FD solution with a reduced interfacial conductivity k given by (4.9).

Discussion and Conclusions for section I

A simple molecular dynamics/continuum coupling algorithm has been developed which does not depend upon the underlying interatomic potential model employed. This should therefore be useful in a wide range of areas from solid crystals to fluids [86][87]. The models in this paper have been developed within the context of a quasi one-dimensional problem in which net heat flow only occurs in one direction. The situation is slightly more complicated for controlling the temperature in more than one direction but the method is still applicable. The thermostatting region will no longer necessarily drive the boundary of the TDR towards a uniform temperature, i.e the temperature can vary around the boundary. In the case of local thermostats such as Langevin and stadium damping this is not a problem as each atom can be driven towards a different temperature. Li and E [23] have shown that global thermostats such as Nosé-Hoover can also achieve this by dividing the thermostatted region in to sections along the periphery and having a different target temperature in each. However, this sectioning can complicate the algorithm and leads to an increase in the number of

thermostatting variables ξ . This is not the case for the local thermostats, where atoms can be individually labelled as thermostatted or unthermostatted. Hence the TR can be easily changed to accommodate an expansion/contraction of the TDR over time by altering which atoms are considered to be part of it. The steady state boundary conditions considered in sections 3.5 and 4.2 require a relationship between the target temperatures on the BR/TDR boundary and the thermostat temperatures with the TR, similar to that given in (3.6). This is obvious if a boundary has a uniform temperature (as considered here) but more complicated if the temperature varies continuously along the boundary, where a particular thermostatted atom will affect the temperature of a number of target temperature atoms and vice versa. In this case the thermostat temperatures can be related to the target temperatures through a compact local proximity weighting kernel. The more general transient boundary conditions considered in section 4.3 however only require a knowledge of the heat flux which can be readily calculated without additional complications to the algorithm.

The issue of thermal expansion has been briefly mentioned in section 4.1. Coupled thermoelastic boundary conditions need to be considered to correctly model this phenomenon. Qu et al. [19] have proposed an isothermal model for coupled continuum/atomistic elasticity problems using a stadium damping thermostat at the boundary to avoid adverse phonon reflection. This method is compatible with the approach here for combined coupled atomistic/continuum thermoelastic problems. This would be a beneficial extension, as stresses due to thermal expansion will always be an issue in complete NEMD simulations. The quasi-elastostatic continuum far field is coupled to the quasi-static atoms at the outer edge of the TR. The force on the nodes of the finite element field must balance the time-ensemble average force on their equivalent atoms in the TR. The only significant difference between the thermal equilibrium case [23] and the non-equilibrium case considered here is that now the temperature in the TR surrounding the TDR is non-uniform. Away from equilibrium the temperatures in the TR are unphysical and chosen to achieve the desired target temperatures at the TDR boundary. This leads to unphysical thermal expansion or contraction of the TR and BR. To implement the method of Qu et al. [23] it would be necessary to subtract this unphysical thermal expansion from the average atomic positions in the TR by calculating the difference between the physical temperature in the finite element field and the unphysical temperature at the same point in the TR.

In summary, a method for controlling the thermal boundary conditions of non-equilibrium molecular dynamics (NEMD) simulations has been presented in section 3. The method is simple to implement into a conventional molecular dynamics code and independent of the atomistic model employed. The body is thermostatted at the boundaries to control the temperature at the edges of the true dynamics region (TDR). A small buffer region lies between these two regions to avoid the TDR being corrupted by boundary effects. This simple feedback control has been shown to work for a quasi one-dimensional example of heat flow down a three-dimensional rod of uniform cross-section. These boundary conditions are of use for analysing the heat transfer across nanoscale features [88] such as grain boundaries [73], nanowires [72][89][90] and nanoconstrictions [70]. As the model does not rely on the potential, there is a possibility that the potentials could be altered to incorporate quantum effects [91]. The method for controlling the boundary conditions of NEMD simulations has been extended in section 4 to allow atomistic/continuum models to be thermally coupled concurrently for the analysis of steady state and transient heat conduction problems. The effectiveness of this algorithm has been demonstrated through a number of examples.

Section II

In this section we look at mass transport in solids by surface diffusion. The timescale of surface diffusion is far beyond the capabilities of the MD simulation method used in section I because it consumes far too much computer time. The Kinetic Monte Carlo (KMC) method is introduced in chapter 5, as well as a comparative continuum KMC method. In chapter 6, the simple case of the zero strain on-lattice models are investigated in order to compare the models using simple systems. Single component ellipses are simulated using both methods. Since quantum dots are formed by the interaction between two different materials, the KMC algorithms are then updated to include multiple atom types in subsection 6.2. Two atomic species are differentiated by changing the jumps rates per bond in the KMC algorithm. Two component ellipses are considered for a variety of interesting cases. The results of these simulations are used to update the continuum KMC method.

Lastly in chapter 7, elastic strain is added to the KMC algorithms. This is a requirement in order to see quantum dot formation. An investigation of surface roughening is performed first in order to ascertain the optimal conditions for performing the atom deposition simulations. The formation and arrangement of quantum dots upon a surface during molecular beam epitaxy experiments are simulated. Finally multiple layers of material and substrate are deposited in order to investigate vertical ordering of quantum dots into a super-lattice.

5 Introduction to surface diffusion and the KMC method

Solid surfaces evolve over time to reduce the free energy of the solid if there is an active kinetic mechanism for it to do this. The dominant diffusion process is assumed to be surface diffusion, which generally has an activation energy lower than lattice or grain boundary diffusion. At the nano-scale, the surface area is large compared to the volume and hence interfacial energies act as a strong driving force for surface diffusion. Also, at this scale the lattice structure of the solid has a strong influence on the shape of crystalline solids, with a preference for forming angular surfaces aligned along particular low energy facet directions.

Atomistic processes such as surface diffusion occur on a relatively large timescale compared to the oscillation of the atoms themselves. In order to simulate such processes using MD, one would have to solve many thousands of atomistic oscillations before even one surface atom hops to its neighbouring site. As has been said many times, such simulations would take prohibitively too long to run.

A much faster method than MD is required to fully investigate real nano-scale devices such as nano-wires, nano-whiskers and quantum dots (QDs) etc (see section 1). The formation and dynamics of these devices occurs via surface diffusion at time-scales many orders of magnitude longer than MD is currently capable of simulating. Furthermore, information of every atom oscillation is not required for understanding of the structure of nano-devices like QDs. The Kinetic Monte Carlo (KMC) method is most suited to surface diffusion problems and can handle many thousands of atoms for much longer time scales than MD. However KMC still has its limits and drawbacks. An atom on a surface will perform many jumps meandering around before finding its minimal energy position. A comparative continuum KMC method is therefore introduced which attempts to reduce unnecessary kinetics. The goal of this section is to develop these KMC methods in order to simulate the formation of QDs on surfaces and to investigate vertical alignment of multiple layers of QDs.

5.1 Quantum dot formation

The KMC models developed will be used to simulate QD systems. It is therefore important to review real QD systems in order to know what to expect from the computer simulations.

QDs are produced by depositing a layer of material (such as GaAs) upon a substrate material (such as InAs) by a process called molecular beam epitaxy (MBE). The substrate is heated to constant temperature in a vacuum, while a crucible containing the deposition material is also heated until it vaporises. This vapour rains down upon the substrate forming a compressed film on surface. Figure 5.1 illustrates this process.




Fig. 5.1 Diagram of heteroepitaxial growth on a surface. (Left) a vapour of the deposition material rains down on the substrate, (right) illustration of the mismatch in strain causing compressive strain in the film.

The important point is that the two materials have a different lattice spacing. If the deposition is epitaxial (such that the deposit adopts the atomic configuration of the substrate forming a coherent interface) this introduces an interfacial mismatch strain $\epsilon_T = \frac{a_s - a_f}{a_s}$ in the deposited film, where a_f is the lattice spacing of the film and a_s is the lattice spacing of the substrate. A high strain state is a state of high energy. The film seeks to minimise its energy by roughening the surface in order to reduce the total elastic strain energy, however this increases the surface energy of the film.

In the Stranski-Krastanov growth mode, a wetting layer (typically a few monolayers thick) forms first. At a critical thickness, the flat layers become unstable and the atoms begin to cluster together to form QDs. QDs form as a result of the competition between elastic strain energy and surface energy. Figure 5.2 shows some QDs on a surface with a thin wetting layer. The atoms at the peak of the QD have relaxed back to their equilibrium spacing whereas atoms near the substrate are under high compression (e.g. -4.2% for Ge/Si (110)).




Fig 5.2 A film of Germanium deposited upon a Silicon substrate. QDs form due to the competition between surface energy and elastic strain energy.

The QDs formed are all roughly the same size as seen in experiments (figure 5.3). This preferred size arises from the competition between the surface energy (which prefers large dots) and the elastic strain energy (which prefers small dots). QDs are therefore considered to be highly self organised in that they have a very narrow size distribution.




Fig 5.3 AFM topographies ($0.9 \times 0.5 \mu\text{m}^2$) for (a) 1.54 Mono-layers (ML), (b) 1.57 ML and (c) 1.64 ML of InAs coverage. Panel (d) shows the number density dependence on InAs coverage of small and large QDs [92].

Quantum dots are useful structures for nano-scale devices. Since the substrate and QD have a different valence band energy and conduction band energy, the QDs formed can confine electrons in a similar way to atoms. By controlling the size and shape of the quantum dots, one has direct control over the energy levels of the confined electrons and holes. Hence this provides a means of "tuning" the resulting wavelength of the nano-scale device simply by changing the material composition of the substrate and QD.

The deposition rate of material on a surface strongly affects the formation and distribution of QDs. For example, for InP islands on a GaInP/GaAs surface, a high deposition rate (2.6 Mono-layers per second (ML/s)) yields islands with a height of $290 \pm 12 \text{Å}$ whereas for low (0.1 ML/s) and moderate (0.8 ML/s) deposition rates the islands are slightly shorter at $250 \pm 16 \text{Å}$ [93]. In general, a high deposition rate yields many QDs forming close together, whereas a low deposition rate results in fewer QDs that are distributed sparsely over the surface.

Once the QDs have formed, more substrate material can be deposited on top of them. This results in some intermixing of the material from the QD with the substrate. Primarily, the material at the peak of the QD mixes with the substrate material that is raining down,

resulting in the removal of the top of the QD. This occurs because these atoms that were once relaxed become increasingly strained as they come into direct contact with the substrate material as it builds up around them.

A buried QD causes tension in the capping layer above it. This is then a favourable site for another QD to form, and hence a vertical stack of QDs (as seen in figure 5.4) can be made by depositing alternating layers of substrate and QD material.




Fig. 5.4 Vertical stacks of Germanium quantum dots [94].

This is also an example of self-organisation, where a system exhibits long range order. The ‘holy grail’ of nano-engineering is to be able to obtain useful nano structures like arrays of transistors simply by mixing some chemicals.

5.2 The KMC algorithm

The Kinetic Monte Carlo method is an extension of the Monte Carlo (metropolis) algorithm to model the kinetics of a system. The metropolis algorithm is used to find the minimum energy state of a system only, whereas the kinetic Monte Carlo method can compute the path taken and the timescale required to get there.

The KMC method has been used extensively over the years. It has been used for simulating MBE experiments, where a material is rained down upon a flat substrate [95].

Using accurate chemical kinetic data and allowing for adsorption and desorption from a surface, KMC has been used to investigate the interaction of a surface with a hot vapour [96].

The KMC method is a stochastic method, which does not simulate the atomistic vibrations, but instead allows atoms to jump to the neighbouring sites based on given transition rates. In the simplest on lattice method, the rates just depend on the number of nearest neighbours at the start position and the temperature. The details of the kinetic algorithm are described below.

