

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/94350>

Copyright and reuse:

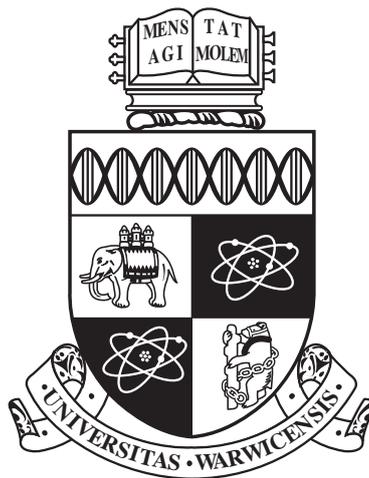
This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



**Abstract surface growth modelling with application to
graphene**

by

Gwilym Enstone

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Centre for Complexity Science

June 2017

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	iv
List of Figures	v
Declarations	vii
Abbreviations	viii
Abstract	x
Chapter 1 Introduction	1
Chapter 2 Background	4
2.1 Modelling surface growth	4
2.1.1 Density functional theory	5
2.1.2 Molecular Dynamics	7
2.1.3 Monte Carlo methods	8
2.1.4 Phase-field modelling	10
2.1.5 Rate equations and scaling	10
2.2 Experimental comparison: chemical vapour deposition of graphene on copper	11
2.2.1 Challenges involved in observing a growing system	13
2.2.2 Growth on liquid substrates	16
2.2.3 Wrinkling	16
2.2.4 Structural feedback	16
2.2.5 Other systems	17
2.3 Summary	18
Chapter 3 Lattice Monte Carlo simulations of graphene growth with a dynamic substrate	19
3.1 Introduction	19

3.2	Methods	20
3.3	Results	24
3.4	Discussion	28
3.5	Conclusion	32
Chapter 4	The random field Ising model with a dynamic field	34
4.1	Introduction	34
4.2	Methods	35
4.3	Results	37
4.3.1	The random field Ising model	37
4.3.2	The dynamic random field Ising model	38
4.3.3	Spin exchange in the dynamic random field Ising model	41
4.3.4	Magnetisation reversal	46
4.4	Conclusion	50
Chapter 5	Off-lattice molecular dynamics simulations of graphene growth on an effective model substrate	53
5.1	Introduction	53
5.2	Methods	54
5.3	Results and discussion	60
5.3.1	Island orientation	65
5.3.2	Growth statistics	67
5.3.3	Diffusion	71
5.3.4	Flake strain energy	73
5.4	Limitations	75
5.5	Conclusion and outlook	76
Chapter 6	Cluster moves in lattice Monte Carlo simulations	78
6.1	Introduction	78
6.2	Methods	80
6.2.1	Domain decomposition in Monte Carlo simulations	80
6.2.2	Incorporating diffusion	83
6.2.3	Incorporating cluster diffusion	84
6.3	Testing and performance	87
6.4	A lattice Monte Carlo graphene growth model with cluster moves	92
6.4.1	Simulation protocol	92
6.4.2	Results	93
6.5	Clusters, dimers and graphene growth	96

6.6 Conclusion	99
Chapter 7 Conclusion	100
Appendix A Supplementary videos	103

List of Tables

3.1	Kinetic and thermodynamic simulation parameters.	23
3.2	Goodness of fit for theory curve on scaling plots	28
5.1	Simulation parameters for MD growth on an effective substrate	58
5.2	Fitting parameters for the four Gompertz type curves	70
5.3	Monomer and dimer diffusion coefficients	72
5.4	Substrate and flake strain energies	75

List of Figures

2.1	Epitaxial growth processes	5
2.2	Surface growth modelling techniques	6
2.3	Typical CVD apparatus	12
2.4	Graphene island orientation from FFM	15
2.5	LEED patterns from graphene on copper	17
3.1	Diffusion moves and island morphology	22
3.2	Effect of varying surface roughness	25
3.3	Time evolution of island size	26
3.4	ISD scaling at different coverages	27
3.5	Contribution to energy from rough substrate	29
4.1	Snapshots from IM and RFIM growth	39
4.2	Absolute magnetisation evolution in the RFIM	40
4.3	Snapshots from growth in the dynamic RFIM	42
4.4	Absolute magnetisation evolution in the dynamic RFIM	43
4.5	Final magnetisation in the dynamic RFIM	44
4.6	Largest domain in the dynamic RFIM	45
4.7	Absolute magnetisation evolution in the dynamic RFIM with spin exchange	47
4.8	Final absolute magnetisation in the dynamic RFIM with spin exchange	48
4.9	Evolution of the magnetisation for domain reversal in the RFIM	49
4.10	Evolution of the magnetisation for domain reversal in the dynamic RFIM	50
5.1	Effective substrate potentials	57
5.2	Dendritic to compact morphology transition	59
5.3	Snapshots from end of growth	61
5.4	Snapshot of a multi-orientation cluster	62
5.5	Evolution of island densities	64
5.6	Graphene island orientation distributions	65

5.7	Angular distribution for (110)	66
5.8	Graphene atomic position on effective potential	68
5.9	Comparative statistics of growth	69
5.10	Flake distortion energy calculation example	74
6.1	Diagram of SL algorithm operation	82
6.2	Diagram demonstrating diffusion in the SL algorithm	85
6.3	Diagram of domain reallocation for a cluster spanning subdomains	86
6.4	Comparison of serial and parallel mean island size	88
6.5	Comparison of serial and parallel autocorrelation function	89
6.6	Simulation times for the SL algorithm	90
6.7	Simulation times for the SL algorithm with cluster diffusion	91
6.8	Snapshots from growth with and without cluster moves	94
6.9	Time evolution of average island size for different proportions of cluster moves	95
6.10	ISDs for different proportions of cluster moves	97
6.11	Diagram of dimer movement in the cluster movement algorithm	98
6.12	Evolution of island densities	98

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

The work presented (including data generated and data analysis) was carried out by the author except in the cases outlined below:

Figure 2.1 was reproduced from reference [120]

Figure 2.2 was reproduced from reference [120]

Figure 2.4 was reproduced from reference [79]

Figure 2.5 was reproduced from reference [141]

Parts of this thesis have been published by the author:

The work presented in Chapter 3 has been published by the author [28].

Abbreviations

CVD Chemical vapour deposition

MBE Molecular beam epitaxy

STM Scanning tunnelling microscopy

SEM Scanning electron microscopy

FFM Friction force microscopy

LEED Low energy electron diffraction

RHEED Reflection high energy electron diffraction

LEEM Low-energy electron microscopy

UHV Ultra high vacuum

XRD X-ray diffraction

DFT Density functional theory

MD Molecular dynamics

REBO Reactive empirical bond order potential

BOP Bond order potential

MC Monte Carlo

kMC Kinetic Monte Carlo

MCMC Markov chain Monte Carlo

VMMC Virtual move Monte Carlo

1NN First near neighbour

2NN Second near neighbour

3NN Third near neighbour

IM Ising model

RFIM Random field Ising model

ISD Island size distribution

Abstract

Graphene is one of a number of layered materials which have tremendous applications in future devices due to their interesting electrical and structural properties. Key to material synthesis is understanding growth mechanisms for thin film formation, and the elimination of defects; consistently creating smooth layers. A number of interesting effects seen in graphene growth highlight the importance of understanding the interaction between the substrate and the growing layer, such as growth on liquid substrates, wrinkling and structural feedback. Chemical vapour deposition (CVD) on a copper substrate has been hailed as a promising scalable route to synthesis, however there are restrictions relating to *in situ* observation. Recent experimental advances have led to graphene grain production of increasing size and purity, yet the early stages of growth remain relatively unexplored.

Current modelling techniques tend to either investigate small systems in a high level of detail, or neglect the substrate detail in a coarser grained model. In this thesis we present results from abstract surface growth modelling which incorporate substrate effects into larger scale simulations. Firstly, we present a lattice Monte Carlo (MC) model of graphene growth on a rough substrate, and show the introduction of a dynamic roughness energy leads to an increase in island size. Secondly, we introduce a dynamic field into the standard random field Ising model, and observe a domain size increase with a dynamic field, but also observe a lower temperature field ordering effect. Thirdly, we construct geometric effective potentials in a Molecular Dynamics model of graphene growth, reproducing experimental island orientation distributions and growth behaviour. Finally, we detail the construction of a cluster moving algorithm for lattice MC simulation, and demonstrate that its implementation leads to an enhancement of island size. Together, these results highlight the importance of substrate roughness and geometry in the early stages of growth.

Chapter 1

Introduction

The synthesis of 3D thin film semiconductors such as GaAs has been fine tuned through over 30 years of research, and can now be grown with a high degree of precision [13], and is supported by a well developed modelling field. They exhibit strong binding between the substrate and the growing layer, with dangling covalent bonds on the substrate providing specific sites for the deposited species to bind to. A good lattice match between the two layers is thus required for epitaxial growth.

Recently there has been increasing interest in 2D materials such as graphene and BN [110] due to their interesting electronic and structural properties. These 2D materials have bonding contained in two dimensions, with no dangling covalent bonds. The 2D material can still be grown epitaxially, but the geometrically strict covalent bonds are replaced by a weaker Van der Waals interaction between layers, giving rise to so called Van der Waals epitaxy [64, 38]. In Van der Waals epitaxy, the requirement for a good lattice match is relaxed, and a mismatch between the geometries of the two layers becomes possible.

Understanding the growth mechanisms of 2D thin films, particularly graphene [69, 120], is key to the consistent production of high quality, low defect materials. Experimental techniques are constantly evolving in order to produce grains of increasing size and purity [69, 151, 110]. Of particular focus is understanding graphene growth on low cost and relatively low purity copper foils [141, 145] via chemical vapour deposition (CVD), which is a promising route to mass production. These foils are polycrystalline in nature, exhibiting multiple different symmetries across the spatial extent of the chip, and understanding the effects of the different symmetry regions has been cited as an area of importance [137, 18]. As these systems are far from the ideal flat growth environment [37], issues such as grain boundaries, facet dependent growth [137], wrinkling [69] and nucleation on copper terraces become important and are far from entirely understood. Growth on substrates near their melting temperature [129], or molten substrate [39, 23] have produced highly regular, single

layer grains.

Performing *in situ* experimental techniques on materials grown via CVD is challenging due to the absence of vacuum and typically high growth temperatures. Due to this restriction, modelling has a key part to play in elucidating the mechanisms of graphene layer formation. There are a wide range of modelling techniques available spanning large spatial and temporal ranges. Incorporating a high degree of system specific detail is computationally expensive and limited in scope, whilst a more general approach can explain features across multiple different systems. The well developed field for 3D modelling often treats the growing layer as strongly bound to the substrate layer, which is not necessarily appropriate for the weakly bound 2D graphene layer.

The aim of this thesis is to simulate the effects of the substrate on a growing 2D system. The models we discuss treat the growing layer as 2D, and a particular aspect of the interaction with the substrate is introduced in an abstract way. Whilst this approach allows us to capture particular aspects of the substrate interaction, we take care to discuss the limitations with respect to the wider literature.

In Chapter 2, we give an overview of surface growth modelling techniques, and a description of experimental findings from CVD of graphene on copper. An understanding of the popular computational techniques helps set the abstract models presented in this thesis in their proper context. The experimental background provides the motivation for our modelling work, and identifies some particularly interesting aspects of graphene growth which we explore in this thesis.

Prompted by reports of growth on nearly molten [129] or liquid substrates [39], we introduce a model for growth on a dynamically rough substrate in Chapter 3. We use lattice Monte Carlo (MC) methods for the graphene layer, and introduce a dynamic roughness energy across the lattice for the substrate roughness. We find that a static roughness energy decreases island size whilst a dynamic roughness energy leads to an increase in island size.

In Chapter 4 we generalise the results of the previous chapter to the random field Ising model (RFIM). The static RFIM is well explored in the context of a model of randomly defective magnets, and has been proposed as a model of surface growth. Whilst the static RFIM introduces a pinning effect which inhibits domain formation, the dynamic RFIM exhibits lattice ordering or domain smoothing. The more general model we present in this chapter could be of interest to the wider statistical physics community.

In Chapter 5, we investigate growth on different facets of polycrystalline copper foil [18, 141] with an MD model of growth on an effective model substrate. The graphene layer gets bond energetics from a recent potential [153], whilst we create an effective substrate potential based on atomic positions in Cu(111), Cu(100), Cu(110) and Cu(210). The graphene grown on the effective potential reproduces experimental growth laws [18] and

orientation [141] with a relatively simple model and we discuss the implications of this for the experimental and modelling communities.

Finally, motivated by recent work reporting the importance of dimers and small islands to the early stages of graphene growth [36, 147, 92], and our own observations of growing clusters from Chapter 5, we introduce a cluster diffusion algorithm to the lattice MC model of graphene growth in Chapter 6. We describe the construction and implementation of this algorithm, and show that its inclusion leads to an increase in island size in simulations of growth.

In each of the four results chapters in this thesis, supplementary videos are supplied alongside the text. These videos are highly informative for showing characteristic features of growth, and we refer the reader to them as appropriate. Short descriptions of each video are supplied in Appendix A.

Chapter 2

Background

2.1 Modelling surface growth

Surface growth is a conceptually simple idea. The growing species is incident upon a flat substrate; small units of the growing species bond to the surface and to each other, eventually forming islands and full layers. An example of this for graphene on copper is shown in Figure 2.1, where we see an incident gas depositing ethene upon the substrate. These molecules eventually form into graphene islands, with hydrogen molecules desorbed from the system. This gives an example of some of the possible mechanisms by which surface growth can happen, and which any model must consider.

The simulation and modelling of atomic processes faces numerous scientific and technical challenges. In order to model any atomic system or process, one has to determine the level of detail, system size and timescale on which to look. From atomistic first principles methods of simulation to Monte Carlo methods to rate equations, there is a rich tapestry of modelling techniques available. A diagram of the spatial and temporal extent of some of these techniques is shown in Figure 2.2. As the level of atomic detail reduces, the extent and scope of simulations increases; from highly detailed *ab initio* methods focussing on small systems, to rate equation methods which neglect atomic positions entirely.

In this thesis we focus on the region labelled as mesoscopic in Figure 2.2, with atomistic detail but without the exact Schrödinger equation solutions from *ab initio* methods. We use abstract modelling techniques in Monte Carlo and Molecular Dynamics models to represent the substrate interaction, rather than directly simulating it. Abstract models are well suited to investigating a particular aspect of a system, but often introduce an unphysical dependency on a particular parameter. It is important therefore to set our work in the proper context and we compare to results from all ends of the modelling spectrum, which we give an overview of here. Some outstanding review articles discuss many of the topics

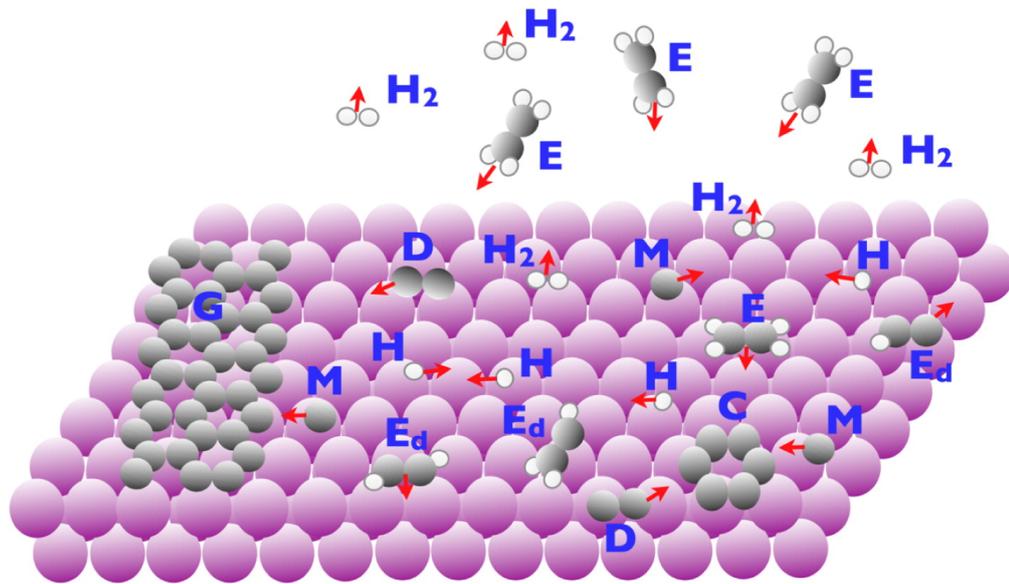


Figure 2.1 Some of the processes undergone during epitaxial growth. The large circles at the bottom represent the substrate, the smaller black circles carbon atoms and the smallest white circles hydrogen atoms. Ethane species (E) diffuse toward the surface, depositing atoms as monomers (M), dimers (D) or other species. Hydrogen atoms can also bond to the substrate, or be evaporated. Carbon atoms can then diffuse until they form islands (G). Reproduced from [120].

we introduce in this section, for example multiscale surface growth simulation [131] and all aspects of graphene growth [120], to which we refer the interested reader.

2.1.1 Density functional theory

First principles methods are popular in accurate material modelling. The solutions to a Schrödinger equation for the electrons and nuclei govern all observable properties of a material, and hence if these solutions can be found, so can the properties. In practice of course this is a many-body, many-termed equation which is difficult to solve, and so some reduction of it is studied in order to capture particular elements of a system.

One such reduction is density functional theory (DFT), which is formulated loosely as follows. The work of Hohenberg, Kohn and Sham [46, 63] showed that the total energy of an electron system in an external potential is a unique functional of the total electron density. The density which minimises the total energy is the ground state density, and this corresponding energy is the ground state energy of the system. This principle of recasting system properties as functionals of the electron density underpins DFT and allows precise calculation of system properties. This early development, and more recent improvements in algorithmic efficiency and scope, has made DFT the go to method for calculating material properties such as band structures, band gaps, and activation barriers for various kinetic

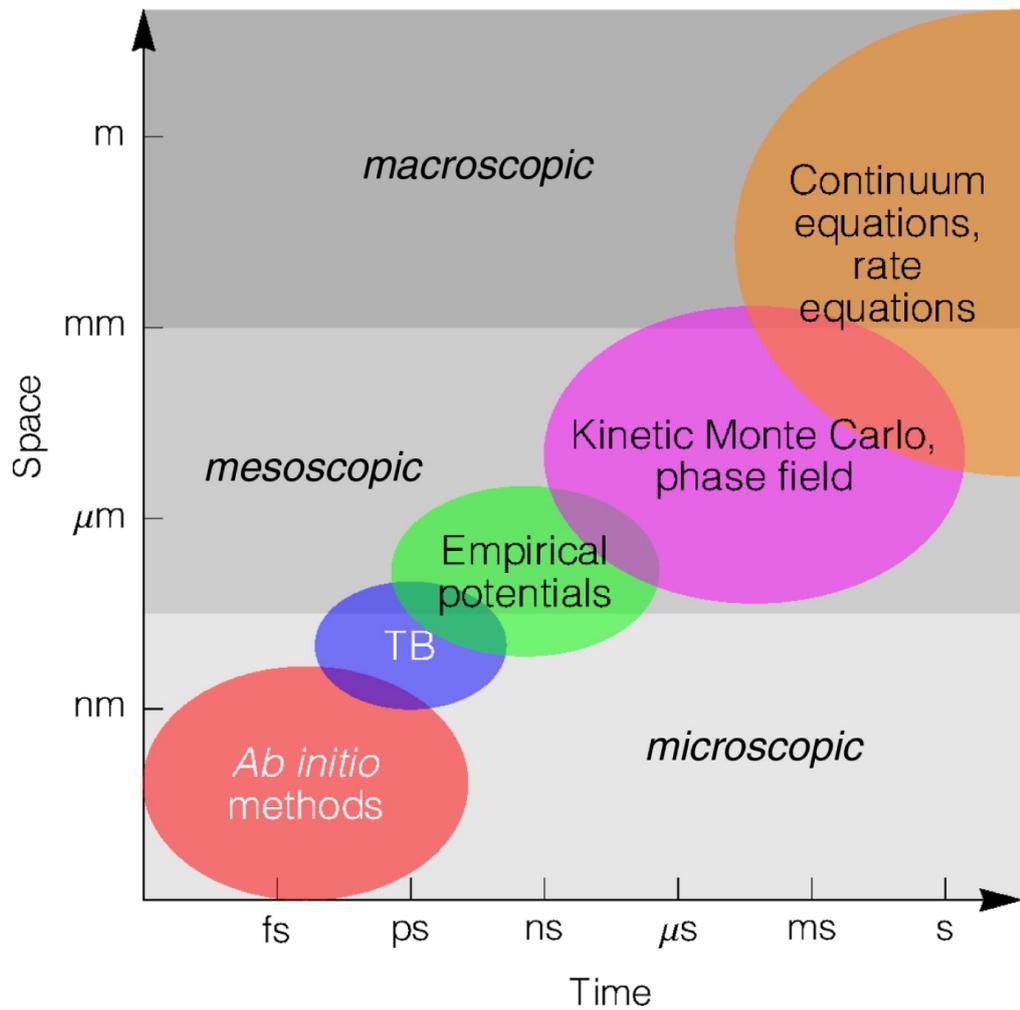


Figure 2.2 Diagram comparing different techniques modelling surface growth by the time and length scales they cover. Tight-binding and semi-empirical models are represented by TB. Reproduced from [120].

processes [62].

There are limitations to DFT; there are errors far from equilibrium, it is computationally expensive, and tends to focus on small systems of up to a few hundred atoms. In the specific context of growth, this presents a barrier at first impression—the scale of growing graphene grains we explore in this thesis for example reach up to millimetre size [148]. However, DFT is effective at elucidating the fundamental mechanisms of growth. In the context of graphene growth this includes investigating the stability of carbon clusters on different copper facets [83], obtaining graphene band gap [41], investigating standard defects [77] and studying early stages of growth [36, 104, 147].

2.1.2 Molecular Dynamics

Molecular dynamics (MD) simulations determine velocities by numerically solving Newton's laws of motion. Whilst *ab initio* methods can be used to determine inter-atomic forces, this becomes progressively more expensive computationally as simulations are performed over longer spatial and temporal scales. This level of detail is frequently removed, and inter-atomic potentials are approximated. As with DFT, large scale simulation packages [99, 121] give users pre-written implementations of pre-parametrised potentials and numerical integrators, which has increased the ease with which MD methods can be implemented.

Timescales are important in MD growth—the integration timestep must be small enough to capture atomic vibrations on a femtosecond timescale, but the simulation length must be long enough to capture atomic and cluster movements which could take microseconds. Entire simulations of aggregation could then take seconds or minutes. Whilst cumulative errors in sampling can lead to errors in MD simulations, averages across the ensemble give good predictions. Ensemble quantities at equilibrium are good estimators to the thermodynamic properties of a system due to the concept of ergodicity—the ensemble averages are equivalent to temporal averages.

The selection of the potential used to define inter-atomic interactions depends upon the overall purpose of the simulation. Potentials fall across a spectrum of complexity, empirical basis and computational cost. For simulations which aim to model large system sizes, or which require a high degree of computation efficiency, pair potentials such as the much studied Lennard-Jones potential [54] give a description of the interactions between any two given atoms, without taking into account the wider atomic locations. These efficient potentials are not necessarily system specific and are good at simulating general properties of systems. Empirical many-body potentials parametrise quantum mechanical calculations with experimental or computational data, and commonly take into account how multiple atoms would interact. These potentials do an excellent job of describing the sys-

tems for which they are parametrised, but often fail to capture systems outside the scope of the original parametrisation. Popular examples include the Tersoff [119] and Brenner [15] potentials, or more recently the reactive bond order potential (REBO) [16]. Quantum mechanical potentials are derived directly from quantum mechanical principles, and are often better at predicting systems outside their original scope. Such potentials have become increasingly relevant as computational capacities increase, with early examples [85] evolving into the more complex ones in use today [153].

2.1.3 Monte Carlo methods

In Monte Carlo (MC) methods, deterministic equations of motion based on real atomic positions and characteristics are replaced with stochastic equations which sample possible system events [32, 100]. This allows events which take a relatively long time in MD or DFT simulations, such as the diffusion of a particle or cluster through space, to be performed based on the energetic favourability of the endpoints. This removes the system time in deterministic modelling spent waiting for an atom to move between local minima which drastically improves the efficiency of simulation and possible temporal scale which can be simulated. This set of transitions between different systems, or *moves* as they are often referred to, could be defined based on experimental or computational observations.

In Markov chain MC modelling (MCMC), we construct a model system which obeys the Markov property. Loosely speaking, this means the evolution of the simulation can be entirely predicted based upon its current state, and has no dependence upon the history of the simulation. The connection of multiple iterations of such a simulation gives it the name chain, and hence MCMC. Considering the time reversibility of the Markov property leads to a concept known as detailed balance: every MC move made must have a reverse move possible at equilibrium. Whilst it is not possible to prove a MC model is ergodic—that is that all points of the corresponding phase space would be visited if the algorithm were to run for infinite time—it is safe to assume a system which obeys detailed balance is ergodic. Note that reversibility refers specifically to the model system and not to the physical system, which does not obey time reversal due to the second law of thermodynamics.

Detailed balance has implications for the transition probabilities between any two states. From elementary thermodynamics, the probability of occupying a given state p_i is given by the Boltzmann probability

$$p_i \propto e^{-E_i/k_B T}, \quad (2.1)$$

where k_B is the Boltzmann constant and T is the system temperature. For a system transi-

tioning from state i to state j the condition of detailed balance must be satisfied,

$$p_i T_{ij} = p_j T_{ji}. \quad (2.2)$$

where T_{ij} is the probability of transition between state i and j . As we know the probabilities from thermodynamics, we must choose the transition probabilities such that

$$\frac{T_{ij}}{T_{ji}} = e^{-(E_i - E_j)/k_B T}. \quad (2.3)$$

One popular choice for these probabilities is the Metropolis choice [82]

$$T_{ij} = \begin{cases} 1 & \text{if } E_i < E_j; \\ e^{-(E_i - E_j)/k_B T} & \text{if } E_i \geq E_j. \end{cases} \quad (2.4)$$

This selection automatically accepts moves which transition the state to a system of lower energy, and accepts moves which transition it to a state of higher energy based on a probability dependent on temperature and the energy difference involved in the transition. Computationally, this is the basis of the Metropolis algorithm [82] for MCMC simulations. Following some initialisation, a series of moves are attempted which evolve the system. The energy cost for each move is calculated, and it is then accepted based on the transition probabilities defined above.

Simulations are said to be in the *grand* canonical ensemble if they are connected to an external reservoir of particles which allows the particle number to change during simulation, and in the *semi-grand* canonical ensemble if the total particle number is fixed but its composition is allowed to change. Care must be taken with simulations in the grand or semi-grand canonical ensemble to ensure that the dynamics of the modelled system are correctly compared to the real system. The number of moves performed does not necessarily correspond to a real system time, and is likely to be a poor representation of one [56]. Rather, these simulations are best viewed as representative of behavioural patterns in real world systems, rather than an exact reproduction. For more detail on the implementation of these methods we refer the reader to one of the excellent texts available [32, 100], and discuss specific implementations as appropriate in this thesis.

In kinetic Monte Carlo (kMC) simulations [9, 34], events are assigned rates according to their favourability, rather than being assigned an acceptance probability. These rates can be derived from barriers for certain processes; for example, a monomer moving between two wells would face an energy barrier to leave the first well. This is in contrast to MC modelling, where the probability of moving between the wells would only be dependent upon the energy of the initial and final state. These rates can be derived from transition

state theory or parametrised by experimental or more detailed work [131]. When an event occurs, the system clock is incremented according to the rate of the processes, as opposed to MC which samples all possible events at every timestep. This provides a link to experimental timescales. The removal of event rejection increases the efficiency of kMC simulations in comparison to MC. This increase in detail, experimental justification and computational efficiency make kMC simulations more common in the context of surface growth modelling than simulations performed in the grand or semi-grand ensemble.

2.1.4 Phase-field modelling

Phase-field modelling is a mathematical modelling technique which has been used to model microscale and mesoscale growth structure evolution in many settings [21, 111]. The principle is to define an order parameter which varies according to the phase of the system across the region. The evolution of the boundary between phases can then be investigated as the system evolves. As well as capturing much of the salient physics in nucleation and grain boundary evolution [123], phase-field models have also been used to capture substrate anisotropic diffusion [80], and substrate step motion [72].

More recent phase-field crystal modelling incorporates crystal structure into the phase field whilst retaining a single order parameter. This style of modelling has proven successful at capturing defect stability and stacking faults [12]. The introduction of a three point correlation function in structural phase-field crystal modelling [112] allows the investigation of structures with a specific bond angle, such as graphene, which then gives atomic level detail on the grain boundaries in an evolving system.

2.1.5 Rate equations and scaling

The most macroscopic way to model growth is to completely neglect atomic positions and model growth via a rate equation approach. A set of differential equations governing the density of different species on the substrate is constructed, most simply of monomers and islands [126]. For example the evolution of monomer density could be formulated as the deposition rate less the rate at which monomers join larger islands. Rate equations can be used to clarify key processes seen in more detailed modelling, which inform the rates used in the construction of the differential equations [51].