Firstly, simulations in section 6 were performed using the simple *on-lattice* model at zero strain. All the simulations were performed in 2D using a simple triangular lattice where each atom has six nearest neighbours all at an equal distance from the centre site and each other. This represents a cross-section through a simple HCP crystal lattice and is depicted in figure 5.5.




Fig. 5.5 Labelling of the six nearest neighbour sites (A to F) for an atom at a site (denoted O) in a fixed triangular lattice.

The simulations are performed in 2D. In the *on-lattice* model, the lattice sites are fully determined before the simulation. Using a simple 2D array, the neighbour sites can be easily derived. Figure 5.6 illustrate the numbering system and neighbour site locations for an atom at site (i, j) .




Fig. 5.6 Relative coordinates of neighbours for an atomic site (i, j) in an: (a) even column and (b) an odd column, for a fixed triangular lattice.

The first step in the KMC algorithm is to identify all the possible atom jumps that can occur. Obviously each atom can only jump to a neighbour site if that site is currently empty. Allowed jumps are defined by a set of rules:

- 1 Atoms are prevented from isolating themselves from the others as a result of a jump. E.g. Atom 5 cannot move to site 1 in figure 5.7 because site 1 has no other neighbours.
- 2 An atom cannot isolate any others atoms. E.g. atom 7 cannot move to site 6 as this will leave atom 3 isolated from the others.
- 3 If there is a site directly below an atom which is available, the atom is forced to move there. This prevents high energy structures such as overhangs from occurring and also prevents clusters of three or more atoms from escaping from the surface. This effectively simulates gravity. This rule overrides rules 1 and 2. For example, atom 8 will move down to site 12 even if doing so isolates other atoms as this rule will force those to fall on the next step.




Fig. 5.7 Diagram of a typical high energy surface of a KMC simulation. The filled hexagons (blue) are occupied sites and unfilled sites are white. Some sites are numbered for reference.

Once a valid atom jump has been determined, it is assigned an activation energy E equal to Kn , where K is the bond energy and n is the number of nearest neighbours (1 to 6). This energy depends only upon the initial state of the model. The jump rate of event k is then calculated from:

$$R_k = v_0 \exp\left(\frac{-Kn_{(i,j)}}{k_b T}\right), \quad (5.1)$$

where, $v_0 = 1$ is an attempt frequency and $n_{(i,j)}$ is the number of neighbours of site (i, j) . k_b is Boltzmann's constant and T is the temperature.

The rates are then listed in no particular order, forming the events table. A cumulative sum and total sum of the rate is then calculated from:

$$R_c = \sum_{k=0}^c R_k, \quad (5.2)$$

where $c = 1, \dots, N$ and N is the total number of atom jump events. Denote R_T as the total sum of the rates. This cumulative sum is saved in the event table, which is now automatically sorted by this parameter. This is important as it allows the use of the binary search algorithm on the events table. After this setup stage, the KMC algorithm proceeds as follows:

- 1 A uniformly distributed random number U is selected where $0 \leq U \leq R_T$ using the Mersenne Twister algorithm [97].
- 2 The corresponding event is found in the list via the binary search algorithm. This algorithm proceeds by picking the event in the middle of the table and asking if the selected event is the correct one or higher/lower in the table. At each step half of the table is eliminated and the next mid point in the valid range is tested. This continues until the correct event is found and will take at most $\log_2 N$ tests to find the correct event.
- 3 The selected event is then executed by updating the occupation flags of the atom sites.
- 4 The events that have now changed as a result of the transition must be recalculated and added/removed from the event table as appropriate. Furthermore the cumulative sum and total sum of the rates will have to be updated.
- 5 The elapsed time is then updated by an amount:

$$\Delta t = -\frac{\ln U}{v_0 R_T} \quad (5.3)$$

- 6 Repeat step 1 - 5 for as long as is needed.

This algorithm is known by various names such as the *residence-time algorithm* or the *n-fold way* or the *Bortz-Kalos-Liebowitz (BKL)* algorithm or just the *kinetic Monte Carlo (KMC)* algorithm.

Optimisation of KMC code

All the random numbers used in this work were generated using the Mersenne Twister (MT) algorithm [97]. This was found to be faster than the standard C++ random number generator. The C++ random number generator only picks an integer between 0 and 32767, whereas the MT algorithm can pick any number between 0 and 4 billion. MT is capable of providing fast and high quality random numbers.

Once an event has been picked and executed, the events table needs to be re-calculated. One could simply cycle through the whole simulation and re-calculate all the rates,

but this is rather inefficient. A better way is to maintain a list of which events correspond to each atom such that they can be found and easily deleted once they become invalid. Then the list just needs to be updated with new events related to the atoms that moved, and their nearest neighbours. This data is then slotted into the events table, overwriting invalid events. Since the KMC model is an initial state model, the movement of one atom will affect at most 11 neighbouring atoms (not 12 because one site needs to be empty for a movement in the first place).

Therefore we now have an updated table without having to re-calculate events that have not changed. This lead to a vast increase in computational speed in our zero strain on-lattice KMC.

In the KMC simulations with strain, this is not as easily implemented because the movement of one atom will affect the strain field in the whole crystal. This will consequently change the jump rates of all local atoms.

5.3 Continuum KMC model

Given that an atom on a surface has a 50/50 chance of moving left or right, it is observed that a surface atom does not move directly to its minimum position. But rather meanders around on the surface until it finds its minimum by chance. Therefore, a large amount of the simulation time is taken up by atoms doing this.

A continuum model [98] replaces this with an atom flux which depends on the same KMC rules. Care must be taken to ensure the timescale is still correct between simulation methods. The continuum model is derived from the atomic model and the similarities and the differences between them are compared in section 6.

The Continuum Model

In this model an occupation probability, $p(i, j)$, is assigned to each lattice site (denoted by 2D coordinate indices i and j) and allow it to evolve as a continuous variable between 0 and 1. Therefore a lattice site can be partially occupied and the solid surface is no longer necessarily sharply defined but can be diffuse. It is found that the width of the diffuse interface (distance between fully solid ($p=1$) and vacuum ($p=0$) regions) increases with

temperature, reflecting the increase in the randomness of the surface. Referring to figure 5.5, if atom O is at (i, j) and atom A is at position (i_A, j_A) etc. then:

$$\frac{dp(i, j)}{dt} = \left[\begin{aligned} &w_D(i_A, j_A) + w_E(i_B, j_B) + w_F(i_C, j_C) \\ &+ w_A(i_D, j_D) + w_B(i_E, j_E) + w_C(i_F, j_F) \end{aligned} \right] - [w_A(i, j) + w_B(i, j) + w_C(i, j) + w_D(i, j) + w_E(i, j) + w_F(i, j)] \quad (5.4)$$

where $w_A(i, j)$ is the net transfer rate from site O to site A (in the direction OA), and $w_D(i_A, j_A)$ is the net transfer rate from site A to site O (in the direction OD) etc. The net transfer rate is the sum over all the possible configurations of the individual transition rates weighted by the configuration probability. The individual transfer rates are determined from the occupancy probabilities of the neighbouring sites and the number of nearest neighbour bonds, e.g. the transfer rate from O to A if only B and D are occupied is: $\nu_0 p(i, j)(1 - p_A)(p_B e^{-2K})(1 - p_C)(p_D e^{-2K})(1 - p_E)(1 - p_F)$, where $p_A = p(i_A, j_A)$ etc. The net transfer rate is the sum of all such individual rates over the 64 (2^6) possible configurations, and can be written as:

$$\begin{aligned} w_A(i, j) &= \nu_0 p(i, j) n_A m_C m_D m_E (m_F m_B - n_F n_B) \\ w_B(i, j) &= \nu_0 p(i, j) n_B m_D m_E m_F (m_A m_C - n_A n_C) \end{aligned} \quad (5.5)$$

etc, where $n_A = 1 - p_A$ is the probability that site A is unoccupied and $m_A = p_A \exp(-K/k_B T) + n_A$ is the effect on the transfer rate due to the bond OA. In each case the final bracketed term enforces the rule that the contribution to the transfer rate due to hopping from site O to A (say) when the adjacent sites (B & F) are empty is removed. In the integer limit of all $p=0$ or 1, the transfer rates reduce to those of the kMC algorithm.

6 KMC Simulations

In this section, the results of the *on-lattice* KMC simulations are presented and compared to the continuum KMC method. Firstly single component clusters are simulated and then the more interesting case of two component clusters is investigated.

6.1 Single component simulations

The initial configuration (figure 6.1) is a two dimensional ellipse of 1629 atoms, with a height of 16 atoms and a width of 130 atoms. This system was simulated using the KMC algorithm described in section 5.2 and the continuum algorithm presented in the previous section (5.3). The ellipse floats in free space in the computer simulation and hence rule 3 (Gravity) is not in effect. The attempt frequency, ν_0 , is set to 1 in these simulations. This parameter has no effect on the evolution of the system, only the timescale associated with it.

The surface atoms have a higher energy than atoms in the bulk simply due to having fewer bonds. It is therefore expected that a collection of atoms will minimise their surface energy by forming a circular nano-cluster.




Fig. 6.1 A depiction of the initial configuration of the single component KMC simulation. The starting shape in all simulations in this section is an ellipse consisting of 1629 atoms (width 130, height 16).

The exponential term in equation 5.1 can be varied in order to simulate the behaviour of the system at different temperatures. The bond energy and temperature are wrapped up into

a single variable, the rate constant, $R_C = K/k_b T$. The total jump rate for each atom now becomes:

$$R_k = \exp(-R_C n_{(i,j)}), \quad (6.1)$$

The KMC ellipse was simulated at a high temperature ($R_C = 2.0$) and a low temperature ($R_C = 4.0$). Snapshot of the KMC simulations at high and low temperature are presented in figure 6.2.

The shape of the nanocluster is tracked quantitatively by measuring the second moment of area of the system in both the x and y directions. The second moment in the x direction is given by:

$$I = \sum_{n=0}^N (x_n - \langle x \rangle)^2. \quad (6.2)$$

where x_n is the x position of the n^{th} atom, $\langle x \rangle = \frac{1}{N} \sum_{n=0}^N x_n$ is the average x (the centre) and the sum is over all N atoms in the simulation. A similar equation applies for the y direction.




Fig. 6.2 Snapshots of a KMC simulation of an elliptical nano-cluster evolving via surface diffusion to its equilibrium shape which is a circle. The high temperature ($R_C = 2.0$) results are shown in panels (a)-(d) and the low temperature ($R_C = 4.0$) results are shown in panels (e)-(h). The snapshots are taken at 10%, 25%, 50% and 100% of the total simulation time for both cases. At the high temperature a total of 40 million time steps (equivalent to a time of 3.11×10^7 in arbitrary units given by equation 5.3) were required to obtain the final state. At low temperature the total simulation time required was 15 million time steps (equivalent to a time of 6.0×10^9 in arbitrary units).

High temperature KMC simulation, $R_c=2.0$

The high temperature simulation results are presented in panels (a) – (d) in figure 6.2 (left column). The high curvature edges of the ellipse initially become rounded in order to

reduce the surface energy. The system requires 40 million discrete atom jumps to converge to its lowest energy state which is a circle. At this high temperature the atoms are highly mobile and can explore high energy states. As expected, there is a large amount of thermal noise indicated by the highly irregular surface. Furthermore, some clusters of 2 or more atoms can break away from the surface at these high temperatures.

The plot of the second moment of area vs. time for the high temperature case is shown in Figure 6.3a. The initial ellipse has a large moment in the x direction but a small moment in the y direction, as expected due to the elongated starting shape. Both values converge to the same value of 110, indicating that the shape formed is a circle.




Fig. 6.3 Plots of the 2nd moment of area vs. time for the KMC simulations shown in figure 6.2 at, (a) high temperature and (b) low temperature. The plots show the second moment in both the x (bright red) and y (dark red) directions. In both cases the 2nd moment in x and y converges to a constant value of 110 indicating that the final state is a circle. The high temperature ellipse converges to a circle 200 times faster (using calculated elapsed time from equation 5.3) than the low temperature ellipse. However, twice as many discrete atom jumps are required.

Low temperature KMC simulation, $R_c=4.0$

The low temperature simulation results are presented in panels (e) – (h) in figure 6.2 (right column). The KMC ellipse required 15 million discrete atom jumps to converge to its lowest energy state which is also a circular nanocluster. The high curvature edges of the ellipse become rounded as before. At this low temperature there is very little thermal noise compared to the high temperature case. Some flat regions are visible in the low temperature snapshots; these are the crystal facets which are fairly stable at this temperature.

The plot of the second moment of area vs. time for this low temperature case is shown in figure 6.3b. The curve is similar to the high temperature case. Due to the low temperature each discrete atom jump takes a greater amount of time (calculated using equation (5.3)) to occur. Despite only needing half the number of discrete atom jumps, the low temperature simulations take a real time of 200 times longer to converge. But due to the lower amount of atom jumps taking place, the computer simulation is faster here.

Comparison of the atomistic and continuum models

The initial elliptical body shown in figure 6.1 is shown here again in the centre of figure 6.4. This body has a high surface-to-volume ratio and so evolves into a shape with a smaller surface area (circle) to reduce its surface energy by surface diffusion. The final shapes for the continuum and standard KMC models are also shown in figure 6.4. Both the high temperature case ($R_C = 2.0$, close to melting point) and the low temperature case ($R_C = 4.0$) are shown.




Fig. 6.4 An elliptical body evolves into a (near) circular one to reduce its surface area. The initial elliptical body is shown along with the final equilibrium states for two different temperatures ($K/k_B T = 2.0$ (top) and 4.0 (bottom)) and three different models (atomistic (left), basic continuum (middle) and adapted continuum (right)). In the continuum models, partially occupied sites are shown by atoms with a size proportional to their occupation probability.

The surface of the atomistic model is very rough at the high temperature, whereas the basic continuum model (defined by equations (5.4) and (5.5), shown in the middle images) is smooth. Fig. 6.5 shows the evolution over time of the second moment of area (in the horizontal and vertical directions for each temperature and model. Figure 6.5a illustrates that, although the atomistic surface is more dynamic, its average shape evolves in a similar way to the basic continuum model. However, at the low temperature, the behaviour of the two models is quite different. The atomistic model reaches a near circular shape as before, but the basic continuum model gets trapped in a higher (elongated) metastable energy state. This is because the facets are stable without any thermal fluctuations. For this reason, an adapted continuum model is proposed. Firstly, the transition rates in (5.5) are given the same random distribution that occurs in the KMC algorithm, i.e. w_A is replaced by $-w_A \ln U$ in (5.4), where $0 \leq U \leq 1$ is a uniformly distributed random number. The random number is updated after the total corresponding flux due to this event is equivalent to the movement of a single atom.




Fig. 6.5 A quantitative assessment of the evolution (shown in Fig. 6.4) shows the change in the second moment of area of the elliptical body against time for (a) high temperature ($R_c = 2.0$) and (b) low temperature ($R_c = 4.0$). The upper and lower sets of curves are the second moments along the horizontal and vertical axes respectively. At high temperatures all three models behave in a consistent way. At lower temperatures the basic continuum model gets trapped into a faceted state. The adapted continuum model artificially introduces thermal fluctuations and promotes nucleation of new facets. This agrees very well with the atomistic model. Note that the stochastic simulations are for one run and hence there will be some statistical variation in these curves.




Fig. 6.6 The probability bias function, $h(p)$, as a function of site occupation probability, p .

This is used in the adapted continuum model to promote nucleation.

However, adding this randomness to the algorithm is not sufficient to achieve a continuum model that is comparable with the atomistic one. Although the randomness destabilizes the facets and the transition to a (near) circle is achieved, the process is very slow. It is thought that this is because nucleation of new steps/facets on top of the long flat facet is very important for the transition to a circle at low temperatures. To introduce this “discrete” feature into the continuum model, all the probabilities used in the calculation of the transition rates, p , are pre-processed by a biasing function, $h(p)$, i.e. p_A is replaced by $h(p_A)$ in (5.4), where $h(p_A) = \frac{1}{2}((2p-1)^n + 1)$, and n is an odd, positive integer. For $n=1$ the model is as before. For $n=3,5,7$, the occupation probability is biased as shown in Fig. 6.6. This accentuates the occupation probability for sites that have a small probability of being occupied, destabilizes the uniformity of occupation probabilities along facets and hence promotes nucleation. As with randomness, it is interesting to note that on its own this biasing contribution does not help the continuum model to reproduce the average behaviour of the atomistic model. In fact, biasing on its own can make the predictions of the continuum model much worse. However, when the randomness and biasing are applied together, the features of the atomistic model are reproduced by this adapted continuum model at high and low temperatures. This is shown in Fig. 6.4 and Fig. 6.5 for $n=3$. Encouragingly, the adapted continuum model is not sensitive to the value of n , with very similar results being obtained for $n=7$.

Summary

The simple case of the single component ellipse was studied. It was found that in both methods the ellipse evolves towards the minimum energy state which is a circle in 2D. At high temperatures, both algorithms are in agreement (Figure 6.5a shows that the second moment of area vs. time curves are in good agreement with each other for both algorithms). However, at low temperatures the continuum method becomes trapped in a metastable state. Since the continuum method does not contain any noise it is unable to obtain the correct solution. It is found that randomness and biasing need to be added to the continuum method to reproduce the behaviour of the atomic model at low temperatures.

Given the success of the single component system, we proceed to compare the atomistic and continuum models for two component systems.

6.2 Two component simulations

Quantum dot formation is the result of differing material properties between substrate and deposition material. The KMC simulations are therefore extended in this subsection to include multiple atom types. In this subsection, simulations of two component ellipses are performed in order to compare the KMC and continuum KMC methods.

Introducing multiple atom types is a simple case of updating the transitions rates, based on pre-defined rules. With two atom types, there are three kinds of bonds between them. K_{11} is the bond energy of atoms of type 1 with similar atoms. K_{12} is bond energy between atoms of type 1 and 2. K_{22} is bond energy of atoms of type 2 with similar atoms. If all these constants are equal, then the result is the same as the one component simulations shown earlier.

In the two component simulations, the jump rates for an atom of type 1 are calculated using:

$$R_k = v_0 \exp\left(-\sum_{n=0}^5 K_{1n_{ID}}\right). \quad (6.3)$$

Here the rate constant depends on the identity of the atom moving and its current neighbours. The sum is over the six neighbour sites, n , and n_D is the identity (1 or 2) of neighbour n . Obviously, this constant is zero for unoccupied neighbour sites. A similar equation applies for the jump rates of atoms of type 2.

The initial configuration (figure 6.7) of the two component simulations is an ellipse the same size as the single component simulations (130 x 16 atoms). However, this time the left half (red in fig. 6.9) are type 1 atoms, and the right half (blue) are type 2 atoms.




Fig. 6.7 The initial configuration of the two component simulations is shown here. In all cases in this section the starting state is an ellipse as before, with the left half (red atoms) are type 1 atoms and the right half (blue atoms) are type 2 atoms.

The behaviour of the two component system is strongly dependent on the choice of the values for the three different rate constants. Table 6.1 shows the rate constants used for each case presented in this subsection. The values in the table are specific to the high temperature case ($T=1$). The temperature of the low temperature simulations is half that of the high temperature simulations ($T=\frac{1}{2}$). Since the rates are inversely proportional to the temperature, the rates used at low temperature can be obtained by doubling the values in table 6.1.

	K_{11}	K_{12}	K_{22}
Case A	2.4	2.0	2.4
Case B	2.0	2.4	2.0
Case C	2.0	2.4	2.8

Table 6.1 Values of the rate constants for each case at the high temperature ($T=1$) regime. For the low temperature simulations ($T=\frac{1}{2}$) these values are simply doubled.

Case A at high temperature, $T=1$

In case A, the bond energy of each atom type with similar atoms (K_{11} and K_{22}) is lower than the bond energy between mixed atoms (K_{12}). Therefore each atom type preferentially bonds to itself rather than mixing. The rate constants are shown in table 6.1.

The high temperature results for the standard KMC are shown in Figure 6.8 panels (a) – (d) while the final state obtained by the continuum KMC model is shown in figure 6.8.e. The final states of both models are very similar (see pictures in figures 6.8.d and 6.8.e). Here the edges of the ellipse quickly become more rounded in order to minimise the surface energy. The ellipse collapses inwards in a similar way to the single component simulations to form a circular shape.

Due to the high temperature, an atom that exists on the opposite side is not necessarily picked to move again. As a result of this, other atoms can diffuse around the surface and bury the “foreign” atom. Only the surface atoms can move, therefore a buried atom can remain buried forever. Some mixing of the internal areas of the cluster occurred due to voids propagating through the cluster. The standard KMC simulation exhibits a higher degree of mixing of the atom types than the continuum case.




Fig. 6.8 Case A snapshots of a standard KMC and continuum KMC simulation of an elliptical nano-cluster evolving via surface diffusion to its equilibrium shape are shown. The left column (panels (a) – (e)) shows the high temperature results and the right column (panels (f) – (j)) shows the low temperature results. High temperature standard KMC snapshots are shown at (a) 10% (b) 25% (c) 50% and (d) 100% of a total simulation of 60 million time steps (time = 1.08×10^8 in arbitrary units given by (5.3)). Panel (e) shows the state of the continuum KMC simulation at high temperature at time of 5×10^7 . Low temperature standard KMC snapshots are shown at (f) 10% (g) 25% (h) 50% and (i) 100% of a total simulation of 20 million time steps (time = 4.52×10^{10}). Panel (j) shows the metastable state of the continuum KMC simulation at low temperature at a time of 2.5×10^{10} .

Figure 6.9 shows the second moment of area vs. time of the whole system for both models at high temperature. The curves are of a similar form, but the standard KMC has a higher second moment in the y direction. This higher moment is a direct consequence of the greater amount of mixing that occurs in the standard KMC simulations.




Fig 6.9 Comparison between KMC and the continuum KMC of the second moment of area vs. time for case A at high temperature.

Case A at low temperature, $T=1/2$

The right column of figure 6.8 (panels (f) – (j)) shows snapshots of the low temperature case A simulation. As before, the system evolves towards its minimum energy shape which is a circle, with the atoms not mixing much in this case. The high curvature edges of the ellipse quickly become more rounded in order to reduce their surface energy. The final state appears less rounded than the higher temperature simulations due to the crystal facets being more stable at this temperature. Very little mixing of the atom types has occurred. As can be seen from the above pictures, a small number of atoms did become buried on the opposite side. This is simply because buried atoms cannot move, and once they are buried by at least two layers of atoms, they no longer affect the transition rates.

The time scale of the simulation, as measured by (5.3) is many orders of magnitude higher than the high temperature case. However, the final state was obtained with only a third of the discrete atoms jumps taking place.