Such an approach can accurately capture macroscopic properties of growth such as island size distributions (ISDs). A scaling law for ISDs of islands in epitaxial growth takes the form

$$N_S = \frac{\theta}{\bar{S}^2} f(S/\bar{S}), \quad (2.5)$$

where S is the island size, N_S is the island number density distribution, \bar{S} is the average

island size and θ is the fractional monolayer coverage [65]. The exact form of this scaling function has been analytically investigated using rate equations [2, 98, 25]. This form is often dependent on the critical island size i , defined as one less than the smallest stable island on a surface. If ISDs follow this scaling law across the whole coverage scale, it indicates the same process is occurring at all coverages. Deviations from or changes to the scaling can indicate a change in behaviour. For example in InAs/GaAs, a transition from $i = 0$ to $i = 1$ behaviour of an ISD scaling law was observed due to low quantum dot number density following saturation [65].

2.2 Experimental comparison: chemical vapour deposition of graphene on copper

We focus the experimental comparison for the theoretical work in this thesis on growth of graphene on copper via chemical vapour deposition (CVD), and in this section we describe some aspects of this system which make it a particularly strong candidate for simulation. A typical CVD experimental apparatus is shown in Figure 2.3. The sample is placed in a quartz tube, and influx of hydrocarbon and hydrogen gases into the system is controlled. Heaters then enclose the tube, heating the system during growth. Imaging of this system during growth is made challenging by the lack of access through the heaters which enclose the sample throughout. A zoomed image of the exit tube shows visible sublimated copper deposits, demonstrating that growth conditions correspond to a highly active copper substrate.

This substrate can be films with a highly uniform single orientation [37], or multi orientation foils [49], but it is most commonly Cu(111), due to its similar structural geometry to graphene, and Cu(100) as it is the most common facet in cheaply available copper foils. CVD has been established as the most promising method for cheap scalable material syntheses [110], and significant advances [67, 70] have been made towards production of uniform graphene sheets ideal for new devices. In early stages of growth, graphene forms islands of various shapes and sizes [128, 120] depending upon growth conditions, before growing into monolayer or more than monolayer coverage [50].

After growth, a graphene sheet has almost entirely planar bonding, with no dangling bonds. This is in contrast to more traditional 3D material growth, where the covalent bonds between each layer enforce a rigid lattice matching between layers. The weaker nature of the bonding between the graphene and the copper layer—so called Van der Waals epitaxy [64, 38]—relaxes the lattice matching condition from 3D material growth and allows graphene to grow even on facets without a good geometric match.

This weak interaction after growth allows ascension of surface steps and prompts

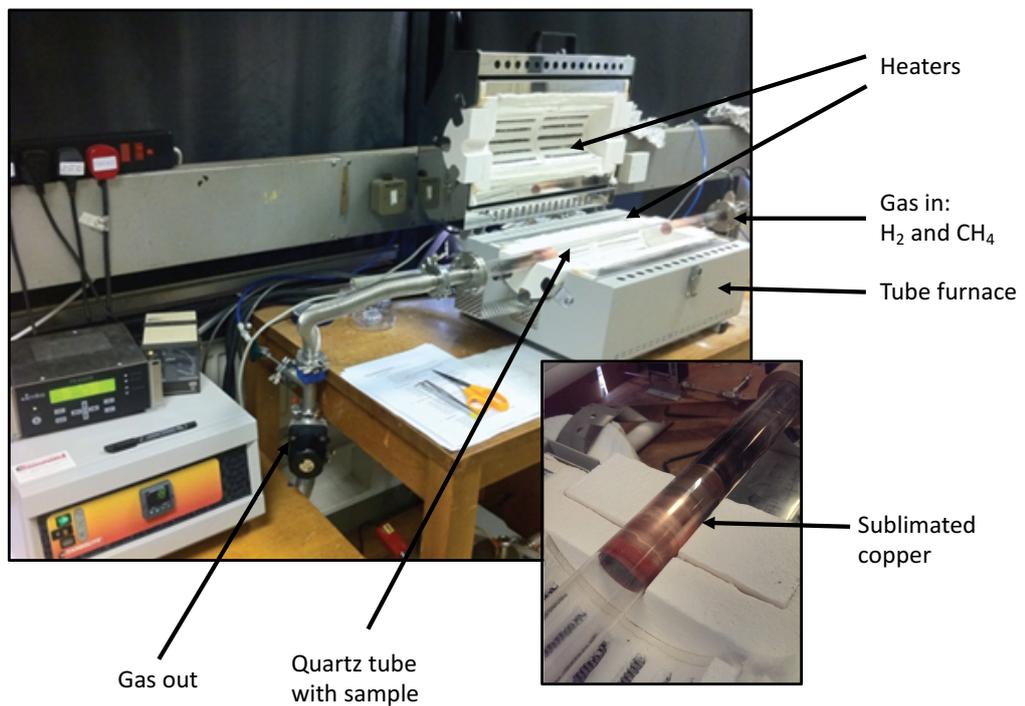


Figure 2.3 Picture of typical experimental CVD apparatus. The sample is placed within the tube furnace, with the heaters enclosing the tube during growth. Methane and hydrogen then pass through the tube, exiting on the left hand side. The zoomed picture shows exit tubes after growth: there is a visible deposit of copper sublimated during the growth. When the top heater is lowered there is no room for the access required for *in situ* techniques.

graphene to grow with many different rotational orientations [69]. Single graphene flakes are able to span substrate grain boundaries [151, 146], and have minimal carbon-carbide formation [29]. Regardless of the weaker bonding, the orientation of graphene flakes has been observed as determined by the copper substrate. Substrate effects have been observed to be important in early stages of growth, with different growth rates, nucleation densities and orientations reported on different facets of polycrystalline foils [141, 145, 137]. The effects of small island formation and diffusion are particularly influenced by substrate effects [147].

2.2.1 Challenges involved in observing a growing system

There are numerous complementary experimental techniques available for the identification and analysis of graphene grown on copper films. Two main categories of techniques are *diffraction* and *microscopy*. Diffraction techniques reconstruct symmetry properties of materials, and allow the geometric structure of the sample to be explored. Microscopy techniques involve imaging or sampling the real space surface properties, from which island sizes, orientations and morphologies can be determined. Most properties of a grown film can be obtained through directed use of these techniques, but they are not necessarily suited to imaging during growth.

Low-energy electron diffraction (LEED) is a diffraction technique for determining crystallographic structure of films, where a columnated beam of electrons is fired at the thin film, and the periodic structure of the surface is reconstructed based on the diffraction pattern of the electron beam. Diffraction patterns are observed in spots or rings, with different patterns corresponding to different identified symmetries. LEED, and similar techniques such as reflection high energy electron diffraction (RHEED) and low-energy electron microscopy (LEEM), probe a surface using a beam of electrons in ultra high vacuum (UHV) conditions ($<10^{-9}$ mbar). CVD typically takes place at atmospheric pressure, and even low pressure CVD is typically performed at 1 mbar [78], making *in situ* measurements challenging as electrons cannot penetrate the gas. Nevertheless, LEEM has been used *in situ* to image graphene grown on copper [143] and Pt [22]. Incremental LEEM images during early stages of growth on copper foils revealed that graphene islands had a four lobed morphology, each lobe with a different rotation about the film normal.

X-ray diffraction (XRD) fires a well focussed X-ray beam at a sample, and reconstructs structure by observing the intensity of the outgoing beam. It is not surface sensitive—X-ray photons have a large penetration depth into matter, and interact only weakly with electronic clouds of atoms. It is also not necessarily well suited to graphene growth, as carbon is a relatively light and difficult to detect atom. Nevertheless, it has been implemented *in situ* for growth using molecular beam epitaxy (MBE) [117], and for

graphene grown on copper [59].

Scanning electron microscopy (SEM) is a microscopy technique which involves scanning a focussed beam of electrons across a surface sample, and constructing a real space image of the surface from back scattered electrons. It has been widely used to perform *in situ* observations of growing materials and has been used to view graphene grown on copper [136] and Pt [137]. The nucleation and orientation of graphene grown on copper studied with SEM was shown to have strong dependence upon substrate dynamics by viewing incremental images and constructing the time dependent coverage. Such *in situ* observations have shown that growth at high temperature occurs on a highly mobile copper surface.

Scanning tunnelling microscopy (STM) is an example of a microscopy technique, which allows for the imaging of a surface down to single atom resolution [55, 52, 120]. STM works through quantum tunnelling—an electrode tip is brought close to a surface, and the composition of the surface is reconstructed based on the tunnelling current from tunnelling electrons. The output comes in the form of an image of the local density of states, and allows the pinpointing of atomic positions. STM requires high levels of precision—the tip must be kept very still and surfaces need to be relatively clean and pure [102]. Imaging at typical growth temperatures for graphene on copper (700 °C-1000 °C) *in situ* is therefore challenging due to issues such as thermal drift, or damage to the probe at high temperature. Despite these restrictions, *in situ* high temperature STM is possible [101], and low temperature STM has highlighted the importance of dimers and small clusters in early stages of graphene growth [92].

Friction force microscopy (FFM) can identify graphene grains at high contrast due to the low friction of graphene [79]. In Figure 2.4 we show an example of how island orientation could be reproduced from FFM mapping of graphene grown on copper, reproduced from [79]. Such a process clearly shows graphene islands, and the quality of this image allows their properties to be determined.

The *ex situ* measurements of graphene grown via CVD on copper can give extensive information about graphene islands, for example determining morphologies [128] or epitaxy on polycrystalline foils [145]. The only way to investigate growth using *ex situ* methods however, is to interrupt a sample grown up to different coverages, and image it after it has cooled. This limits the scope of *ex situ* measurements for investigating growth—to a certain extent the early stages of growth are a black box which remains unilluminated. Measurements taken *in situ* are therefore preferable for understanding the fundamental mechanisms behind island formation. The significant challenges associated with such measurements for CVD growth of graphene on copper described here indicate the potential benefits of simulation, particularly in the early stages of growth.

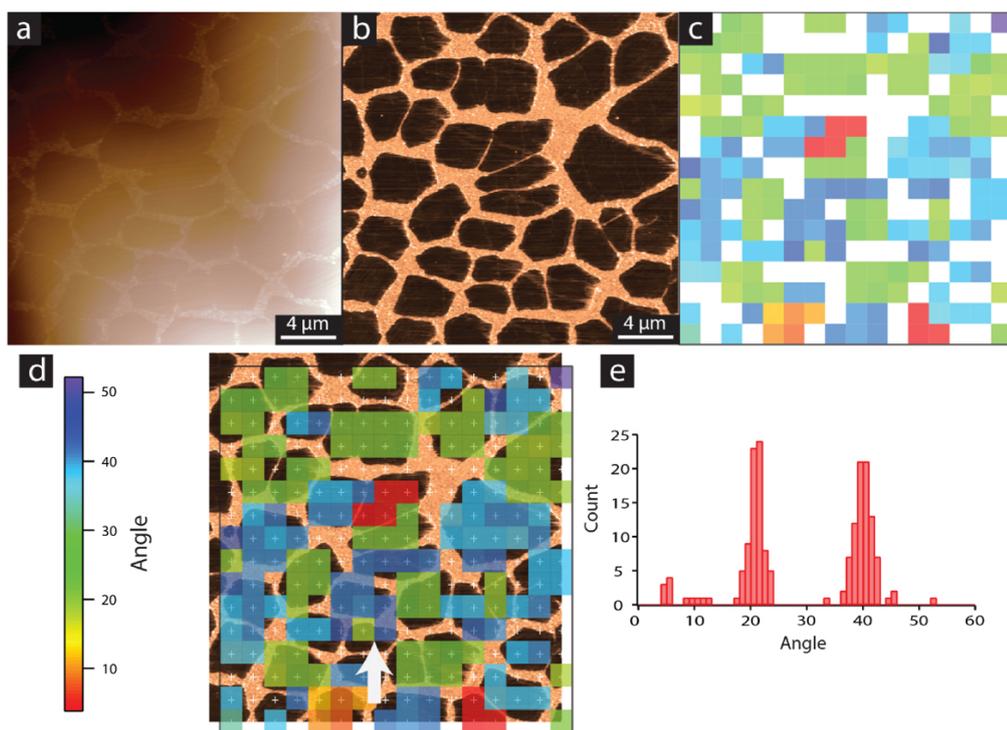


Figure 2.4 Example of reconstructing graphene island orientation from FFM images taken after growth: (a) topography, (b) FFM map of graphene on copper after growth, (c) orientation map taken from fast Fourier transform analysis of grid of images, (d) orientation map overlaid on FFM map, with white crosses marking the grid and a colour scale showing angle and (e) the distribution of island orientation obtained from this process. Reproduced from [79].

2.2.2 Growth on liquid substrates

Copper has a melting point of 1085 °C, which is close to the temperatures used for CVD growth of graphene on copper. The copper deposits in Figure 2.3 point towards the activity of the copper substrate during growth, and growth on nearly molten copper substrates have been reported [95]. Recently, graphene has been grown on liquid Cu [39] and Ga [23], the liquid copper substrate resulting in uniform hexagonal graphene flakes. The liquid substrate results in a lower carbon atom diffusion rate and removes grain boundaries in the copper substrate, resulting in fewer nucleation sites, and a lower nucleation density. Higher island growth rates were thus reported compared to growth on a solid Cu substrate, alongside the formation of fewer, larger islands due to the lower nucleation density. These islands retained the hexagonal shape and self-assembled to a compact, ordered monolayer surface. The purity of the first layer is speculated to remove the nucleation sites for the formation of the second layer.

2.2.3 Wrinkling

Graphene grown via CVD tends to be grown at high temperatures, and then cooled to room temperatures. The graphene and copper have different thermal expansion coefficients, with the thermal expansion coefficient of graphene reported as negative at some temperatures [150]. The epitaxy established at high temperature must be maintained after the sample is cooled. Upon cooldown the different lattices contract in different proportions, leading to a strain. Strain must be absorbed in some way, either as stress or local decoupling of the graphene and copper layers. This excess strain goes into graphene wrinkling [69]. Wrinkling is a macroscopic effect which we do not aim to account for in the modelling in this thesis.

2.2.4 Structural feedback

A feedback effect between the growing graphene grains and the copper substrate has been reported [141] in growth on cheap, predominantly Cu(100) foils. Graphene formation drives a restructuring to (n10) facets in the copper layer, which would not otherwise be present in the copper. These facets in turn affect the orientation of islands in the graphene from the onset of growth. This indicates the importance of incorporating C–Cu interaction for a highly mobile copper substrate into simulations of graphene growth. Figure 2.5 shows LEED patterns reproduced from this paper at different beam energies, before and after ion bombardment and annealing. The ion bombardment strips off the graphene layer, and reveals the crystallography of the underlying substrate. Different energy beams resolve

different symmetry on the substrate. The presence of distinct spots in this LEED pattern indicate the graphene islands are oriented with the surface.

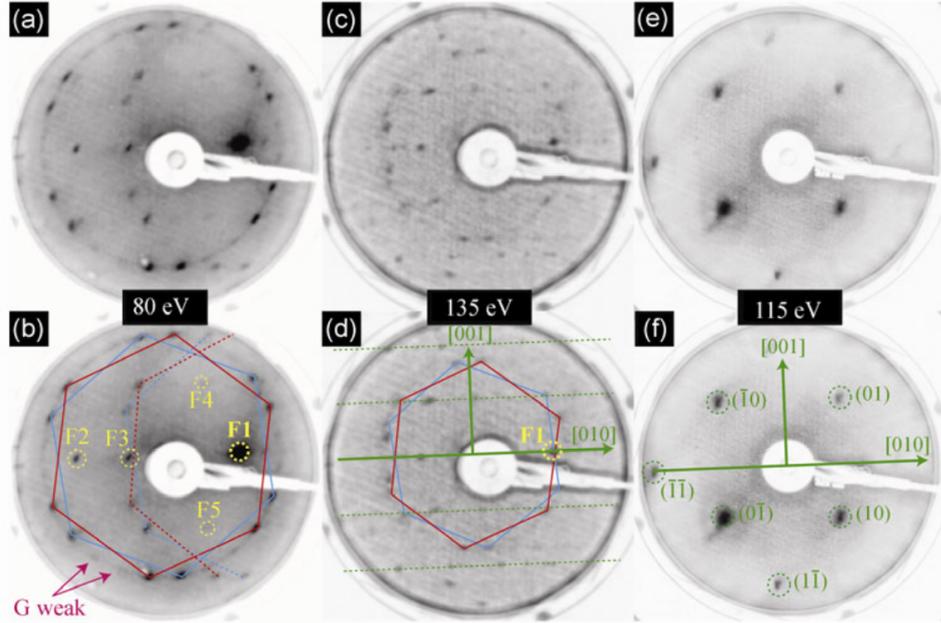


Figure 2.5 LEED patterns from graphene grown on copper before, (a) and (c), and after, (e), ion bombardment and annealing. The images in (b), (d) and (f) are the same as (a), (c) and (e), with annotations added to highlight the symmetries observed. Beam energies are as labelled. The disappearance of the hexagonal spots in (b) and (d) after bombardment in (f) indicate the orientation of islands with the substrate. Reproduced from [141].

2.2.5 Other systems

The models applied to graphene growth on copper in this thesis could with little effort be transferred to models of different growing systems which exhibit similar behaviour. Whilst the effects of growth on copper make it a particularly good candidate for simulation, graphene has been grown on a variety of substrates using different methods, for a comprehensive overview of which we refer the reader again to an excellent review which deals with all aspects of graphene growth [120]. Outside of graphene growth, the formation of other 2D materials such as BN [144, 40], $\text{Bi}_2(\text{Te},\text{Se})_3$ [53, 45] and $\text{Mo}(\text{S},\text{Se})_2$ [113, 127] have attracted similar interest. These systems face many of the same challenges, of orientation, mismatch and substrate roughness during growth.

2.3 Summary

In this chapter we have reviewed the theoretical methods available to model growing systems and discussed some aspects of graphene grown on copper via CVD which make it a good candidate for modelling. In particular we note that graphene forms in a two dimensional plane, and interacts weakly with the substrate. This formation of a 2D material under constraints from a substrate is at odds with the traditional modelling approach which considers 3D material growth. Motivated by this, the models we present in this thesis are all 2D, but contain additional constraints from the substrate, which we introduce in an abstract way. Constraints which we explore in this thesis include growth on a rough substrate, motivated by observations of graphene island formation on liquid copper and on a high temperature copper substrate which is highly mobile during growth. The highly mobile substrate also undergoes a structural feedback effect, which demonstrates the importance of the interaction between the copper and graphene layers.

Limitations on *in situ* experimental techniques leave the early stages of graphene growth largely unexplained. DFT simulations can explore very early stages of growth and identify potentially important species in island nucleation. Bridging this gap between islands of a few atoms and islands of 1000s of atoms is a task well worthy of computational modelling.

Chapter 3

Lattice Monte Carlo simulations of graphene growth with a dynamic substrate

3.1 Introduction

The first characteristic feature of graphene growth on copper we investigate in this thesis is surface roughness. A number of the effects we described in the previous chapter indicate the high level of interaction between a mobile copper surface and the growing graphene layer. We discussed recent findings of graphene growth on nearly molten [129] and on liquid surfaces [39, 23], particularly noting that the latter led to formation of highly regular, single layer graphene islands. We also saw from a typical experimental CVD environment (Figure 2.3) that a great deal of copper sublimation occurs during a typical surface preparation and CVD growth run. Finally, a structural feedback effect has been noted, whereby the copper surface restructures by faceting *only* after CVD growth of a graphene overlayer [141]. Combined, these observations strongly imply that the Cu substrate can be far from equilibrium during CVD growth, and hence cannot be considered as a perfectly static crystal facet. The role of substrate roughness in controlling graphene nucleation has been described as pivotal [60], but this role has yet to be included in any kinetic growth model based on rate equations.

While recent density functional tight binding (DFTB) simulations [68] have probed the early stages of graphene nucleation on semi-molten copper, these cannot access the wider range of time and length scales over which important processes occur [131, 8]. These range from atomistic events on a timescale around 10^{-12} s to the scale of hundreds of microns and minutes for graphene grain completion. Monte Carlo (MC) models allow micro-

scopic events to be aggregated efficiently so that, for example, the nucleation, growth and coalescence of 2D islands [17, 130, 86, 47] or 3D nano-clusters [33, 109, 154] on a surface can be studied. The key ingredient of a MC model is the list of microscopic events which can occur and their rates or probabilities. Such models are often constructed on a static lattice: monomers (atoms) can occupy discrete sites, which are identical in the substrate, so that only occupancy by monomers in the dynamic growing layer differentiates the sites. This type of model does not seem to be appropriate for a growth system where the substrate is highly active during growth, such as copper in graphene CVD.

A great deal of insight into fundamental surface growth processes can be gained by studying growth well below monolayer coverage, i.e. when monomers have aggregated to islands which do not completely fill the layer and have not begun to coalesce. In the previous chapter we discussed how deviations from the dynamic scaling behaviour of the island size distribution (ISD) at this stage of growth can uncover scale dependent behaviour [65]. Experimentally, one can use a technique such as scanning tunnelling microscopy (STM) to study island formation and growth with atomic-scale precision, normally quenching the sample after sub-monolayer growth and working *ex situ* [55, 52, 120]. Atomic resolution STM is well suited to ultra-high vacuum growth techniques, and can even be performed *in situ* [124]. Performing STM in a CVD growth environment [101], is very challenging: more generally, *in situ* CVD growth monitoring is far from routine [74, 102] especially on polycrystalline and non-planar substrates such as Cu foil. This makes modelling work particularly appropriate to graphene growth on a mobile copper substrate.

In this chapter, we report simulations on the early stages of graphene growth using a minimal MC model constructed to mimic the semi-molten dynamically rough nature of a hot copper substrate. While the effect of dynamic roughness on individual ad-atom diffusion has been studied in lattice-gas models [118], the effect on growth has not been previously quantified. We find that dynamic substrate disorder actually *enhances* the growth of large, regular islands. By contrast, static disorder hinders the growth of large islands, compared to a uniform lattice, and destroys ISD scaling. Our dynamic substrate approach is applicable to many surface growth systems.

3.2 Methods

We study an abstract lattice-gas growth model, with parameters motivated by graphene CVD growth on copper. This consists of Metropolis MC on a periodic, two dimensional honeycomb lattice, simulated in the semi-grand ensemble [61]. We use a honeycomb lattice to replicate the coordination number of carbon in graphene. In such models, precise structural details are abstracted into an effective picture. Our choice of lattice symmetry

is hence essentially arbitrary and is not in any way intended to reflect the symmetry of preferential absorption sites on a faceted copper surface.

Lattice sites are occupied by either hydrogen (H) or carbon (C) atoms. Energetics are captured via nearest neighbour interactions with Hamiltonian

$$H_0 = \sum_{\langle i,j \rangle} E_{CC} s_i s_j + \sum_i \delta_{s_i,0} \min_j (E_{CH} \delta_{s_j,1}) \quad (3.1)$$

where $s_i = 0$ ($s_i = 1$) if site i is occupied by H (C). E_{XY} is the effective bond energy between two atoms of species X and Y. The first term in equation 3.1 runs over all nearest-neighbour pairs, and in the second term j runs over the nearest neighbours of each lattice site i . We assume hydrogen is in excess, and that there is no overall system gain from H–H bonds by setting E_{HH} to 0 and have no term corresponding to it in the Hamiltonian. The final term only has non-zero value when a hydrogen site has carbon neighbours, and then has value E_{CH} regardless of the number of carbon neighbours, corresponding to the hydrogen site only interacting with a single Carbon neighbour. This selection of parameters, where $E_{CH} < E_{HH}$, effectively leads to hydrogen terminated carbon clusters. We adopt this construction to reflect realistic valence behaviour in an abstract fashion, but note that detailed bonding constraints and topology are not included.

The relevant thermodynamic potential is

$$G = H_0 - \mu_C N_C - \mu_H N_H, \quad (3.2)$$

where μ and N are the chemical potentials and species number respectively. We work in reduced units such that $E_{CC} = -1$ corresponds to the strength of a C–C bond in graphene relative to the H–H bond in an adsorbed H_2 dimer, i.e. $E_{HH} = 0$. On this energy scale, a C–H interaction strength E_{CH} of -0.1 captures an energetic penalty to forming interfaces between graphene flakes and the hydrogen saturated surface.

Provided $E_{CH} \ll E_{CC}$ the exact choice of this parameter does not significantly alter the characteristics of simulated growth. In these units, a temperature of $T = 1000$ K scales to a lattice temperature of ~ 0.01 .

Simulations are initialised with hydrogen (assumed to be in excess), occupying every lattice site. All simulations reported here were performed on a lattice with $N = 36864$ sites, and consist of growth and annealing phases. During growth, carbon is inserted into the lattice via transmutation of H into C, capturing the displacement of molecular hydrogen by hydrocarbons during CVD. In addition to transmutation, our simulations model diffusion of carbon via exchange with H atoms on lattice sites within the same hexagonal unit. Possible diffusion moves are represented in Fig. 3.1. Inclusion of moves which are equivalent by

symmetry results in 12 move targets, which are selected with equal probability.

To motivate this model selection, Fig. 3.1 also shows representative snapshots from simulations with nearest-neighbour moves, as well as using all 12 diffusion targets. In the former case, carbon islands form disjointed fractal shapes. The greater isotropy of the latter case generates smoothly terminating regular islands, representative of experimental graphene islands.

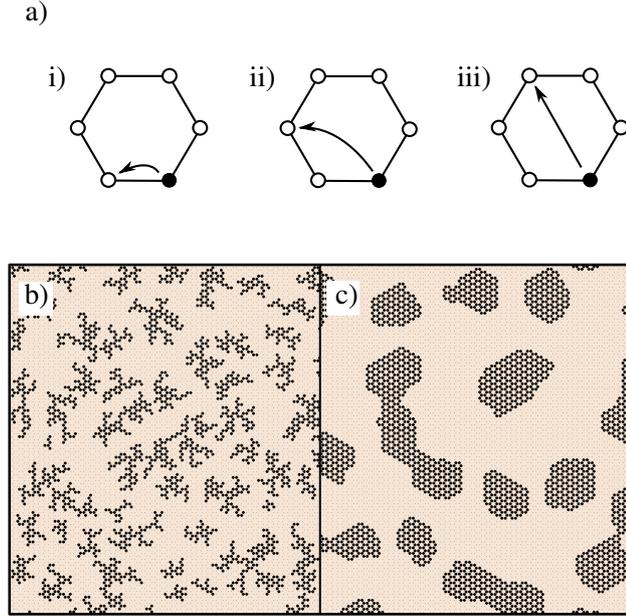


Figure 3.1 Diagram of permitted diffusion moves, and their effect on island formation. Panel a) shows three possible moves, to first, second and third nearest neighbours (1NN, 2NN, 3NN respectively) within local hexagons, in i), ii) and iii) respectively. Panel b) shows a representative snapshot of growth with only 1NN moves, and panel c) a snapshot with 1NN, 2NN and 3NN moves. Note that with only 1NN moves carbon islands are fragmented, but regular in shape when all three moves are included.

A single MC sweep consists of N trial moves, each attempted on a randomly selected lattice site. We interpret our simulations on a notional time scale by connecting the mean square displacement $\langle \delta r^2 \rangle$ of single C atoms over an MC sweep to a timestep via $\delta t = 4D \langle \delta r^2 \rangle$. For convenience, we choose D such that $\delta t = 1$, i.e. time is incremented by one unit per MC sweep.

During growth, transmutation and diffusion moves are attempted with ratio $F = 10^{-5}$ up to a fixed coverage θ . With this choice of F , the timescale of diffusion is far longer than the timescale of insertion. During the annealing phase only diffusion moves are permitted, for a further t_A sweeps of the system. In principle, results can be scaled to real time units via experimental measurements of effective D and F , however such data are not typically accessible.

Chemical potentials are chosen such that insertion moves are effectively always accepted ($\mu_C > \mu_H$). As the simulation temperature is low compared to the effective carbon-carbon bond energy, events which involve breaking these bonds once formed, are rare. Growth is irreversible under these conditions.

Surface roughness is introduced into the model by assigning each lattice site i a “roughness energy” ε_i drawn from a rectangular distribution of width ξ , centred about 0.

$$H_r = H_0 - \sum_i s_i \varepsilon_i \quad (3.3)$$

This reflects the spatial variation in carbon attachment energy expected of a disordered substrate. The exact shape of distribution is inconsequential to the results described in this chapter, provided it is symmetric about zero.

As well as the roughness amplitude captured by ξ , we model the *mobility* of the substrate roughness. Sites are permitted to exchange their roughness energies ε_i with that of their neighbours. This roughness move is subject to the same Metropolis acceptance criteria as the diffusion moves in Fig. 3.1, but using a separate ‘substrate temperature’ T_s , interpreted as a roughness mobility parameter. Decoupling these moves from the simulation temperature allows us to vary this parameter from a totally static surface roughness ($T_s = 0$) through to a freely diffusing molten substrate ($T_s = \infty$). These roughness moves occur with the same frequency as carbon diffusion moves, and varying this frequency was seen to have little to no effect on the final island size distribution.