The second moment of area of the whole system as well as each atom type individually is plotted in figure 6.10a. The pink and blue lines which represent the whole

system converge to a single value as expected. The second moments of the individual components do not converge to the same value, which is expected as each atom type forms a semicircle which has a greater second moment in y than x . Figure 6.10b shows the comparison of the second moment of area of the whole system for the continuum KMC and standard KMC models. This plot shows that the continuum KMC model has not converged as much as the standard KMC model did in the time it was run, and the gradient of the graph indicates that it will not attain the final state for a long time, if at all. Figure 6.8j shows that the continuum model at low temperature has become trapped in a metastable state once again. In the single component model, the additions of randomness and biasing allowed the continuum model to avoid becoming trapped in a metastable state. In this case, these additions have not been sufficient to avoid this. This indicates that two-component systems are more complex than single component ones, and that the continuum model is not likely to be directly comparable with the atomistic model at low temperatures. This conclusion is reinforced by the results for the other cases in the following subsections.




Fig. 6.10. (a) Graph of second moment of area of each component in x and y for the KMC simulation. Blue and pink curves are the second moment of the whole system. These both tend towards a constant value of 110, which is the expected result for a circle. The second moment of each component becomes lower in x and higher in y , which is expected as these form semi-circles. (b) Comparison between KMC and the continuum KMC of the second moment of area vs time.

Figure 6.11 shows the result of a KMC simulation of case A at very low effective temperature ($T = \frac{1}{6}$). Here the system is highly faceted and exhibits no mixing at all. An atom that exists on the opposite side (in this case the blue atom at the top left corner) will

have an energy far higher than any other atom in the system, and hence will always be picked to move again until it is back on its own side.




Fig. 6.11 Snapshot of the final state after 30 million discrete atom jumps of a case A system at very low temperature $T = \frac{1}{6}$. Here absolutely no mixing occurs at all.

Case B at high temperature, $T=1$

Case B is the opposite of case A, where intermixing of different components is preferred. Snapshots of the simulation at high temperatures are shown in the left column of figure 6.12. As soon as a surface atom moves past the interface between atom types, it enters a far lower energy state. This results in a well mixed cluster of atoms forming a bulge in the centre (figure 6.12a). The final equilibrium state (fig 6.12d) occurs when all the outer atoms become mixed leaving a small area of ‘pure’ material buried in the centre. Interestingly this is not the minimum energy that this system can achieve. The minimum would be all the atoms completely mixed together. The system never attains this as the centre becomes frozen in by the surrounding material.




Fig. 6.12 Case B snapshots of a standard KMC and continuum KMC simulation of an elliptical nano-cluster evolving via surface diffusion to its equilibrium shape are shown. The left column (panels (a) – (e)) shows the high temperature results and the right column (panels (f) – (j)) shows the low temperature results. High temperature standard KMC snapshots are shown at (a) 10% (b) 25% (c) 50% and (d) 100% of a total simulation of 28 million time steps (time = 4.5×10^7 in arbitrary units given by (5.3)). Panel (e) shows the final state of the continuum KMC simulation at high temperature. Low temperature standard KMC snapshots are shown at (f) 10% (g) 25% (h) 50% and (i) 100% of a total simulation of 4 million time steps (time = 3.6×10^9). Panel (j) shows the final state of the continuum KMC simulation at low temperature.

Both simulation methods yield an outer fully mixed region with an unmixed central region (see figures 6.12d and 6.12e). Slightly more mixing occurred in the standard KMC simulations than the continuum model. The second moments of area vs. time of each method are shown in Figure 6.13. The standard KMC results have a slightly higher second moment in y due to the slightly increased mixing; however both plots are of similar form.




Fig. 6.13 A plot of the second moment of area vs. time for case B high temperature simulations is shown. Results from the KMC method and continuum KMC method with biasing and noise are plotted.

Case B at low temperature, $T=1/2$

The low temperature simulations proceeded in a similar fashion to the high temperature case. Again the low temperature simulation required less discrete atom jumps but a higher simulation time was measured. A fully mixed outer region with an unmixed centre was obtained by both models. However the outer surface of the continuum KMC simulation was highly faceted (figure 6.12e) whereas in the standard KMC model it was more circular. This lead to a vast difference between the plots of the second moment of area for the two models (figure 6.14). The plots of the second moment do show that the final state was obtained in a similar time for both models. However, comparison of the final states in Figures 6.12i and 6.12j show that they are qualitatively similar, even if the quantitative comparison is not.




Fig. 6.14 A plot of the second moment of area vs. time for the case B low temperature simulations is shown. Results from the KMC method and continuum KMC method with biasing and noise are plotted.

Case C at high temperature, $T=1$

In both case A and B, each atom type had the same properties when interacting with atoms similar to itself. Here, the most stable bonds are between type 2 atoms only. The bonds between mixed atom types have a lower stability and the bonds between type 1 atoms are the least stable of all. Therefore, type 1 atoms are far more mobile than the type 2 atoms and will seek to mix with the type 2 atoms to minimise their energy.

Figure 6.15 (left column) shows that the type 1 atoms (red) quickly surround and mix with the type 2 (blue) atoms and the high curvature edges become more rounded as before. Both models attain a similar final state with a mixed outer region containing mostly red atoms and a pure blue centre.




Fig. 6.15 Case C snapshots of a standard KMC and continuum KMC simulation of an elliptical nano-cluster evolving via surface diffusion to its equilibrium shape are shown. The left column (panels (a) – (e)) shows the high temperature results and the right column (panels (f) – (j)) shows the low temperature results. High temperature standard KMC snapshots are shown at (a) 10% (b) 25% (c) 50% and (d) 100% of a total simulation of 60 million time steps (time = 6.8×10^7 in arbitrary units given by (5.3)). Panel (e) shows the final state of the continuum KMC simulation at high temperature. Low temperature standard KMC snapshots are shown at (f) 10% (g) 25% (h) 50% and (i) 100% of a total simulation of 30 million time steps (time = 1.3×10^{10}). Panel (j) shows the state of the continuum KMC simulation at low temperature at a time of 1×10^{10} .

The plot of the second moment (figure 6.16) shows an initial rise for the type 1 atoms as they spread out around the system. The plots for both models are again very similar at high temperature.




Fig. 6.16 A plot of the second moment of area vs. time for the case C high temperature simulations is shown. Results from the KMC method and continuum KMC method with biasing and noise are plotted.

Case C at low temperature, $T=1/2$

The right column of figure 6.15 shows that the type 1 atoms (red) quickly surround the type 2 (blue) atoms as they did in the high temperature simulations. Less mixing occurs in both models at low temperature. The continuum KMC model again has problems at this low temperature and does not attain the final state in the time that the standard KMC model did. The plot of the second moment (figure 6.17) indicates that the evolution of the shape is very similar in both models, but the continuum model evolves too slowly.




Fig. 6.17 A plot of the second moment of area vs. time for the case C low temperature simulations is shown. Results from the KMC method and continuum KMC method with biasing and noise are plotted.

A similar set of simulations were also performed for the inverse problem of an elliptical void in a two component material for each of the three cases A, B and C at high and low temperatures. These simulations were only performed using the standard KMC method and therefore were not used to compare the KMC models. These results are presented in appendix C.

Summary

In this subsection the standard KMC and continuum KMC algorithm were used to simulate two component ellipses by surface diffusion. The simulations were performed at high and low temperatures and the results compared in order to validate the continuum method.

Both models attain similar final states in nearly all cases, except case A at low temperature where the adapted continuum model became trapped in a metastable state. However, when the plots of the second moment of area vs. time are compared, there are significant differences in some cases. Low and high temperature case A simulations resulted in similar structures forming, but far more mixing occurred in the standard KMC simulation than the continuum. This caused a slight difference in the second moment of area of each component (see figure 6.9 and 6.10). For case B, the results compare well at high temperature, the central region remains unmixed with an outer coating of fully mixed material. For the continuum at low temperature a similar result occurs but the outer surface becomes highly

faceted which is not seen in the standard KMC. This caused a vast difference in the second moment at this temperature (see figure 6.14). In Case C, the high temperature simulations agreed well with each other (figure 6.16). However, at a low temperature, the continuum method evolved more slowly than the standard KMC. This is most likely due to the system becoming trapped in a meta-stable faceted surface.

In conclusion the continuum KMC model is a good representation of the underlying atomistic KMC model in the high temperature regime. It attains the correct final state in all cases and evolves in a similar way as the standard KMC simulations. This is where it is most useful in eliminating unnecessary atom jumps and hence reducing simulation time. At low temperatures the agreement is not so good. The continuum model can become trapped in metastable states, and has crystal facets that are more stable than they should be. However, the final states obtained are often qualitatively very similar. Now, we proceed to introduce another driving force, elastic strain energy, in to the model in the next section.

7 Surface diffusion with elastic strain

Real mechanical systems contain stresses and strains, which significantly alter its behaviour. In order to simulate processes like heteroepitaxial deposition, strain needs to be included in the KMC algorithm in order to see effects like the formation of quantum dots (QDs) (see chapter 5).

In this chapter elastic strain is added to the KMC models in order to investigate the formation and arrangement of QDs upon a surface during molecular beam epitaxy experiments. An investigation of surface roughening is performed first in order to ascertain the optimal conditions for these simulations. Finally multiple layers of material and substrate are deposited in order to investigate vertical ordering of QDs into a QD super-lattice.

7.1 Algorithm

The *on-lattice* model is modified to allow the atoms to have a real position in space which is within half the lattice spacing around each site. The transition rates are then dependant upon the energy of each atom, which is dependant on the overall structure and net strain. The jump rate for each atom is calculated using:

$$R_k = v_0 \exp\left(-\frac{\varepsilon \sum_n^N E_n}{k_b T}\right), \quad (7.1)$$

where the sum is over all neighbouring atoms and E_n is the bond energy between the atom and its n^{th} neighbour. For the simulations presented here the argon Lennard Jones potential from section 1 is used. Therefore $\varepsilon/k_b = 120$ and equation (7.1) becomes:

$$R_k = v_0 \exp\left(-\frac{120 \sum_n^N E_n}{T}\right). \text{ Without strain } E = 1 \text{ and this equation reduces to equation (5.1)}$$

for the zero strain case.

The effect of the strain is enhanced by altering the energy used in the jump rate calculation. The energy is modified according to: $E = E_{\min} + \alpha \Delta E$, where $E_{\min}$ is the bond energy without any elastic strain, and ΔE is the modification of the bond energy due to the strain. The parameter α controls the artificial strain enhancement. A value of $\alpha = 1$ represents no artificial strain enhancement. Figure 7.1 below illustrates where each parameter comes from. This is done separately for each atomic bond. Artificially choosing enhancement values of $\alpha > 1$ is reasonable as ab initio simulations have shown that small strains can have a very strong effect on the activation barrier for diffusion. For example, 1-2% strain within a silicon substrate can lead to an order of magnitude increase in surface diffusion [7].




Fig. 7.1 Schematic showing the origin of parameters in the strain enhancement calculation.

All the KMC simulations with strain used an enhancement of $\alpha = 2$. This increases the mobility of strained atoms and enhances the effect of strain and will be shown to lead to the formation of QDs.

Calculation of minimum energy of atoms in simulations

Since the jump rate of an atom now depends on the exact location of the atom and its neighbours, these atoms must be at their equilibrium positions (zero net inter-atomic forces on all particles) before the calculation of the rates. When an atom moves to a new site, it will invariably be in the un-relaxed position. Therefore it and the surrounding atoms must be allowed to move slightly (relax) to find their minimum energy positions.

The simplest way to do this is to use the MD algorithm already presented. However, the atoms should be heavily damped so that they do not oscillate significantly around the minimum but go straight to it. This can be done by neglecting the inertia of the atoms and simply updating positions after each MD time-step. However, this can lead to a greater computational time required for convergence compared to retaining inertia and heavily damping the atomic motion. Alternatively, the conjugate gradient method can be used to find the minimum energy positions of a collection of atoms. However, in the small systems studied here the MD method used was sufficiently fast since all atoms in the simulation were very close to their equilibrium positions.