Parameter		Value
N	Grid size	36864
F	Deposition ratio	10^{-5}
t_A	Anneal time	50000 sweeps
T	Temperature	0.01
E_{CC}	C–C bond energy	−1.0
E_{CH}	C–H bond energy	−0.1
E_{HH}	H–H bond energy	0.0
μ_C, μ_H	Chemical potentials	5, 3

Table 3.1 Kinetic and thermodynamic simulation parameters.

For typical examples of growth in three key cases ($\xi = T_s = 0$, $\xi = 1.2, T_s = 0$ and $\xi = 1.2, T_s = \infty$), we refer the reader to the supplementary videos, details of which are provided in Appendix A. For reference, a full list of kinetic and thermodynamic parameters is shown in Table 3.1.

The code for this model was written in parallel in C with MPI, using the Synchronous Sublattice (SL) algorithm. This algorithm divides the full lattice into domains,

and executes MC moves on subdomains of those domains at the same time, allowing for faster simulations, and for the investigation of larger lattices. For a discussion of the implementation of the SL algorithm with atomic diffusion and a wider discussion of parallel lattice simulations, see Chapter 6.

3.3 Results

Varying the roughness amplitude and mobility parameters has a dramatic effect upon island size and formation. In the static case ($T_S = 0$), increasing the roughness amplitude ξ leads to formation of small, fragmented islands at concentrations of favourable lattice sites. In the maximally dynamic case ($T_S = \infty$), islands are destabilised, causing them to move, transform, and dissociate freely. In a critical range of ξ values, from roughly 1 to 1.5, this leads to formation of fewer, larger islands with a significant fraction of carbon mobilised in monomers or small clusters. Example surfaces at different roughness parameters are shown in Fig. 3.2. Quantitative analysis is based on mean island size \bar{S} and ISDs. With dynamic roughness, many small islands and monomers appear, making ISDs difficult to visualize. Hence we count only islands above a threshold size of $S = 10$ for these surfaces.

Figure 3.2 shows the effect of varying ξ on \bar{S} . Dynamic cases show peaks in island size beyond $\xi = 1.0$, the effective C–C bond energy, and peaking at $\xi = 1.2$ for $T_S = \infty$. A snapshot is shown for $T_S = \infty$ at $\xi = 1.2$, and large islands are clearly visible. At higher values of ξ , even large islands are no longer stable, and the surface becomes dominated by smaller fragmented islands, as shown for $T_S = \infty$ in a snapshot at $\xi = 1.5$.

Even small values of ξ see significant reduction in island size for the static case, with any regular island structure disintegrating. There is no significant difference in behaviour above or below $\xi = 1.0$, the islands just get smaller and more localised to favourable surface regions.

Typical evolution of \bar{S} during growth and annealing is shown in Fig. 3.3. The smooth and static cases both show a continuous increase in island size during growth, and minimal changes during annealing, as a stable structure is reached. The dynamic case has growth up to larger island sizes, but during the annealing phase islands continue to grow. When the annealing phase is extended, islands continue to shift and reform on the dynamic substrate, steadily increasing island size.

A dynamic scaling relation for island size distribution can be obtained by assuming that at a given stage of growth there is only one length scale relevant to the problem. If we assume this length scale to be the mean island size, \bar{S} , we can write the number density of islands of size s as a function of coverage θ , $N_S(\theta)$, as a function of coverage in terms of a scale independent distribution function, $f(S/\bar{S})$ and a function containing the scale and

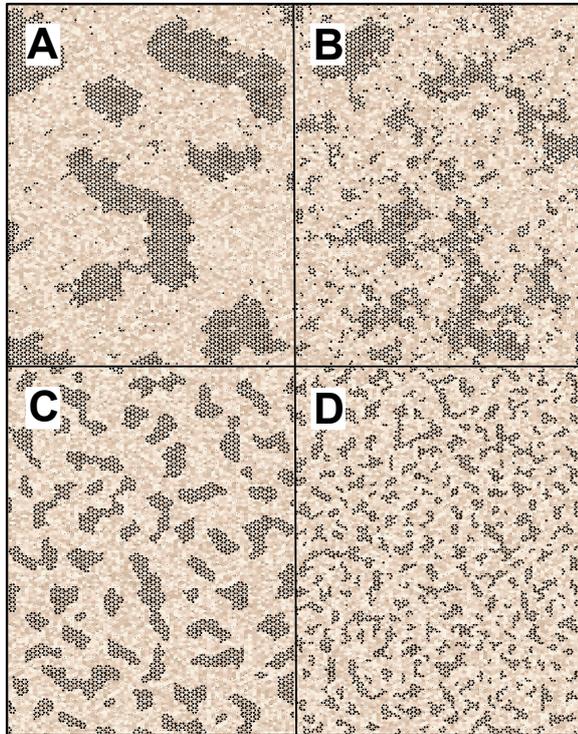
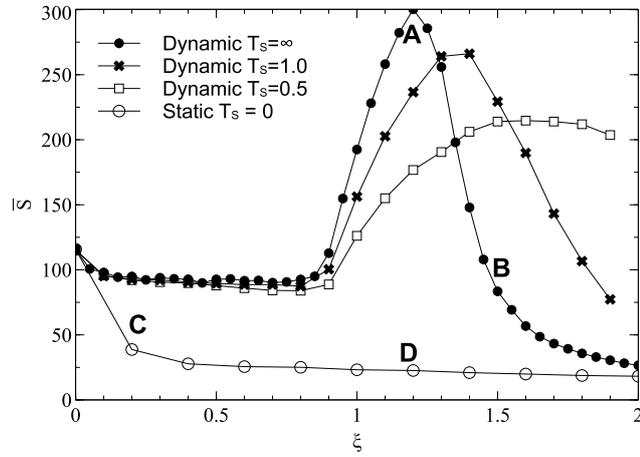


Figure 3.2 Varying strength of surface roughness parameter ξ and mobility parameter T_s against the average island size for islands greater than 10 in size. Data are averaged over 10 trajectories each of which contains typically 30-100 islands. Standard error in the resulting mean of \bar{S} is smaller than the symbol size at each point. Four snapshots from simulations, taken after growth to a fixed coverage and annealing, are shown, for $\xi = 1.2$, $T_s = \infty$ (A) $\xi = 1.2$, $T_s = 0$ (D), $\xi = 1.5$, $T_s = \infty$ (B), and $\xi = 1.5$, $T_s = 0$ (C). The colour scale represents the roughness energy, lighter shades representing positive values. Snapshot sizes represent approximately 15% of the simulation area. All simulations used parameters described in Table 3.1.

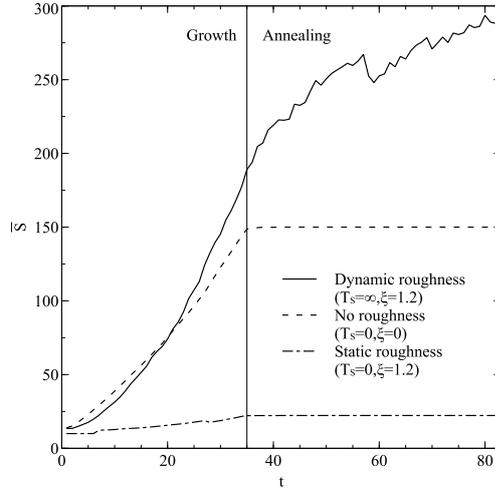


Figure 3.3 Time evolution of average island size for islands greater than 10 atoms in size. Three lines are shown, for no, static, and dynamic roughness. All simulations had roughness strength $\xi = 1.2$, and used parameter values in Table 3.1. A vertical line separates the regimes of growth and annealing.

coverage dependence, $A(\bar{S}, \theta)$

$$N_S(\theta) = A(\bar{S}, \theta) f(S/\bar{S}). \quad (3.4)$$

Integrating both sides over all island sizes yields an expression for the total island density N

$$N \approx \frac{\theta}{\bar{S}} = A(\bar{S}, \theta) \bar{S} \int f(u) du, \quad (3.5)$$

which leads to the identification of $A(\bar{S}, \theta) = \theta/\bar{S}^2$, and thus the dynamic scaling relation

$$N_S(\theta) = \frac{\theta}{\bar{S}^2} f(S/\bar{S}), \quad (3.6)$$

where θ is the fractional surface coverage.

The distribution $f(S/\bar{S})$ has been estimated analytically for non-reversibly aggregation [98] using the Wigner surmise (WS)

$$P_\beta(s) = a_\beta s^\beta \exp(-b_\beta s^2), \quad (3.7)$$

with a_β set by normalisation

$$a_\beta = 2\Gamma\left(\frac{\beta+2}{2}\right)^{\beta+1} / \Gamma\left(\frac{\beta+1}{2}\right)^{\beta+2}, \quad (3.8)$$

and b_β set by unit-mean conditions

$$b_\beta = \left[\Gamma\left(\frac{\beta+2}{2}\right) / \Gamma\left(\frac{\beta+1}{2}\right) \right]^2. \quad (3.9)$$

In our case the fluctuating variable, $s = S/\bar{S}$, and $\beta = (2/d)(i+1)$ is the sole parameter of the WS, set by the critical island size, i^* . The critical island size is one less than the smallest number of atoms required to form a stable island. For example in simulations without roughness we have a critical island size of 1; islands containing two or more atoms are stable. On a dynamic rough surface, the aggregation process is no longer irreversible as islands cleave, move and reform frequently, and thus this theory will not necessarily hold.

The ISDs taken from smooth, static, maximally dynamic after growth and maximally dynamic after annealing cases are shown in Fig. 3.4 in the upper four panels. Applying the dynamic scaling relation in Equation 3.6 to these ISDs gives the distributions shown in the bottom four panels. Also plotted is the curve from the Equation 3.7 corresponding to $i^* = 1$. The goodness of fit of each of the scaled curves against this theoretical curve is shown in Table 3.2.

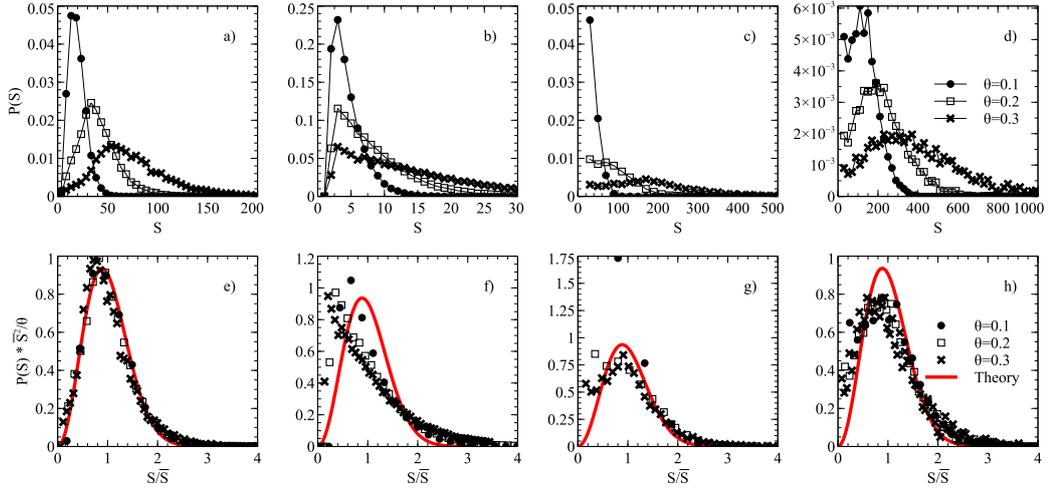


Figure 3.4 Scaling relations of ISDs taken from simulations grown up to coverages of $\theta = 0.1, 0.2, 0.3$ of a smooth surface after annealing (a) and (e), static roughness surface after annealing (b) and (f), maximally dynamic roughness surface after growth (c) and (g), and a maximally dynamic ($T_s = \infty$) roughness surface after annealing (d) and (h). The top pictures (a-d) show unscaled ISDs, and the bottom pictures (e-g) scaled ISDs according to relation described in the text. The red line corresponds to the theoretical form of the $i^* = 1$ curve. All roughness simulations used $\xi = 1.2$, and other parameters described in Table 3.1. Error bars are of the order of the point size.

ISDs from simulations on a smooth surface collapse onto a single curve, which is well represented by an analytical form of the scaling function associated with $i^* = 1$, with p-

Substrate case	Coverage (θ)	Pearson chi-square statistic	p-Value
No roughness	0.1	0.686	0.984
	0.2	6.13	0.963
	0.3	19.7	0.874
Static roughness	0.1	61.2	0.00
	0.2	176	0.00
	0.3	450	0.00
Dynamic roughness after growth	0.1	244	0.00
	0.2	131	0.00
	0.3	207	0.00
Dynamic roughness after annealing	0.1	139	0.00
	0.2	142	0.00
	0.3	270	0.00

Table 3.2 Chi-square test statistics and significance for each ISD shown in Fig. 3.4 against the analytical estimate of an ISD with critical island size $i^* = 1$, shown in Equation 3.7.

values all falling well above significance range to indicate the curve is a good fit. ISDs from static substrate roughness do not collapse onto a single curve, and all p-values in Table 3.2 are 0, indicating the analytical curve does not describe any of the scaled distributions well. The situation for dynamic substrate roughness is more complicated. ISDs from dynamically rough substrates do not collapse onto a single curve immediately after growth termination. However, the post-growth annealing process does result in universal scaling of the ISDs, however the expression in Equation 3.7 is not a good fit to the scaled data. The scaled ISDs are actually broader than the $i^* = 1$ form and have a peak to the left of 1.

To investigate the evolution of surface roughness during growth, the total carbon-substrate interaction energies for dynamic and static roughness are plotted in Fig. 3.5. In the static case, carbon islands form above favourable regions of the lattice, and as such the total carbon-substrate interaction energy is low. In the dynamic case, however, the energy of sites underneath carbon atoms is relatively high, suggesting that the substrate lattice does not reorder itself underneath carbon islands. There is a change in gradient at around $\xi = 1.0$ in both the static and the dynamic cases, corresponds to the beginning of peaks in \bar{S} shown in Fig. 3.2.

3.4 Discussion

Our simple MC scheme, including 2NN and 3NN moves shown in Fig. 3.1 behaves entirely as expected in the absence of roughness. Graphene islands observed in simulations tend to a regular hexagonal or circular shape, although when two islands meet during growth they are often unable to reform into optimal shapes leading to extended anisotropic morpholo-

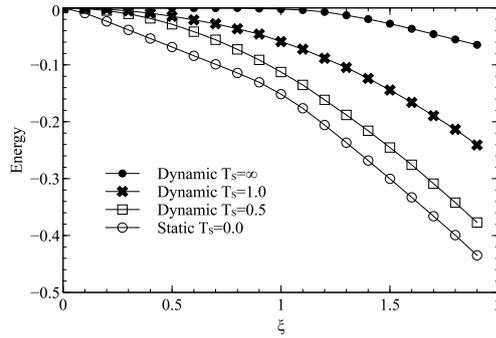


Figure 3.5 Varying strength of surface roughness, ξ , for static and dynamic roughness, against the average substrate energy under a carbon atom. All simulations used parameters described in Table 3.1.

gies. Nonetheless, there is a clear preference for zigzag termination and corners of 120° . Graphene islands on copper have been observed with compact (hexagonal), four-lobed and dendritic shapes, depending on the substrate symmetry and growth conditions [110].

The ISDs (Fig. 3.4 (a) and (e)) produced obey the expected dynamic scaling relation, with ISDs at different coverages collapsing onto a single scaled curve. This demonstrates that the fundamental assumption behind Equation 3.6 is upheld: the same processes of irreversible growth occur across different length scales. The ISDs are well described by a scaling function derived from nucleation theory and observed in a wide variety of surface growth systems [2, 65, 25].

The case of static roughness leads to a drastic reduction in growth, at even low ξ . Islands formed remain small and random in nature, the preference for moving onto a favourable substrate site rather than forming C–C bonds effectively eliminating large island formation. The ISDs (Fig. 3.4 (b) and (f)) do not collapse onto a single curve under a scaling relation, and do not match the theoretical curve. This breakdown of scaling suggests that static roughness has introduced a fixed length scale onto the surface, namely the mean distance between energetically favourable sites, loosely defined by the shape of the energy distribution.

The total carbon-substrate energy decreases linearly with ξ , suggesting that increasing roughness simply makes the sites which carbon atoms select more favourable, rather than affecting the underlying mechanism of growth. The change in gradient at $\xi = 1.0$, the effective C–C bond energy, suggests an increased preference for substrate sites over C–C bonds, but does not produce a noticeable difference in resulting island morphologies.

The case of dynamic roughness, by contrast, leads to an enhancement of island size during the annealing stage of simulation. Islands grow to a larger size than on a smooth surface, even after only the growth stage, and continue to grow during annealing. The

constantly shifting surface roughness prevents kinetic trapping, allowing regular islands to reform, cleave, and move across the surface.

ISDs after growth but before annealing (Fig. 3.4 (c) and (g)) have substantial amounts of carbon atoms as small islands coexisting with large islands. At low total coverage θ there is a large contribution to the ISD from such small islands which reduces in weight as θ increases, destroying ISD scaling. In semiconductor heteroepitaxy, similar deviations from ideal scaling behaviour have been interpreted as due to scale-dependent interactions imposed on a system by surface reconstruction [52], elastic strain [65] or both [11]. By contrast, in the present case the origin of the loss of scaling is purely dynamic because the dynamic roughness has a disproportionate effect on the smaller islands.

ISDs after annealing (Fig. 3.4 (d) and (h)) see much of the mobile carbon being agglomerated into islands, with more defined peaks and heavier tails. Here the ISDs do collapse onto a single curve. Once islands reach a certain size they become resistant to the cleaving effects of dynamic roughness and so scale-free behaviour is recovered. However, the scaled ISDs do not follow the conventional $i^* = 1$ distribution, with a slightly broader and flatter shape. This is not surprising given that dynamic roughness enhances *both* island cleaving and island growth, leading to a broader distribution. Since increasing the value of i^* typically sharpens the peak of the scaled ISD and $i^* = 0$ ISDs are typically monotonically decreasing [2, 65, 25], this altered scaling form suggests that the effects due to a combination of dynamic roughness and annealing could not be captured in any standard irreversible aggregation picture. One approach to characterising heavy tailed ISDs which fall outside of scaling forms is to fit empirical functions to the low S/\bar{S} and heavy tail behaviour [24], but such an approach did not give statistically satisfactory results here.

These conclusions are not greatly affected by the choice to measure ISDs neglecting the smallest islands ($S < 10$) for the dynamic roughness case. This choice simply allows us to display the peak of larger islands more easily and the value of the cutoff simply changes the coverage θ at which deviations from scaling become apparent. When comparing to the static roughness case, the central qualitative point is that there is *no* broad tail of large islands for static roughness.

Increasing ξ for dynamic substrates sees a peak in mean island size (Fig. 3.2). At higher values of ξ , the substrate prevents even large islands having stability on the surface, whilst lower values of ξ are unable to motivate islands to move or morph in any way. In the case of maximal disorder, the total carbon-substrate energy is approximately 0 until $\xi = 1.0$, at which point it shows a small linear decrease with ξ . This suggests a minor coupling between the substrate and the carbon islands, but not large scale reordering. Indeed, examining snapshots of the substrate after annealing shows no inclination to reform underneath carbon islands. This demonstrates the effect described is motivated by thermal

energy and substrate disorder, rather than some sort of feedback and reordering.

Simple diffusion and deposition models can be mapped to Ising-like lattice gas models. Lattice gas models have been extensively studied and examined in the context of equilibrium behaviour [19], as well as during growth. In many cases exact results are possible, and by drawing analogue to these models we can gain insight into the model of graphene growth presented here. A simple NN lattice gas would observe similar phenomenological behaviour with increasing temperature to the behaviour of the substrate driven island size enhancement effect [138]. In that case, islands formed at a lower temperatures would be unable to assume optimal morphologies and higher temperatures would lead to an increase in long range ordering. However, the behaviour we describe here is driven solely by the evolution of the dynamically rough substrate.

The introduction of disorder into such a system is often treated in one of two limits, quenched and annealed disorder. In a system with quenched disorder, the disorder is frozen for the duration of the simulation and leads to a breakdown of long range order [1]. This is similar to the behaviour exhibited by simulations with static roughness presented here. A system with annealed disorder allows the disorder to occupy the state which minimises the free energy, and does not break down long range order as quenched disorder would [42]. This has analogue to the dynamic roughness energy in our model—allowing the randomness in the system to evolve leads to an increase in island size.

The model we have presented in this chapter has been selected to represent atomic behaviour in an abstract fashion. This has led to a large parameter space, and introduces the danger of over-fitting. In order to investigate a system without this danger, and to observe if the island size enhancement effect observed here is more universal, in the following chapter we present a minimal random field Ising model (RFIM) with a dynamic field similar. Growth in the RFIM has been widely studied [3, 31], and exact results from analytical work allows for the potential of more in depth comparison. Experimental results concerning growth on surfaces with static defects have been explained in the framework of the RFIM[27, 155], which also helps make it a good candidate for study

In the case of graphene specifically, our analysis of growth on dynamically disordered substrates has focussed on the extreme case of unlimited surface mobility. It is clear from Fig. 3.2 that substantial enhancement of island size can be achieved with lower mobility. If however one interprets T_s literally, i.e as the temperature of the copper substrate, it is clear that achieving enhanced growth requires unphysical high temperatures and low heat transfer between copper and graphene. In addition, our model cannot capture the structural feedback effect observed experimentally for graphene CVD growth on Cu(100), namely nano-faceting to (210) + (100) morphology [141]. Further, detailed atomistic/electronic studies are required to establish the extent of substrate mobility at experimental growth

temperatures, and the effect this has on destabilisation of high energy aggregates. We note that most existing theoretical studies at higher levels of detail have focussed on perfect copper facets [152, 132, 34, 95] or (for Ni substrates) well defined ideal surface steps [35].

We are presently investigating spatial correlation of the surface roughness to examine effects of faceting. This will also allow us to address short range correlations, for example by one Cu site affecting multiple neighbouring C atoms. Some other experimentally observed aspects of graphene growth are not replicated in our simple lattice model. For example, islands formed with dynamic roughness contain vacancies, such as the ones in the larger islands in Fig. 3.2. These are mostly formed for single, or small clusters of unfavourable sites. They propagate through the islands throughout the simulation, being swiftly incorporated or removed through the jagged edges. There has been much investigation into the behaviour of defects in monolayer graphene [105, 6], including their formation and possible healing. This can never truly be interpreted in a lattice model where grains cannot be oriented differently and it is impossible to consider rings of anything other than 6 carbon atoms.

3.5 Conclusion

In this chapter we have presented an abstract lattice MC model for surface monolayer growth. Motivated by the fact that the copper substrate used in CVD is close to its melting point, far from a smooth and regular surface, we have introduced a random roughness energy to each site on the lattice. These sites were then fixed (static roughness), or allowed to exchange energies locally (dynamic roughness) with varying degrees of mobility.

The static roughness inhibits island formation, leading to a fragmented surface. Dynamic roughness, at optimal roughness strength ξ , increases the mobility of graphene islands, substantially enhancing the observed grain size. As has been established in a number of studies, and cogently summarised in a recent review [140], optimal conditions for self-assembly occur when interaction energies between components are delicately balanced by thermal noise. In this regime, aggregates can be restructured by bond-breaking and reformation, preventing the formation of kinetically trapped high energy structures. In our model, dynamic substrate roughness plays the role of thermal noise, allowing structures which would otherwise form irreversibly, to anneal. This mechanism is entirely consistent with the “defect healing” mechanism induced by Cu surface mobility reported in the more detailed simulations of Li et al. [68]. We believe that our ability to capture this effect in a simple lattice-gas model suggests the phenomenon may be quite general, to the understanding of which could have dramatic effects on nanomaterial production. The next steps for this exploration could include looking at a rough substrate in greater detail, perhaps by

including correlation in the substrate energies, or more complex interaction energy calculations. An off-lattice model could also be explored, which would allow investigation into local epitaxial effects through a more realistic substrate interaction.

Chapter 4

The random field Ising model with a dynamic field

4.1 Introduction

In the previous chapter we introduced a lattice MC model of graphene growth, and observed an enhancement of island size through the introduction of a dynamic roughness energy. We noted the similarity in the model we constructed to the random field Ising model (RFIM). In the RFIM, the constant external field of the standard Ising model (IM) is modified with a random field energy on each site [4, 115, 31]. The introduction of this random field results in domain pinning and roughening of domain walls, with spins becoming trapped in local minima in the field and preventing long range order. The model we used had a number of differences to the conventional RFIM: the honeycomb lattice, second and third near neighbour diffusion, modified energetics and the introduction of an annealing phase. In this chapter we present results from MC simulations of the RFIM with a dynamic field analogous to the dynamic roughness energy of the previous chapter, and observe the effect on observable thermodynamic quantities.

The IM is often used as a simple mathematical model of ferromagnetism which, since its inception, has seen extensive analytical and computational exploration as well as direct application in a wide array of fields. The introduction of randomness in a magnetic system models a defective magnet which produces an inconsistent field. The RFIM has been used to investigate pattern growth [116], and to classify bulk properties of heterogeneous materials [122]. Random Ising models are well suited to investigating surface growth, due to the analogue between the field lattice and the growth substrate. The RFIM has been suggested for modelling static defects in surface growth [27, 155], with the random static field sites representing random static defects in the bonding between substrate and growing

layer.

The RFIM can be treated in both quenched and annealed limits. In the quenched limit, the random field is defined at the start of the simulation for each site and frozen for the remainder of the simulation. This is an analogue to the static roughness presented in the previous chapter. In the annealed limit, the random field is allowed to come to thermal equilibrium, with the distribution of spins assigned optimally to reduce the system free energy. It can be analytically shown that the quenched RFIM does not exhibit long range order and does not see a phase transition in two dimensions [1, 87], whereas the annealed RFIM does exhibit long range order [42].

The structure of this chapter is as follows. Firstly, initialising from a randomly assigned spin lattice, we look at the standard RFIM, reproducing results from other MC simulations [115]. With this same initial conditions we then introduce a dynamic field, allowing the field to exchange values with neighbours, whilst retaining spin flip dynamics. We find that the island enhancement effect described in the previous chapter also increases domain size in the RFIM, and also observe a substrate ordering effect at lower temperature. We then introduce exchange moves in the spin lattice, and observe that the effect this has on the magnetisation of the system is minimal. Finally we look at magnetic domain reversal [93], demonstrating that the introduction of a random field decreases the domain reversal time, but that the dynamic RFIM has similar behaviour to the standard RFIM. We discuss the applicability of such models to surface growth.

4.2 Methods

The RFIM has Hamiltonian:

$$H = -J \sum_{\langle i,j \rangle} s_i s_j + \sum_i \xi_i s_i + H_{\text{ext}} \sum_i s_i \quad (4.1)$$

where s_i is the spin at site i , J is the coupling constant, ξ_i the random field at site i and H_{ext} is a constant external magnetic field. Spins, as in the IM, can take the value of ± 1 . The first term, referred to as the exchange term, encourages neighbouring spins to align with each other. The second term, the pinning term, encourages spins to align with the value of the random field at that site. Finally the third term encourages the spin to align with an externally applied magnetic field. The random field at site i , ξ_i , is drawn from a uniform distribution of width ξ , centred about 0. The RFIM has been simulated with normal distributions [88] or bimodal distributions [3], but the exact choice of this distribution has no effect on the results due to universality [89].

We perform simulations on a square lattice of size 240×240 using the Metropo-

lis algorithm [82] for MC simulation, as in the previous chapter, and discussed further in Chapter 2. In addition to spin flip dynamics, we perform near neighbour Kawasaki style spin exchange moves [57] in the spin and field lattices, with probabilities p_D and p_S respectively. The probability of selecting a spin flip move is then $1 - p_D - p_S$. Note that if one views spin flip moves in the context of surface growth as deposition or adding particles to a system, the introduction of Kawasaki moves relative to spin flip moves will alter the growth rate of the system. We discuss the impact of this on system times as appropriate later in this chapter. Initially however, we define all time scales as relative to the total number of sweeps performed. We define a single sweep as performing a number of moves equal to the system size.

All simulations are performed at fixed temperature $T = 0.5$. Unlike the previous chapter, which operated in a low temperature regime with very low spin flip (or carbon atom insertion) probability, this temperature choice allows moves which result in an increase in system energy to take place. The temperature selected in the previous chapter was well suited to a model of irreversible insertion and aggregation. The selection of a higher temperature in this model allows for direct validation against experimental data from simulations of the RFIM with a static field [115].

As in the previous chapter, we introduce a temperature governing spin exchange moves in the random field, T_S . This additional temperature parameter affects the Metropolis probabilities generated for spin exchange moves proposed in the field lattice. This had direct comparison with nearly molten [129] and liquid substrates [39] in the graphene growth model, and in a model of magnetisation we tentatively suggest this decoupling could be induced by independently heating a defective magnet to provide the random field.