Using MD to minimise the whole system after each atom jump is very computationally demanding. Therefore after each atom jump a *local* MD minimisation occurs. Here, a boundary of atoms that are 4 times the lattice spacing away from the jump site is fixed in space. Now all the local atoms within this boundary are allowed to find their new minimum energy positions.

To ensure that the system as a whole remains close to its minimum energy, a *global* MD minimisation occurs every 5000 atom jumps. Here all atoms in the simulation (apart from fixed boundaries at the edge of the simulation region) are allowed to move during the minimisation.

The model is still an *on-lattice* model, but the lattice is allowed to deform slightly, i.e. the lattice sites undergo displacements from their original positions. All the lattice sites remain in the same position relative to each other, i.e. the neighbour lists are unchanged.

7.2 Surface roughening

The smooth surface of metals appears very flat to the naked eye, but at the atomic scale it can be quite rough due to the accumulated effect of thermal fluctuations. Above a critical temperature (the roughening transition temperature) a surface will roughen naturally. Elastic strain increases the mobility of atoms on a surface and increases surface roughening. Thermal and strain induced roughening have been investigated on a variety of surfaces, including Gallium Arsenide [95] and helium crystals [99]. An example of a thermal roughening is shown in figure 7.2.




Fig. 7.2 Example pictures by Leamy *et al.* [100] of a surface above and below the roughening transition temperature. Here a simple cubic lattice is used where each atom is also represented by a cube. At low temperatures, there are very few defects such as adatoms, surface vacancies and steps. As the temperature increases the surface becomes far more roughened. The temperature is expressed as a function of the bond energy J . The roughening transition occurs at $T_R=0.632J$.

The roughness of a surface is characterized by the RMS roughness, w , which is simply a measure of the absolute deviation of each atom from the average surface height and is given by:

$$w = \left[\frac{1}{L} \sum_{i=0}^L (h_i - \langle h \rangle_L)^2 \right]^{1/2}, \quad (7.2)$$

where, L , is the length of the surface and $\langle h \rangle_L = \frac{1}{L} \sum_{i=0}^L h_i$ is average surface height. Typically w increases over time.

Due to the competition between surface energy and strain energy, the roughening of a strained surface favours a particular wavelength [101][102], λ given by:

$$\lambda = \frac{2\pi\gamma_0}{3\omega_0}, \quad (7.3)$$

where $\omega_0 = \frac{1}{2} E \varepsilon_0^2$ is the strain energy density in the initially flat film and γ_0 is the surface energy.

The dominant frequency of the roughened surface is measured from the simulations by counting the peaks. The frequency is related to the wavelength by, $l = f\lambda$ since the surface width l is fixed. Equation 7.3 can be used to predict the ratio of wavelengths of the roughened surfaces to the ratio of strain in the substrate between two simulations at different compressive strain. These are related by:

$$\frac{f_1}{f_2} = \frac{\varepsilon_1^2}{\varepsilon_2^2} \quad (7.3)$$

where subscript 1 and 2 refer to parameters from the two different simulations.

Simulations of surface roughening

In this section standard KMC simulations of a surface under various compressive strains at a range of temperatures were performed. Compressive strains only are considered here because the vast majority of heteroepitaxial systems are in compression (systems are unstable under tension leading to formation of cracks). Each simulation consists of 7250 atoms (250 x 29) and has mirror boundary conditions at the edges. The effect of the elastic strain is enhanced in all of the KMC simulations presented here ($\alpha = 2$). All simulations start with the jagged (10) surface shown in figure 7.3. This surface is an unstable one and was found to roughen at any temperature given enough time. This is expected as the surface is not a stable crystal facet. The flat (01) surface which is a stable crystal facet was found to have a roughening transition temperature of about 15K. Roughening on this surface has been thoroughly investigated by Tong Wang [102].




Fig. 7.3 Diagram of the initial configuration of the surface roughening simulations on the (10) surface. The system consists of 7250 atoms (250x29).

In the following simulations, the atoms are coloured according to their local strain. This is simply an average of the strain in each of the atomic bonds. The compressive strain is applied in the x direction only. The atoms at the left and right boundaries are fixed in the x direction but can move freely in the y direction. The atoms at the base are fixed in the y direction and are free to move in the x direction. The corner atoms are fixed in both x and y to ensure the system does not drift as a whole. Results for a range of compressive strains at high and low temperature are presented in figure 7.4 below. Plots of the RMS roughness vs. time at various constant temperatures are shown in figure 7.5 and plots at constant strain are shown in figure 7.6.




Fig. 7.4 Final state of surfaces subjected to compressive strain at high and low temperature. Snapshots are drawn after (a) 4.5 million discrete atom jumps (time= 1.5×10^{56}), (b) 10 million discrete atom jumps (time= 7×10^{14}), (c) 1 million atom jumps (time= 1.7×10^{55}), (d) 38.5 million discrete atom jumps (equivalent time 2.18×10^{15}), (e) 2 million atom jumps (time= 1.9×10^{55}) and (f) 4 million atom jumps (time= 17.99×10^{13}). Each column is at the constant temperature and each row is at the constant strain indicated in the figure.


Figures 7.4a and 7.4b shows a snapshot of the KMC simulation of an unstrained surface. The surface profile is highly dynamic, i.e. there is no persistent desired wavelength of the roughening. The plots of the RMS roughness at various temperatures in figure 7.6a show that the unstrained surface quickly roughens to a value between 1 and 2 at all temperatures.

From KMC simulations with compressive strain at very low temperatures (3K), the crystal facets were found to be very stable (figure 7.4c and 7.4e). The system is below the roughening transition temperature of the faceted surfaces and therefore once they form, they remain stable. The RMS roughness plots are shown in figure 7.5d. Greater compressive strain


results in faster initial roughening, however the highly strained surfaces reach a maximum roughness caused by the stability of the crystal facets. Each side of the peaks is a stable crystal facet; in order for the roughness to increase these sides would no longer be stable facets. This means a higher roughness here is equivalent to a much greater surface energy; hence the system remains trapped in this local energy minimum.

At the higher temperature (15K), the atoms are far more mobile. Furthermore a particular frequency of surface roughness is favoured as expected from the competition between surface energy and strain energy. Figures 7.4d and 7.4f shows that at 15K the ratio of frequencies of the roughened surface (2%:3%) is: $f_2 / f_3 = 0.4$. Equation 7.3 predicts this ratio to be: $f_2 / f_3 = 0.44$ which is a fairly good agreement considering the low frequencies involved. A more accurate figure could be obtained by simulating a surface at least 10 times longer. This is not done here as the computational time would be too long. The RMS roughness at 15K (figure 7.5b) also reaches a maximum once stable facets occur on each side of the peaks. This maximum roughness arises simply due to the geometry of the facets and found to be related to the wavelength ($w \propto \lambda/12$) and hence strain.


At 30K the elastic strain appears to increase the mobility of the atoms on the surface, but no stable surface profile ever evolves. This is because the thermal fluctuations dominate the evolution over the elastic strain driving force.




(a)



(b)




(c)




(d)

Fig. 7.5 RMS roughness of the (10) surface under various compressive strains at a temperature of (a) 30K, (b) 15K, (c) 10K and (d) 3K. The compressive strain vastly increases the roughness of the surface compared to the unstrained cases.


Generally increasing the compressive strain increases the mobility of the atoms on the surface, since the atoms have a higher energy which results in faster jump rates. The results in figures 7.4, 7.5 and 7.6 show that there is a narrow temperature window around 15K at which the strain energy dominates. At high temperatures (figure 7.5a) thermal noise dominates, preventing any large structures from forming. At low temperatures (figures 7.6c and 7.4e), the crystal facets are stable and further roughening is restricted. Therefore the quantum dot simulations in the next section are performed at 15K.



(a)



(b)



(c)

Fig 7.6 Plots of the RMS roughness of the (10) surface vs. time for various temperatures at (a) zero, (b) 2%, and (c) 3% compressive strain.

7.3 KMC simulations of Molecular Beam Epitaxy (MBE)

In this subsection the results of the KMC simulations of molecular beam epitaxy are presented. The two material types introduced in chapter 6 now have different lattice spacing, as well as different rate constants. In all simulations, there is a compressive mismatch of -3% between the lattice spacing of the two material types; this causes an interfacial strain between the two materials.

Initially the substrate material (type 1) is unstrained. The type 2 material is deposited upon the surface by picking at random a column which contains a valid surface site. The atom is created in the first empty site above the surface. As the deposition material (type 2) is rained down upon the substrate, QDs spontaneously form. Figure 7.7 shows a typical QD formed in the KMC simulations. Here, the atoms at the peak of the quantum dot have relaxed back to their equilibrium lattice spacing, whereas atoms that are near the substrate are highly strained. Underneath the QD, the substrate is also under tension due to the lattice mismatch.




Fig. 7.7 The picture shows the result of a KMC simulation of a QD on a surface. The colours represent the average strain of each atom with the dark blues representing local strain greater than 1%. The QD shows significant strain relaxation at its peak as well as tension in the substrate underneath the dot.

Effect of the deposition rate

The distribution of QDs on a surface is influenced by the deposition rate of material being rained down. A low deposition rate allows each atom a significant amount of time to find its preferred minimal energy site before the next atom arrives on the surface. However a high deposition rate does not give each atom enough time to find its minimum energy site before it gets buried by further deposit.

Figure 7.8 compares the final state of a low and high deposition rate simulation of MBE. The QDs formed are of a similar size in both cases, yet in the low deposition rate simulation they are separated from each other on the surface.

The system conforms to the Stranski-Krastanow growth mode. A uniform wetting layer of one monolayer forms first. Then the subsequently deposited material roughens to form QDs on the surface, which grow to a preferred size which depends on the lattice mismatch.




Fig. 7.8 Comparison low and high deposition rates. (a) Final state where one atom is deposited every 1000 discrete atom jumps. A total of 1000 atoms are rained down on to a fixed flat surface at 10K with $K_{11}=26$, $K_{12}=19$ and $K_{22}=12$. (b) One atom is deposited every 250 discrete atom jumps at 15K with $K_{11}=9.6$, $K_{12}=8.8$ and $K_{22}=8.0$. A total of 1700 atoms are deposited and they are allowed to mix with the substrate. Plots are coloured to show the local strain on each atom.

Capping – intermixing

Here more type 1 atoms are rained down after the formation of QDs in figure 7.8b from type 2 atoms. The rate constants and temperature remain the same at 15K with $K_{11}=9.6$,

$K_{12}=8.8$ and $K_{22}=8.0$. As the QDS are capped, the peaks of the QDs disappear as this material mixes with the capping layer to reduce the interfacial strain between the two materials. Experimental results (figure 7.11) show that this is expected to occur.

The capping layer is under tensile elastic strain above each of the embedded QDs as can be seen in figure 7.9b. Referring to figure 7.8b, the QDs also stretch the substrate underneath them. Therefore these sites above each QD are the low energy sites for the next layer of QDs to form since the substrate is already stretched here. The dark blue colour of the buried QDs in figure 7.9b indicates that they are now highly compressed by the material all around them.




Fig. 7.9 A capping layer of more type 1 atoms is rained down upon the quantum dots shown in figure 7.8. Snapshots of, (a) the composition and (b) the strain field, are shown.

Vertical alignment

Another layer of type 2 atoms are deposited on top of the capping layer. Figure 7.10 shows that the next layer of QDs form on the low energy sites of the capping layer as expected. This results in a vertical alignment which agrees well with experimental results (figure 7.11).




Fig 7.10 Another layer of the type 2 material is deposited upon the surface shown in figure 7.9. The QDs form at the sites where the substrate is already under tension due to the buried QDs, and hence a vertical alignment of QDs occurs.




Fig. 7.11 An example of an experimental result of the deposition of alternating layers of two materials by MBE. After the deposition, the sample is cut open to reveal vertically aligned quantum dots. Since the QDs are under compression within, they relax slightly by protruding from the surface. Also, clearly visible is the intermixing of the QD material with each capping layer [103].

The continuum KMC method is also capable of simulating this process. Figure 7.12 shows images of a continuum KMC simulation taken at similar points to figures 7.7 to 7.10.

Quantum dot top removal, intermixing in the capping layer and vertical alignment are all demonstrated by the continuum method.




Fig. 7.12 Continuum KMC model of the evolution of an elastically strained two-component system by surface diffusion: the growth, relaxation and capping of alloyed heteroepitaxial QDs [104]. The top picture shows the formation of pyramidal QDs after a period of deposition of (blue) dot atoms. Some intermixing (green) at the substrate-dot interface is observed. The dots are then capped (see middle image) by further deposition of the substrate material (red). As seen experimentally (figure 7.11), the top of the pyramids dissolve into the capping layer and the embedded quantum dots are reduced to truncated pyramids. Another layer of QDs is deposited along with another capping layer. This leads to the development of a vertically-aligned structure, as observed in experiments.

Further work - Extension to multiscale

The simulations presented here require many days of computer time to run. Minimising the system to calculate the jump rates consumes most of the computer time. Ideally greater length and time scales are required to fully investigate the properties of these strained surfaces. A multiscale model would require standard KMC to be applied to the surface atoms only, and a continuum finite element method can be used to compute the strain field in the substrate.

Although it was shown in section 6 that the continuum KMC model did not exactly represent the quantitative evolution of the atomistic KMC model, it was found that there was good qualitative agreement in most cases. Comparison of Figs. 7.10 and 7.12 demonstrate that this is also the case when the additional complexity of elastic strain is included in the two different algorithms. This is encouraging and shows that further development of the continuum KMC model is worthwhile. However, due to lack of absolute quantitative agreement, it is unlikely that these two models could be combined together into a single concurrent multiscale simulation. The continuum KMC is very good at converging quickly at high temperatures which can be used to extend the time scale of these simulations.

Conclusion for section II

In this section the KMC algorithm and the continuum KMC algorithm were introduced. In subsection 6.1, both methods were initially used to simulate the evolution of ellipses by surface diffusion in order to demonstrate the validity of the continuum method. Both methods were found to be in good agreement at high temperatures. At low temperatures, the continuum KMC method becomes trapped in a meta-state. This was resolved by adding thermal noise and biasing to the continuum KMC algorithm. The simulation results for 2 component ellipses for three cases A, B and C, were compared at high and low temperatures. These results were presented in subsection 6.2. The continuum method was found to be in good qualitative agreement in most cases, however, at low temperatures it can become trapped in meta-stable states.

Elastic strain was added to the standard and continuum KMC methods and was introduced in chapter 7. Simulations of surface roughening were performed using the standard KMC algorithm. These simulations show that strain induced roughening is best observed around 15K for this system. At this temperature, the system is hot enough to overcome the stability of the crystal facets, but not too hot such that any structure is lost in the thermal noise. A characteristic wavelength of roughening occurs which depends on the interaction between surface energy and strain energy. The ratio of wavelengths was found to be in agreement with continuum models of surface roughening.

By depositing a material with a lattice spacing 3% higher than the substrate material, an interfacial strain between substrate and deposit occurs. This caused the formation of quantum dots on the surface. A low deposition rate caused a few quantum dots of a characteristic size to form, whereas a high deposition rate caused many QDs of a characteristic size to form. Alternating layers of substrate and deposit material were then rained down upon a flat surface. This caused the formation of vertically aligned quantum dots (QD super-lattice) due to the interaction between embedded QDs and the surface via the strain field. A continuum KMC simulation of MBE also resulted in vertical alignment of QDs. This is a good example of a self organised structure which is very useful in the fabrication of nano-scale devices.

Overall conclusion

This concludes this section and the thesis. Two very different atomistic simulation techniques have been compared with derivative continuum models. The first (dynamic) model considered thermal transport by conduction; the second (kinetic) model was concerned with material transport by diffusion. In both cases it was found that many of the atomistic effects arising from the discrete, stochastic atomistic models could be well-represented by their continuum counterparts. However, equally it was found that some atomistic effects could not be readily represented. For instance, the Kapitza resistance due to phonon scattering at interfaces and defects in thermal problems, and the nucleation of surface steps in diffusion problems. This shows that the need for comparing, benchmarking and validating different models at different scales is of great importance. It also shows the need for the development of effective multiscale coupled atomistic/continuum models for the extension of simulations to larger length and time scales in the future. Hopefully, the work presented in this thesis is a useful contribution to progress in this field.

References

- [1] G.E. Moore, Cramming more components onto integrated circuits, *Electronics*, 38(8), (1965), pp114-117.
- [2] E. Bengu, R. Plass, L.D. Marks, T. Ichihashi, P.M. Ajayan, and S. Iijima, Imaging the Dimers in Si(111)-(7×7), *Phys. Rev. Lett.*, 77(20), (1996), pp4226-4228.
- [3] J. Bansmann, S.H. Baker, C. Binns, et al, Magnetic and structural properties of isolated and assembled clusters, *Surface Science Reports*, 56(5-6), (2005), pp189-275.
- [4] K. Brunner, J. Zhu, C. Miesner, G. Abstreiter, O. Kienzle, and F. Ernst, Self-organized periodic arrays of SiGe wires and Ge islands on vicinal Si substrates, *Physica E*, 7, (2000), pp881-886.
- [5] R.E. Rudd, G.A.D. Briggs, A.P. Sutton, G. Medeiros-Ribeiro, and R.S. Williams, Equilibrium Distributions and the Nanostructure Diagram for Epitaxial Quantum Dots, *J. Comput. Theor. Nanosci*, 4(2), (2007), pp335-347.
- [6] S. Bhowmick, and V.B. Shenoy, Effect of strain on the thermal conductivity of solids, *J. Chem. Phys.* 125, (2006),164513.
- [7] D.J. Shu, F. Liu, X.G. Gong, Simple generic method for predicting the effect of strain on surface diffusion, *Phys. Rev. B*, 64, (2001), 245410.
- [8] B.J. Ohlsson, et al, Growth and characterization of GaAs and InAs nano-whiskers and InAs/GaAs heterostructures, *Physica E*, 13, (2002), pp1126 – 1130.
- [9] J.H. Lee, and J.C. Grossman, Lattice thermal conductivity of nanoporous Si: Molecular dynamics study, *Appl. Phys. Lett.* 91, (2007), 223110.
- [10] C. Santori, D. Fattal, J. Vuckovic, G.S. Solomon and Y. Yamamoto, Single-photon generation with InAs quantum dots, *New Journal of Physics*, 6, 89, (2004).
- [11] <http://cqdeecs.northwestern.edu/research/qdots.php>, Centre for Quantum Dot Devices, Northwestern University, USA.
- [12] Y. Lee, J.Y. Park, S.Y. Kim, S. Jun, and S. Im, Atomistic simulation of incipient plasticity under Al(111) nanoindentation, *Mech. Mat.*, 37, (2005), pp1035-1048.
- [13] K. Capelle, A bird's-eye view of density-functional theory, *Braz. J. Phys.*, 36(4A), (2006), pp1318-1343.
- [14] P. Raybaud, et al, First principles surface thermodynamics of industrial supported catalysts in working conditions, *J. Phys. Condens. Matter*, 20, (2008) 064235.

- [15] W.K. Liu, E.G. Karpov, S. Zhang, and H.S. Park, An introduction to computational nanomechanics and materials, *Comput. Methods Appl. Mech. Engrg.*, 193, (2004), pp1529-1578.
- [16] www.comsol.com.
- [17] J.Q. Broughton, F.F. Abrahams, N. Bernstein, and E. Kaxiras, Concurrent coupling of length scales: Methodology and application, *Phys. Rev. B* 60, (1999), 2391.
- [18] S.P. Xiao, and T. Belytschko, A bridging domain method for coupling continua with molecular dynamics, *Comput. Methods Appl. Mech. Engrg.*, 193, (2004), pp1645-1669.
- [19] S. Qu, V. Shastry, W.A. Curtin, and R.E. Miller, A finite temperature dynamic coupled atomistic/discrete dislocation method, *Modelling Simul. Mater. Sci.* 13 (2005), 1101.
- [20] S.P.A. Gill, Z. Jia, B. Leimkuhler and A.C.F. Cocks, Rapid thermal equilibration in coarse-grained molecular dynamics, *Phys. Rev. B.* 73, (2006), 184304.
- [21] W.A. Curtin and R.E. Miller, Atomistic/continuum coupling in computational materials science, *Modelling Simul. Mater. Sci. Eng.*, 11, (2003), R33-R68.
- [22] X. Li and W. E, Variational Boundary Conditions for Molecular Dynamics Simulations of Solids at Low Temperature, *Comm. Comp. Phys.*, 1, (2006), pp135-175.
- [23] X. Li and W. E, Multiscale modelling of the dynamics of solids at finite temperature, *J. Mech. Phys. Solids*, 53, (2005), pp1650-1685.
- [24] W.K. Liu, H.S. Park, D. Qian, E.G. Karpov, H.Kadowaki, G.J. Wagner, Bridging scale methods for nanomechanics and materials, *Comput. Methods Appl. Mech. Engrg.*, 195, (2006), pp1407-1421.
- [25] H.S. Park, W.K. Liu, An introduction and tutorial on multiple-scale analysis in solids, *Comput. Methods Appl. Mech. Engrg.*, 193, (2004), pp1733-1772.
- [26] H.S. Park, E.G. Karpov, W.K. Liu, A temperature equation for coupled atomistic/continuum simulations, *Comput. Methods Appl. Mech. Engrg.*, 193, (2004), pp1713-1732.
- [27] R. E. Rudd and J.Q. Broughton, Concurrent coupling of length scales in solid state systems, *Phys. Stat. Sol. B*, 217(1), (2000), pp251-291.

- [28] L.E. Shilkrot, R.E. Miller and W.A. Curtin, Multiscale plasticity modelling: coupled atomistic and discrete dislocation mechanics, *J. of the Mech. Phys. Of Solids*, 52, (2004), pp755-787.
- [29] W.K. Liu, E.G. Karpov, and H.S. Park, personal communication.
- [30] In-Ho Lee, S.Y. Kim, and S. Jun, An introductory overview of action-derived molecular dynamics for multiple time-scale simulations, *Comput. Methods. Mech. Engrg.*, 193, (2004), pp1633-1644.
- [31] S.Y. Kim, In-Ho Lee, S. Jun, Y. Lee, and S. Im, Coalescence and T-junction formation of carbon nanotubes: Action-derived molecular dynamics simulations, *Phys. Rev. B*, 74, (2006), 195409.
- [32] R.A. Miron and K.A. Fichthorn, Accelerated molecular dynamics with the bond-boost method, *J. Chem. Phys.*, 119(12), (2003), pp6210-6216.
- [33] A. Chatterjee, D.G. Vlachos, Temporal acceleration of spatially distributed kinetic Monte Carlo simulations, *J. Comput. Phys.*, 211, (2006), pp596-615.
- [34] M.A. Katsoulakis, and D.G. Vlachos, Coarse-grained stochastic processes and kinetic Monte Carlo simulators for the diffusion of interacting particles, *J. Chem. Phys.*, 119(18), (2003), pp9412-9427.
- [35] A. Chatterjee, M.A. Katsoulakis and D.G. Vlachos, Spatially adaptive grand canonical ensemble Monte Carlo simulations, *Phys. Rev. E*, 71, (2005), 026702.
- [36] A. Chatterjee, D.G. Vlachos and M.A. Katsoulakis, Spatially adaptive lattice coarse-grained Monte Carlo simulations for diffusion of interacting molecules, *J. Chem. Phys.*, 121(22), (2004), pp11420-11431.
- [37] C. Chou, and M.L. Falk, Multiscale diffusion Monte Carlo simulation of epitaxial growth, *J. Comput. Phys.*, 217, (2006), pp519-529.
- [38] J.P. DeVita, L.M. Sander, and P. Smereka, Multiscale Kinetic Monte-Carlo for Simulating Epitaxial Growth, *arXiv:cond-mat/0504272v2*, cond-mat.mtrl-sci, (2005).
- [39] D.R. Mason, R.E. Rudd and A.P. Sutton, Stochastic kinetic Monte Carlo algorithms for long-range Hamiltonians, *Computer Physics Communications*, 160, (2004), pp140-157.
- [40] L. Nurminen, A. Kuronen, and K. Kaski, Kinetic Monte Carlo simulation on patterned substrates, *Phys. Rev. B*, 63, (2000), 035407.

- [41] L.M. Dupuy, E.B. Tadmor, R.E. Miller and R. Phillips, Finite-temperature quasicontinuum: Molecular dynamics without all the atoms, *Phys. Rev. Lett.*, 95, (2005), 060202.
- [42] M.P. Allen, D.J. Tildesley, *Computer simulations of Liquids*, Oxford science publications, (1990).
- [43] S. Lepri, R. Livi, A. Politi, Thermal conduction in classical low-dimensional lattices, *Physics reports* 377, (2003), pp1-80.
- [44] M.P. Allen, *Introduction to Molecular Dynamics Simulation*, Computational Soft Matter: From Synthetic Polymers to Proteins, Lecture Notes, John von Neumann Institute for Computing, Julich, NIC Series, 23, (2004), pp. 1-28.
- [45] R. Komanduri, L.M. Raff, A review on the molecular dynamics simulation of machining at the atomic scale, *Proc.Instn. Mech. Engrs.*, 215 Part B.
- [46] K. Binder, J. Horbach, W. Kob, W. Paul and F. Varnik, Molecular dynamics simulations, *J. Phys.: Condens. Matter* 16, (2004), S429-S453
- [47] Y. Wang, D. Tomanek, G.F. Bertsch, Stiffness of a solid composed of c60 clusters, *Phys. Rev. B*, 44(12), (1991), pp6562–6565.
- [48] R. Mahaffy, R. Bhatia, B.J. Garrison, Diffusion of a Butanethiolate molecule on a Au{111} surface, *J. Phys. Chem. B*, 101, (1997), pp771–773.
- [49] J. Tersoff, Modeling solid-state chemistry: Interatomic potentials for multicomponent systems, *Phys. Rev. B*, 39, (1989), pp5566-5568.
- [50] J. Tersoff, New empirical approach for the structure and energy of covalent systems, *Phys. Rev. B*, 37(12), (1988), pp6991–7000.
- [51] J. Tersoff, New empirical model for the structural properties of silicon, *Phys. Rev. Lett.*, 56 (6), (1986), pp632–635.
- [52] J. Tersoff, Empirical interatomic potential for carbon, with applications to amorphous carbon, *Phys. Rev. Lett.*, 61(25), (1988), pp2879–2882.
- [53] M.S. Daw, S.M. Foiles, M.I. Baskes, The embedded-atom method: a review of theory and applications, *Materials Science Reports*, 9, (1993), pp251-310.
- [54] L. Verlet, Computer “Experiments” on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules, *Phys. Rev.* 159(1), (1967), pp98-103.
- [55] D. C. Rapaport, *The art of Molecular Dynamics Simulation*, Cambridge University Press, (2004).

- [56] Z. Huang, Z. Tang, Evaluation of momentum conservation influence in non-equilibrium molecular dynamics methods to compute thermal conductivity, *Physica B*, 373, (2006), pp291-296.
- [57] D. Quigley, M.I.J. Probert, Langevin dynamics in constant pressure extended systems, *J. Chem. Phys.*, 120(24), (2004).
- [58] Wm.G. Hoover, K. Aoki, C.G. Hoover, S.V.De Groot, Time-reversible deterministic thermostats, *Physica D*, 187, (2004), pp253-267.
- [59] A.C. Brańka, Nosé-Hoover chain method for nonequilibrium molecular dynamics simulation, *Physical Review E*, 61(5), (2000).
- [60] S. Nosé, A molecular dynamics method for simulations in the canonical ensemble, *Molecular Physics*, 52(2), (1984), pp255-268.
- [61] H.J.C. Berendsen et al, Molecular dynamics with coupling to an external bath, *J. Chem. Phys.* 81(8), (1984).
- [62] J.R. Hook, H.E. Hall, *Solid state Physics second edition* (J. Wiley & Sons Ltd. 1991).
- [63] www.wikipedia.org, submitted under GNU free documentation license.
- [64] P. Heino, Thermal conduction simulations in the nanoscale. *J Comput Theor Nanosci*, 4(5), (2007), pp896-927.
- [65] J.Y. Murthy, S.V.J. Narumanchi, J.A. Pascual-Gutierrez, T. Wang, C. Ni, S.R. Mathur, Review of multiscale simulation in submicron heat transfer. *Int J Multiscale Comput Engrg.*, 3(1), (2005), pp5-32.
- [66] C. Kittel, *Introduction to Solid State Physics* (John Wiley and Sons, Inc, 8th edition), (2005).
- [67] S. Sinha, K.E. Goodson, Thermal conduction in sub-100nm transistors, *Microelectronics Journal*, 37(11), (2006), pp1148-1157.
- [68] P. Heino, Thermal conductivity and temperature in solid argon by nonequilibrium molecular dynamics simulations, *Phys. Rev. B*, 71, (2005), 144302.
- [69] S. Lepri, R. Livi and A. Politi, Energy transport in anharmonic lattices close to and far from equilibrium, *Physica D*, 119, (1998), pp140-147.
- [70] R. Prasher, Diffraction limited phonon thermal conductance of nanoconstrictions, *Appl. Phys. Lett.*, 91, (2007), 143119.
- [71] P.K. Schelling, S.R. Phillpot and P. Keblinski, Comparison of atomic-level simulation methods for computing thermal conductivity, *Phys. Rev. B*, 65(14), (2002) 144306.

- [72] D. Segal, A. Nitzan and P. Hänggi, Thermal conductance through molecular wires, *J. Chem. Phys.* 119(13), (2003), 6840.
- [73] T. Watanabe, B. Ni and S.R. Phillpot, Thermal conductance across grain boundaries in diamond from molecular dynamics simulation, *J. Appl. Phys.*, 102, (2007), 063503.
- [74] T. Terao and F. Muller-Plathe, A nonequilibrium molecular dynamics method for thermal conductivities based on thermal noise, *J. Chem. Phys.*, 122, (2005), 081103.
- [75] A. Rahman, Correlations in the Motion of Atoms in Liquid Argon, *Phys. Rev.*, 136, 2A, (1964).
- [76] B.L. Holian and R. Ravelo, Fracture simulations using large-scale molecular-dynamics, *Phys. Rev. B.*, 51, (1995), pp11275-11288.
- [77] M.E. Lumpkin, W.M. Saslow, One-dimensional Kapitza conductance: Comparison of the phonon mismatch theory with computer experiments, *Physical Review B*, 17(11), (1978), pp4295-4302.
- [78] Rob Phillips, *Crystal, defects and microstructures – Modelling across scales*, Cambridge University Press, (2001).
- [79] C. Giardiná, R Livi, A. Politi, and M. Vassalli, Finite Thermal conductivity in 1D lattices, *Physical Review Letters*, 84(10), (2000), pp2144 – 2147.
- [80] Bambi Hu et al, Heat conduction in one-dimensional nonintegrable systems, *Physical Review E*, 61(4), (2000), pp3828 – 3831.
- [81] A. V. Savin and O. V. Gendelman, Heat conduction in one-dimensional lattices with on-site potential, *Physical Review E*, 67, (2003), 041205.
- [82] B. Li, H. Zhao and B. Hu, Can disorder induce a finite thermal conductivity in 1D lattices?, *Physical Review Letters*, 86(1), (2001).
- [83] A. Lippi, R. Livi, Heat conduction in 2D nonlinear lattices, *arXiv:chaodyn/9910034 v1*, (1999).
- [84] K. Jolley, S.P.A. Gill, Atomistic/continuum models for multiscale heat conduction, *Proceedings of the Fourth International Conference of Multiscale Materials Modelling*, (2008), P214.
- [85] K. Jolley, S.P.A. Gill, Modelling Transient Heat Conduction at Multiple Length and Time Scales: A coupled Non-Equilibrium Molecular Dynamics/Continuum Approach, To appear in *J. Comp. Phys.*

- [86] W. Ren, Analytical and numerical study of coupled atomistic-continuum methods for fluids, *J. Comp. Phys.*, 227, (2007), pp1353-1371.
- [87] J. Liu, S. Chen, X. Nie and M. O. Robbins, A continuum–atomistic simulation of heat transfer in micro- and nano-flows, *J. Comp. Phys.*, 227, (2007), pp279-291.
- [88] G. Chen, D. Borca-Tasciuc and R.G. Yang, Nanoscale heat transfer, *Encyclopedia of Nanoscience and Nanotechnology* (American Scientific publishers, 2004).
- [89] W. Huang, G. Huang, L. Wang and B. Huang, Phonon-cavity-enhanced low temperature thermal conductance of a semiconductor nanowire with narrow constrictions, *Phys. Rev. B*, 75, (2007), 233415.
- [90] W. Tian and R. Yang, Effect of interface scattering on phonon thermal conductivity percolation in random nanowire composites, *Appl. Phys. Lett.*, 90, (2007), 263105.
- [91] J. Wang, Quantum Thermal Transport from Classical Molecular Dynamics, *Phys. Rev. Lett.*, 99, (2007), 160601.
- [92] F. Arciprete, E. Placidi, V. Sessi, M. Fanfoni, F. Patella, and A. Balzarotti, How kinetics drives the two- to three-dimensional transition in semiconductor strained heterostructures: The case of InAs/GaAs(001), *Appl. Phys. Lett.*, 89, (2006), 041904.
- [93] C.M. Reaves, Effects of Deposition Rate on the Size of Self-Assembled InP Islands Formed on GaInP/GaAs(100) Surfaces, *Journal of Electronic Materials*, 24(11), 1995.
- [94] K. Brunner, Si/Ge nanostructures, *Rep. Prog. Phys*, 65, (2002) pp27–72.
- [95] T. Tiedje, A. Ballestad, Atomistic basis for continuum growth equation: Description of morphological evolution of GaAs during molecular beam epitaxy, *Thin Solid Films*, 516, (2008), pp3705–3728.
- [96] C.C. Battaile, D.J. Srolovitz, Kinetic Monte Carlo Simulation of Chemical Vapor Deposition, *Annu. Rev. Mater. Res.*, 32, (2002). pp297–319.
- [97] M. Matsumoto and T. Nishimura, Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator, *ACM Transactions on Modeling and Computer Simulation*, 8(1), (1998), pp 3-30.
- [98] S.P.A. Gill and P.E. Spencer, Concurrent Multiscale Kinetic Monte Carlo-continuum Models for the Evolution of Solids via Diffusion, *Solid State Phenomena*, 129, (2007), pp19-24.
- [99] S. Balibar, H. Alles, A.Ya. Parshin, The surface of helium crystals, *Reviews of modern physics*, 77, (2005).

- [100] H.J. Leamy, G.H. Gilmer, and K.A. Jackson, Statistical Thermodynamics of Clean Surfaces, Surface Physics of Materials, Ed: J. M. Blakely, Academic Press, New York, (1975), p121.
- [101] D.J. Srolovitz, On the stability of surface of stressed solids, Acta Metall., 37, (1989), pp621-625.
- [102] S.P.A. Gill, and T. Wang, On the existence of a critical perturbation amplitude for the Stranski-Krastanov transition, Surface Science, 602(23), (2008), pp3560-3568.
- [103] D.M. Bruls, P.M. Koenraad, H.W.M. Salemink, et al, Stacked low-growth-rate InAs quantum dots studied at the atomic level by cross-sectional scanning tunneling microscopy, Applied physics Letters, 82(21), (2003), pp3758-3760.
- [104] S.P.A. Gill, A pseudo-atomistic model for nanoscale strain-controlled diffusion problems: the formation of heteroepitaxial alloyed quantum dot superlattices, (invited article) Journal of Computational and Theoretical Nanoscience, 5, (2008), pp80-87.

Appendices

A Derivation of heat flux in 1D

This derivation is similar to the one presented in [15]. The energy per particle is defined as:

$$h = \frac{\dot{x}_n}{2} + \frac{1}{2}[V(x_{n+1} - x_n) + V(x_n - x_{n-1})] \quad (\text{A.1})$$

The first term is the kinetic energy of the particle and the second term is half of the total pairwise interactions. Now consider the continuity equation:

$$\dot{h}(x, t) + \text{div} j(x, t) = 0 \quad (\text{A.2})$$

In 1D this is equivalent to:

$$\frac{\partial h(x, t)}{\partial t} + \frac{\partial j(x, t)}{\partial x} = 0 \quad (\text{A.3})$$

Hence we find:

$$\frac{\partial h(x, t)}{\partial t} = \dot{x}\ddot{x} + \frac{\partial V(x_{n+1} - x_n)}{\partial(x_{n+1} - x_n)} \frac{\partial(x_{n+1} - x_n)}{\partial t} + \frac{\partial V(x_n - x_{n-1})}{\partial(x_n - x_{n-1})} \frac{\partial(x_n - x_{n-1})}{\partial t} \quad (\text{A.4})$$

$$\frac{\partial h(x, t)}{\partial t} = \dot{x}\ddot{x} + \frac{1}{2}[-(\dot{x}_{n+1} - \dot{x}_n)F(x_{n+1} - x_n) - (\dot{x}_n - \dot{x}_{n-1})F(x_n - x_{n-1})] \quad (\text{A.5})$$

where $F = -\frac{dV}{dx}$.

From the fact that the acceleration of the particle depends on the net force we have:

$$\ddot{x} = [-F(x_{n+1} - x_n) + F(x_n - x_{n-1})] \quad (\text{A.6})$$

So, on substitution into (A.5) one obtains:

$$\frac{\partial h(x, t)}{\partial t} = -\frac{1}{2}[(\dot{x}_{n+1} + \dot{x}_n)F(x_{n+1} - x_n) - (\dot{x}_n + \dot{x}_{n-1})F(x_n - x_{n-1})] \quad (\text{A.7})$$

Now, j is defined as:

$$\frac{\partial j(x,t)}{\partial x} \approx \frac{\Delta j_n}{\Delta x} = \frac{j_n - j_{n-1}}{a}, \quad (\text{A.8})$$

and from continuity equation:

$$j_n - j_{n-1} = a \frac{\partial h(x,t)}{\partial t}, \quad (\text{A.9})$$

where j_n is:

$$j_n = \frac{1}{2} a (\dot{x}_{n+1} + \dot{x}_n) F(x_{n+1} - x_n) \quad (\text{A.10})$$

This can be interpreted as a power (force between atoms $\times$ average velocity between atoms, from $P=FV$). The authors of [15], also present a further method involving Fourier integrals to obtain expressions where density fluctuation cannot be neglected.

B Hamiltonian formulation

The derivation of the Hamiltonian and the Hamiltonian equations of motion are presented here. It should be noted that this formulation does not depend on the coordinate system used. Starting with the Euler-Lagrange equation:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_\alpha} \right) - \frac{\partial L}{\partial q_\alpha} = 0 \quad (\text{B.1})$$

where, L is the Lagrangian (Kinetic energy – Potential energy) of the system, and q are the generalised coordinates, and $\alpha = 1, 2, \dots, s$.

The Hamiltonian formulation uses the generalised momentum defined as:

$$p_i = \frac{\partial L}{\partial \dot{x}}, \quad (\text{B.2})$$

which, in Cartesian coordinates leads to the expected result:

$$p_i = \frac{\partial}{\partial \dot{x}_i} \left(\frac{m_i \dot{x}_i^2}{2} \right) = m_i \dot{x}_i. \quad (\text{B.3})$$

From the Euler-Lagrange equation (B.1) $\dot{p}_i$ can also be defined:

$$\dot{p}_i = \frac{\partial L}{\partial q}. \quad (\text{B.4})$$

To obtain the Hamiltonian, start with the total derivative of the Lagrangian:

$$dL(q, \dot{q}) = \frac{\partial L}{\partial q} dq + \frac{\partial L}{\partial \dot{q}} d\dot{q}. \quad (\text{B.5})$$

Upon substituting equations (B.2) and (B.4) in the above (B.5), one obtains:

$$dL(q, \dot{q}) = \dot{p}dq + pd\dot{q}, \quad (\text{B.6})$$

From the product rule, it is known that:

$$d(p\dot{q}) = p d\dot{q} + \dot{q} dp. \quad (\text{B.7})$$

On substitution of (B.7) in (B.6):

$$dL(q, \dot{q}) = \dot{p}dq + d(p\dot{q}) - \dot{q}dp. \quad (\text{B.8})$$

Then rearranging (B.8) gives:

$$d(p\dot{q} - L) = \dot{q}dp - \dot{p}dq, \quad (\text{B.9})$$

$$dH = \dot{q}dp - \dot{p}dq, \quad (\text{B.10})$$

where $H = p\dot{q} - L$.

Therefore the Hamiltonian equations of motion can be found from (B.10):

$$\frac{\partial H}{\partial p} = \dot{q}, \quad (\text{B.11})$$

$$\frac{\partial H}{\partial q} = -\dot{p}. \quad (\text{B.12})$$

From (B.9), the Hamiltonian is:

$$H = \left[\left(\frac{\partial L}{\partial \dot{q}} \dot{q} \right) - L \right]. \quad (\text{B.13})$$

Now given the Lagrangian, L , which is (Kinetic energy – Potential energy), or $L = T - V$, and as V does not depend on the speeds and that T is function of speeds squared, one obtains:

$$\frac{\partial(T - V)}{\partial \dot{q}} \dot{q} = 2T, \quad (\text{B.14})$$

and therefore on substitution of (B.14) in (B.13):

$$H = 2T - (T - V) = T + V. \quad (\text{B.15})$$

That is, the Hamiltonian is the sum of the kinetic and potential energy.

In the case of a system of particles in the Cartesian coordinate system, the Hamiltonian in 1D becomes:

$$H = \sum_{i=1}^N \left[\frac{p_i^2}{2m_i} + V(x_{i+1} - x_i) \right]. \quad (\text{B.16})$$

The first term is the kinetic energy, the second term is an optional on site potential, and the last term is the potential due to interaction with neighbouring particles and finally the sum is over all the atoms in the simulation.

C Elliptical void in a two component material

Presented here are the results of KMC simulations of an elliptical void in a two component material. The starting configuration (shown in figure C.1) is simply the inverse of the 2 component ellipse simulations presented in chapter 6. The bonding energies for each case are also the same as the ones used in chapter 6.




Fig. C.1 The initial configuration of the KMC simulations of a void within a two component material is shown. The type 1 atoms (red) occupy the left half and the type 2 (blue) atoms occupy the right half of the simulation region initially.

Case A low temperature

The bond energies for each case are the same as the equivalent cases in chapter 6 shown in table 6.1. For this case they are: $K_{11} = 4.8$, $K_{12} = 4.0$ and $K_{22} = 4.8$. Each atom type preferentially bonds to itself rather than mixing. The final state (Figure C.2) shows that a small amount of mixing has occurred and the final state of a circular void is obtained.




Fig. C.2 Final state of the case A low temperature KMC simulation. The snapshot is taken after 20 million atom jumps (equivalent to a time of 6.42×10^{10}). A small amount of mixing occurs as the ellipse collapses to a circle. The many straight edges are due to the stability of the crystal facets at this temperature.

Case A High temperature

At the high temperature, the jump rates are: $K_{11} = 2.4$, $K_{12} = 2.0$ and $K_{22} = 2.4$. Figure C.3 shows that vastly more mixing of the atomic species has occurred, and a more rounded void in the centre is obtained.




Fig. C.3 Final state of the case A high temperature KMC simulation. The snapshot is taken after 40 million atom jumps (equivalent to a time of 8.46×10^7). Significantly more mixing occurs compared to the low temperature case.

Case B low temperature

In case B, mixing of atomic species is preferred. A uniformly mixed cluster of atoms grows at the interface between the two materials on both the upper and lower faces. Eventually these clusters meet in the middle causing the void to be split into two. The final state obtained is shown in figure C.4.




Fig. C.4 Final state of the case B low temperature KMC simulation. The snapshot is taken after 20 million atom jumps (equivalent to a time of 1.12×10^{10}). The material builds up in the centre on both sides. Eventually the two mixed areas meet dividing the void in two.

Case B high temperature

The high temperature simulation proceeds in a similar fashion to the low temperature case, resulting in the formation of two voids as before as seen in figure C.5. However, significantly more mixing occurs both in the centre and around the outer edges. Also many single atom voids meander through the simulation region at this temperature.




Fig. C.5 Final state of the case B high temperature KMC simulation. The snapshot is taken after 35 million atom jumps (equivalent to a time of 2.95×10^7). The void is split into two by the mixed material forming in the centre as before. However, significantly more mixing occurs compared to the low temperature case as expected.

Case C low temperature

In this case the type 1 atoms (red) are more mobile than the type 2 atoms (blue). The jump rates in this case are: $K_{11} = 4.0$, $K_{12} = 4.8$ and $K_{22} = 5.6$. The type 1 atoms diffuse across the simulation region to fill the empty half ellipse within the type 2 atoms as seen in figure C.6. The void left behind on the left-hand side then evolves to its minimum energy shape. Crystal facets are relatively stable at this temperature, resulting in a roughly circular faceted void being formed.




Fig. C.6 Final state of the case C low temperature KMC simulation. The snapshot is taken after 30 million atom jumps (equivalent to a time of 1.87×10^{10}). The type 1 (red) atoms diffuse into the half elliptical void in the type 2 (blue) atoms on the right. The void in the left then converges to its minimum energy state, which is a circle.

Case C high temperature

Again the high temperature case proceeds in a similar fashion to the low temperature case. Significantly more mixing between atomic species occurs, both within the filled half of the ellipse on the right, and around the edge of the void on the left-hand side. Furthermore, the initial sharp edges of the type 2 atoms have been eroded by mixing to produce a wider opening of the half ellipse in the centre and a rounded edge at the far right as seen in figure C.7 below.




Fig. C.7 Final state of the case C high temperature KMC simulation. The snapshot is taken after 30 million atom jumps (equivalent to a time of 4.41×10^7). The simulation proceeds in a similar fashion to the low temperature case, yet significantly more inter-mixing occurs along with all edges becoming more rounded.

As with the simulations of ellipses in chapter 6, the final state of the simulation is highly dependant upon the bonding energies of each of the atomic species.