We perform simulations in two different growth regimes: random initialisation with no external field and magnetisation reversal. For the former, the spins are initially assigned $s_i = +1$ or $s_i = -1$ with equal probability, and the external magnetic field $H_{\text{ext}} = 0$. For magnetisation reversal the spins are initially all set to $s_i = +1$, with $H_{\text{ext}} = -2.0$. The large negative choice of H_{ext} leads to a most favourable state with all spins $s_i = -1$. All simulations have periodic boundary conditions.

The code for this model was written in parallel in C with MPI, using the Synchronous Sublattice (SL) algorithm. This algorithm divides the full lattice into domains, and executes MC moves on subdomains of those domains at the same time, allowing for faster simulations, and for the investigation of larger lattices. For full details of the implementation of the SL algorithm with and without atomic diffusion and a wider discussion of parallel lattice simulations, see Chapter 6

In order to characterise the system we define the following parameters. The mag-

netisation

$$M(\xi, t) = \left\langle \frac{1}{L^2} \sum_i s_i \right\rangle, \quad (4.2)$$

and absolute magnetisation

$$M_a(\xi, t) = \left\langle \frac{1}{L^2} \left| \sum_i s_i \right| \right\rangle, \quad (4.3)$$

compute the average magnetisation and absolute magnetisation respectively. The absolute magnetisation provides more information about simulations performed in the quenching regime—neither state is favoured and we are more interested in the transition to the ordered state. The magnetisation provides more information for magnetisation reversal where there is a preferred system state at the end of simulation. To compare simulations in different parameter selections, we define the final absolute magnetisation

$$M_a(\xi, t_{\max}) = M_f(\xi), \quad (4.4)$$

which is a measure of how ordered the spin lattice is at the end of the simulation.

We measure the domain size $S(\xi, t)$ using a recursive cluster detection algorithm. This algorithm incrementally selects sites on the grid, adds the neighbours to a list, and continues to do so until all sites have been visited. This allows statistics on the domains to be calculated. There are more efficient ways to detect a cluster, for example using labelling algorithms [48], but this recursive approach is sufficient here as we are detecting domains relatively infrequently on a relatively small grid. Using this algorithm we report the final size of the biggest cluster

$$S(\xi, t_{\max}) = S_{\max}(\xi). \quad (4.5)$$

4.3 Results

4.3.1 The random field Ising model

In this section we investigate the standard RFIM in the quenching growth regime at different random field strength ξ . No diffusive moves are included in this initial exploration, i.e. $p_D = p_S = 0$. The $\xi = 0$ case corresponds to the standard IM with no random disorder. The IM is paramagnetic above the Curie temperature T_c . Below T_c the system is ferromagnetic—it is energetically favourable for the spins to align with an externally applied field. Over our simulation timescale, the spins align with each other to form a completely oriented grid, or form metastable domains. In a square grid, the metastable domains are vertical, horizontal or diagonal stripes, where the domain walls on either side of the stripe evolve and shift very slowly over time. An example of time evolution to a completely ordered spin lattice for

the IM is shown in Figure 4.1, with the initially formed domains eroding into a completely ordered spin lattice. Video 4.1 shows typical growth in the IM as outlined here, for more information on which we refer the reader to Appendix A.

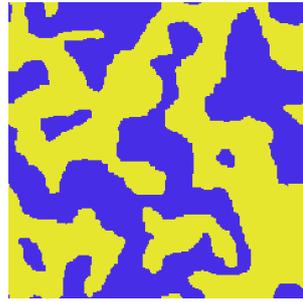
In the RFIM at low values of ξ , the random field slows the formation of domains and eventual ordering, and roughens domain boundaries. At high ξ the domains which initially form have their boundaries pinned in place by the magnetic field, with the pinning term dominating domain evolution. A domain wall is unable to evolve past the random regions of field aligned opposite to the domain wall, and so becomes pinned. An example evolution of the domain pinning at $\xi = 2.0$ is shown in Figure 4.1. Local domains formed in the early stages of growth are pinned in the later stages of growth, with a slowly evolving rough interface. As with the Ising case, Video 4.2 shows typical growth in the RFIM with $\xi = 2.0$. A quenched random field, such as in this case, has been shown analytically to break down long range order [1]. This is in agreement with the results from these simulations.

In Figure 4.2, the time evolution of the absolute magnetisation $M_a(\xi, t)$ is plotted for different ξ . In the standard IM, the absolute magnetisation evolves until reaching a steady state value after roughly 10000 MC sweeps, at which point the system has either fully aligned, or formed striped domains. As ξ increases, the order parameter decreases, resulting in a lack of any ordering at high ξ . The data in this figure is in agreement with similar figures from the literature [115], and indicates our code is a valid implementation of the RFIM. It should be noted that whilst the data from their MC simulations provides a good benchmark for ours, some of the conclusions they draw are incorrect; for example they claim to observe a phase transition in the 2D RFIM, which has been analytically proven to be incorrect [1, 87].

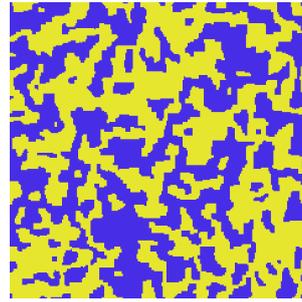
4.3.2 The dynamic random field Ising model

In this section we report results from simulations of the dynamic RFIM with spin exchange moves in the field, $p_D = 0.0$, $p_S = 0.5$. Simulations were performed in the quenching growth regime for different values of field strength ξ and decoupled temperature T_S . This introduction of a dynamic field gives a direct mechanism for coupling and feedback between the field and spin lattices, and we observe two behaviour regimes.

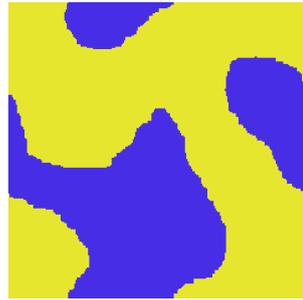
At low T_S , the field orders beneath domains formed in the early stages of growth. This increased the pinning effect; with stable domains in the field beneath the spin grid, the system is trapped in this local minima. At high T_S , the constantly shifting field remains dynamic beneath the spin lattice, and does not order beneath domains. Instead, the spin lattice orders, with the dynamic field providing a mechanism for domains to adjust and realign. Snapshots from simulations at $\xi = 2.0$ for $T_S = 0.5$ and $T_S = 10.0$ are shown in Figure 4.3, with corresponding Videos 4.3 and 4.4 detailed in Appendix A. In these



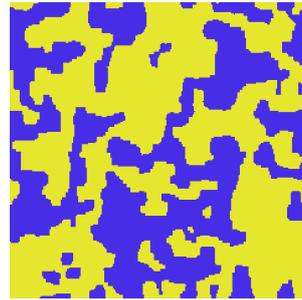
Ising, $t = 100$



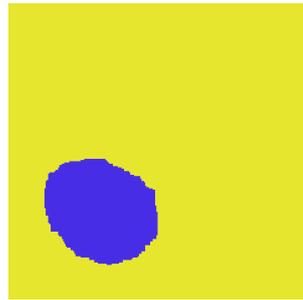
RFIM, $t = 100$



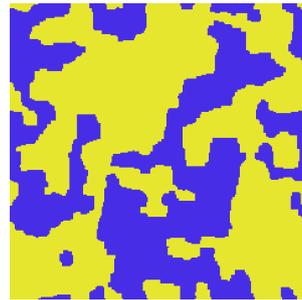
Ising, $t = 1000$



RFIM, $t = 1000$



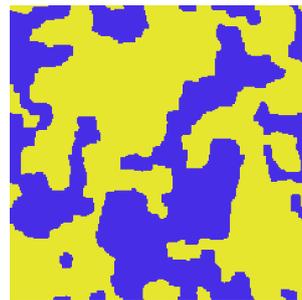
Ising, $t = 5000$



RFIM, $t = 5000$



Ising, $t = 10000$



RFIM, $t = 10000$

Figure 4.1 Snapshots from the evolution of the spin system for the IM and the RFIM with noise $\xi = 2.0$. Spins are represented by coloured pixels, $+1$ are blue and -1 yellow. In the IM, domains form at early time and evolve into a single spin or striped domains (not pictured). The snapshots from the RFIM show domain pinning—the domains formed early in the simulations are pinned by the random field and are thus prevented from growing further.

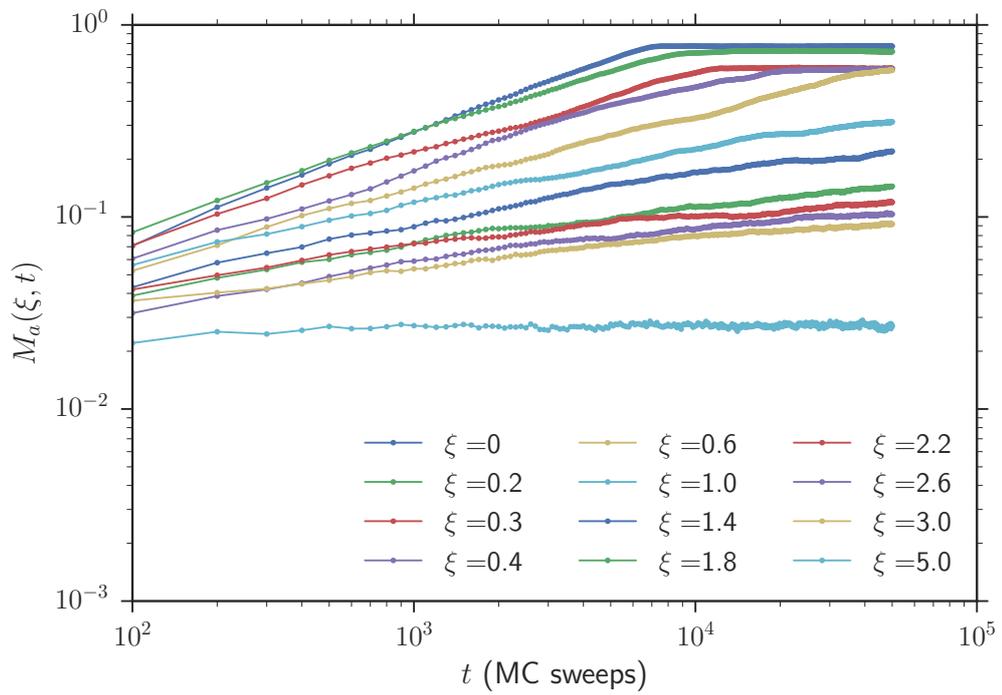


Figure 4.2 The average absolute magnetisation, $M_a(\xi, t)$ against time for different random field strength ξ . Increasing ξ leads to an increase in the domain pinning, and a reduction in $M_a(\xi, t)$. Each point is an average of 20 simulations, and errors are on the order of the point size.

snapshots and videos the field lattice is shown alongside the spin lattice, and the ordering of the field lattice in the $T_S = 0.5$ case is clear.

The time evolution of the absolute magnetisation for selected substrate temperature across a range of ξ values is shown in Figure 4.4. At low T_S the field ordering effect is visible at high ξ . Trajectories never increased over the entire simulation, which is in contrast to the equivalent figure for the RFIM, Figure 4.2. There is a transition between the pinning through field ordering and spin alignment, and at low values of ξ the spin lattice aligns as in the RFIM. At high T_S the dynamic field causes the spin lattice to align for all values of ξ , which results in trajectories for all ξ which are steeper than the equivalent trajectory for the RFIM.

The final magnetisation $M_f(\xi)$ for this same selection of T_S is shown in Figure 4.5 with the final magnetisation for the RFIM also shown. Here the two new effects are demonstrated. Simulations performed with low T_S have final magnetisations lower than the RFIM due to the added stability provided by the ordering of the random field. However at high T_S , final magnetisations higher than the RFIM are observed as the dynamic field unpins domain boundaries. This effect is also demonstrated in the largest domain, $S_{\max}(\xi)$, shown in Figure 4.6, with simulations performed at low T_S giving smaller domains than the RFIM and high T_S larger.

As we remarked upon earlier in this chapter, the definition of time as number of MC sweeps regardless of move selection needs to be considered. Trajectories in Figure 4.4 do not all reach a steady state. The time to reach a steady state is not necessarily only dependent upon the number of attempted spin flip moves however, as the number of attempted field exchange moves also affects growth. The selection of total MC sweeps as system time is therefore acceptable for the comparisons we seek to make in this chapter, but we note that an exact comparison could examine equilibrium magnetisation and domain size, or the critical exponents governing growth [115].

4.3.3 Spin exchange in the dynamic random field Ising model

In this section we explore the effect of introducing Kawasaki style spin exchange moves into the spin lattice, for both the RFIM and the dynamic RFIM described in the previous two sections, corresponding to $p_S = 0.0$, $p_D = 0.5$ and $p_S = p_d = 1/3$ respectively. Simulations are performed in the quenching growth regime. We restrict our observations to two values of the decoupled temperature for the dynamic RFIM, $T_S = 0.5$ and $T_S = 5.0$, which gives three cases to investigate:

1. Dynamic RFIM with spin exchange, $T_S = 0.5$,
2. Dynamic RFIM with spin exchange, $T_S = 5.0$,

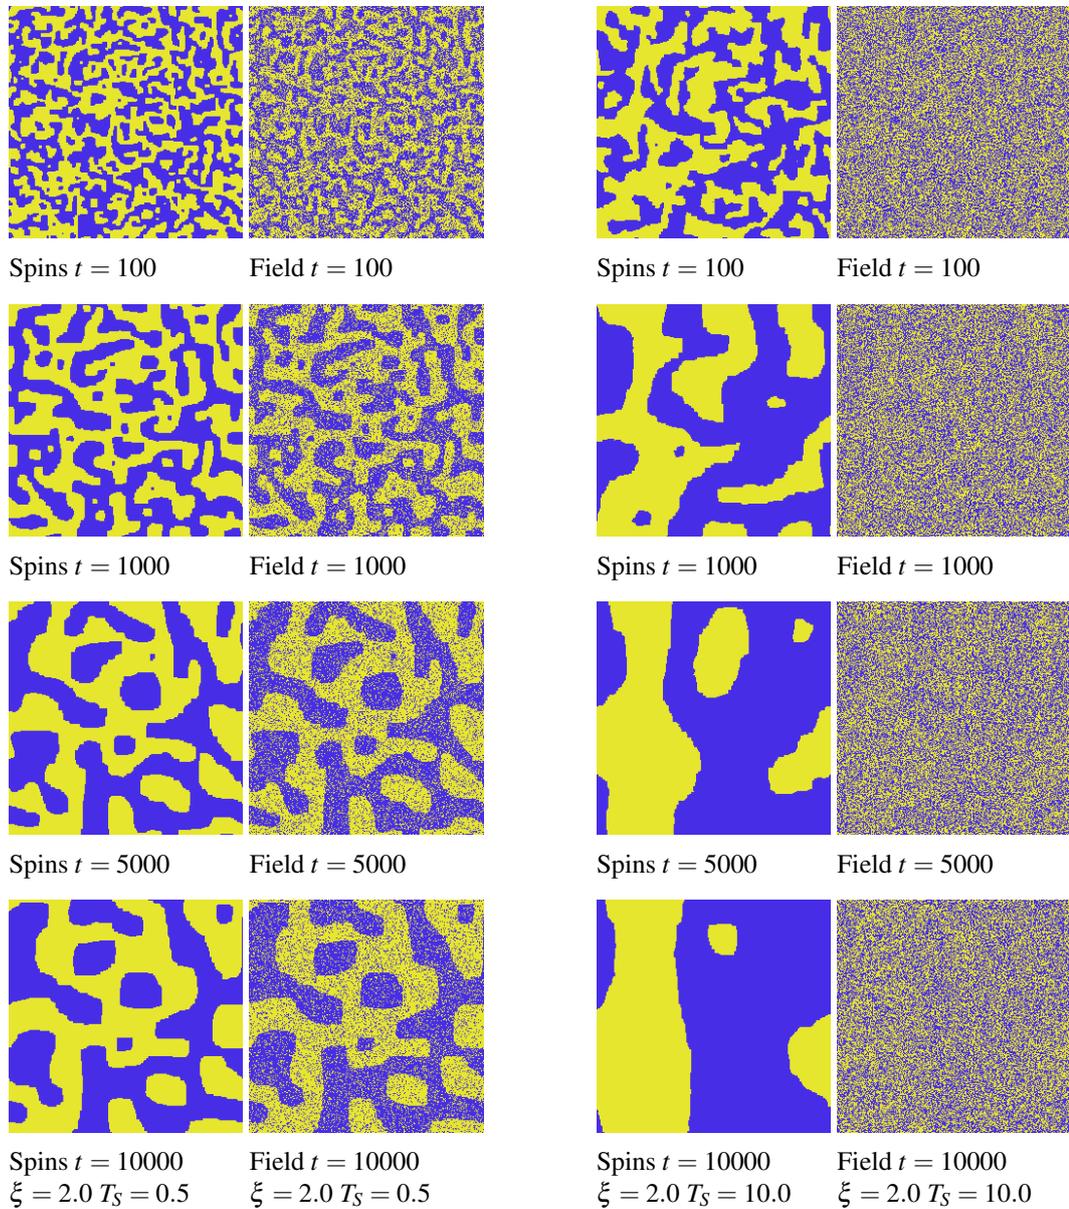


Figure 4.3 Snapshots from the evolution of the spin system for the dynamic RFIM with noise $\xi = 2.0$ and decoupled field temperatures $T_S = 0.5$ (left) and $T_S = 10.0$ (right). Spins are represented by coloured pixels, $+1$ are blue and -1 yellow, whilst field sites with positive field strength are coloured blue and negative field strength coloured yellow. For $T_S = 0.5$, domain walls are smoothed in the spin grid in comparison to the static RFIM, and the random field aligns with the spins. For $T_S = 10.0$ the random field is disordered, but domains in the spin grid are unpinned.

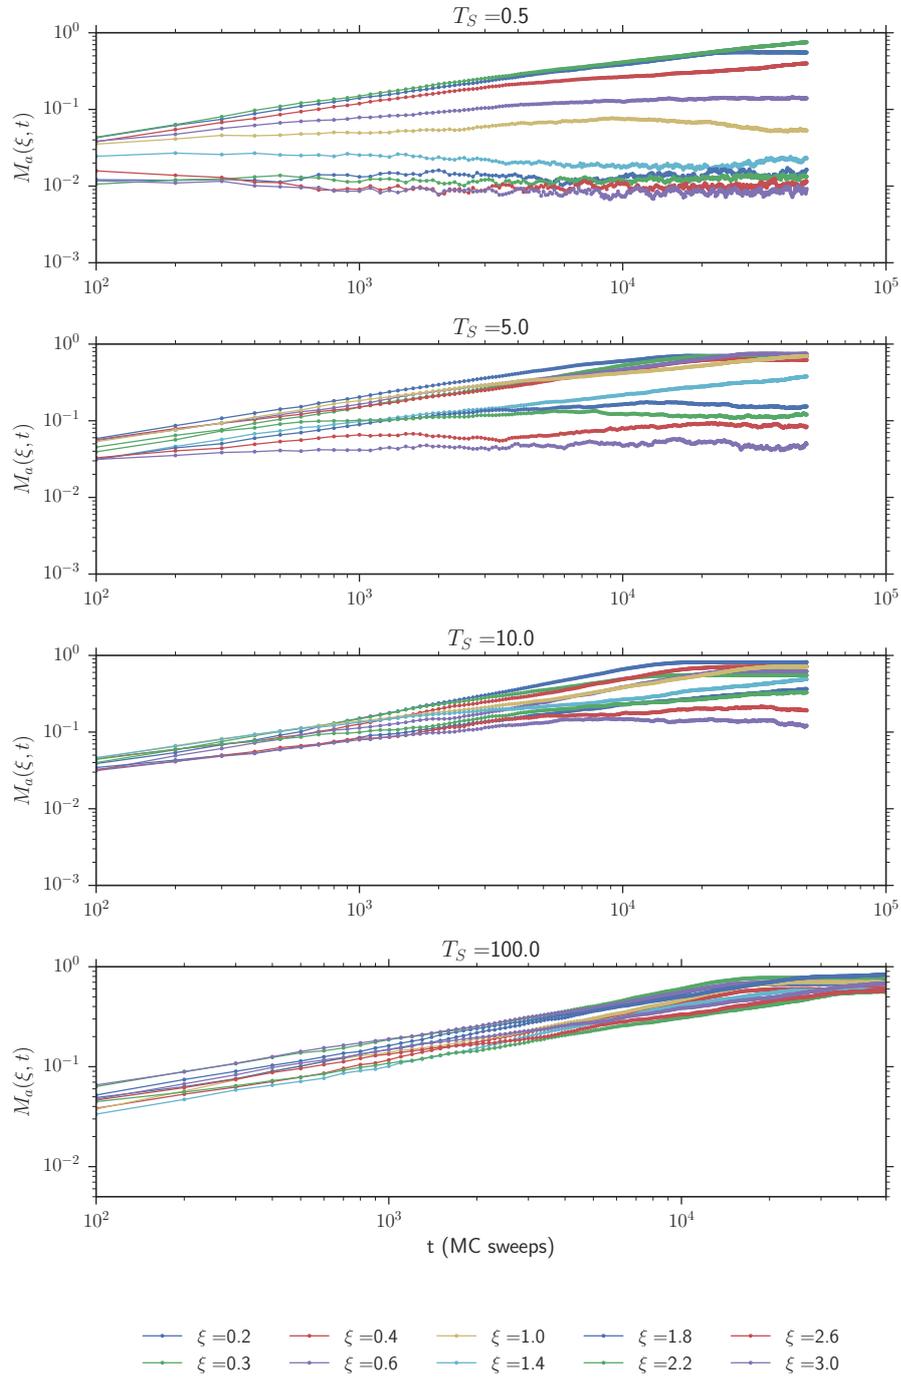


Figure 4.4 The average absolute magnetisation $M_a(\xi, t)$ against time for different random field strength ξ with decoupled field temperatures T_S in the RFIM with dynamic noise. Increasing T_S across the subplots results in an increase in magnetisation. Each point is an average of 20 simulations, and errors are on the order of the point size.

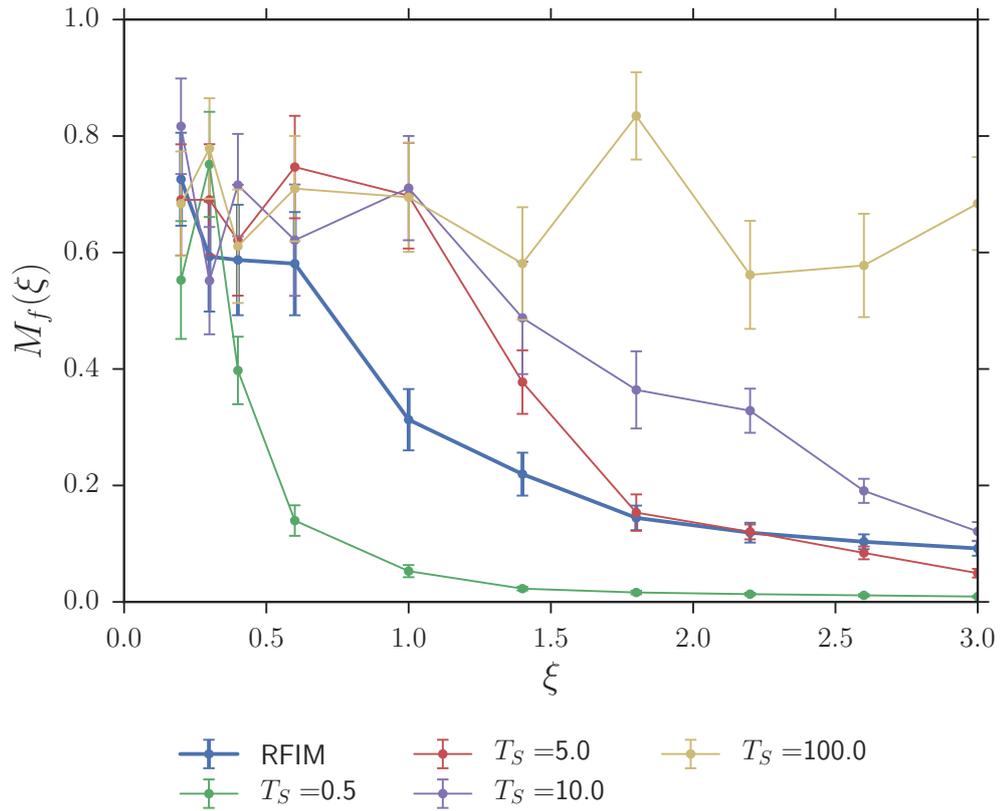


Figure 4.5 The average absolute magnetisation at the end of simulation $M_f(\xi)$ against different noise strength ξ , for different decoupled field temperatures T_S . At low T_S , increasing ξ leads to an increase in the domain pinning over the RFIM whilst higher T_S removes the pinning effect altogether. The RFIM is also plotted for comparison. Each point is an average of 20 simulations, with standard error as shown.

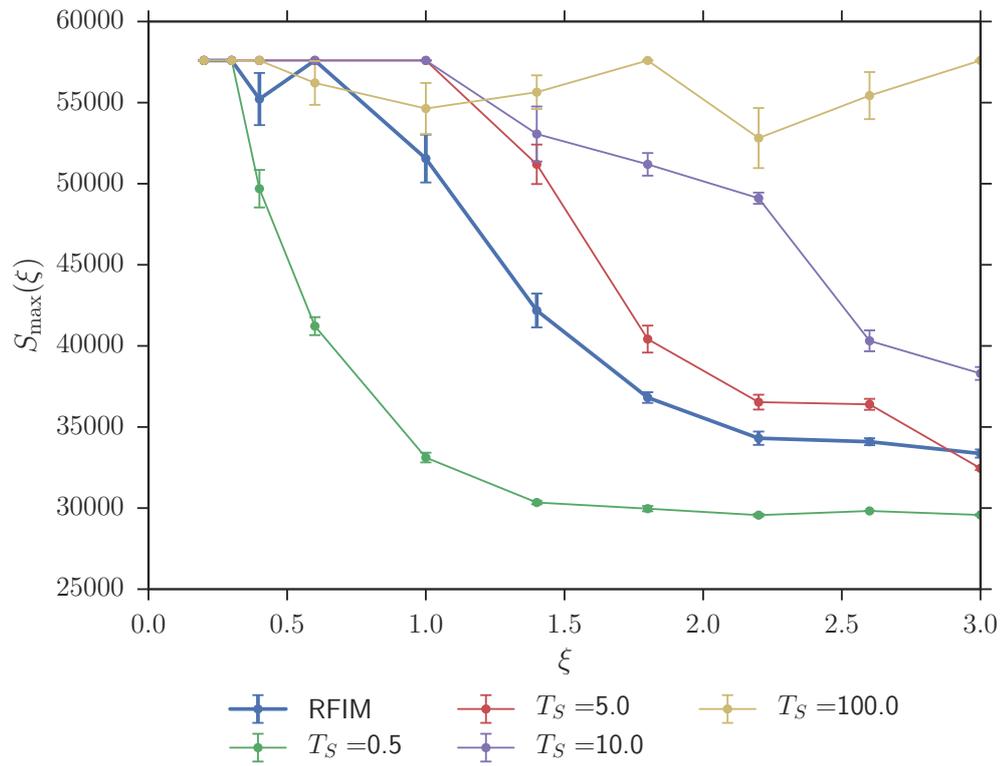


Figure 4.6 The average largest domain size at the end of simulation $S_{\max}(\xi)$ against different noise strength ξ , for different decoupled field temperatures T_S . The RFIM largest domain is also plotted for comparison. Each point is an average of 20 simulations, with standard error as shown.

3. RFIM with spin exchange.

The evolution of the absolute magnetisation $M_a(\xi, t)$ for different ξ is shown in Figure 4.7, for each of these three cases. These graphs can be directly compared to Figures 4.2 and 4.4 for the RFIM and dynamic RFIM respectively.

The final absolute magnetisation $M_f(\xi)$ for the three cases, as well as the equivalent cases in the RFIM and dynamic RFIM without spin exchange, are plotted in Figure 4.8. The final magnetisation in each of the three cases is within standard error of the corresponding case without spin exchange, and so we conclude that the inclusion of spin exchange resulted in no substantial difference in behaviour.

As in the previous section, introducing diffusion moves does prompt an examination of whether total MC sweeps is a good choice for system time. Following the same logic as the previous chapter, we deem it appropriate for the comparison of different parameter regimes discussed in this section, but note that a comprehensive study would examine equilibrium properties or critical exponents.

For phase ordering kinetics with a conserved order parameter, we expect quantitatively different behaviour of critical parameters to that with a non-conserved order parameter [14]. However, we note that in all three cases outlined above we perform spin exchange *in addition* to spin flip moves. As such the order parameter remains non-conserved in all three cases.

4.3.4 Magnetisation reversal

In this section we investigate growth in the magnetisation reversal growth regime for the RFIM and dynamic RFIM. We report results from simulations with random field strength $\xi = 0$, $\xi = 0.5$ and $\xi = 2.0$, with decoupled field temperature $T_S = 0.5$ and $T_S = 5.0$. In each individual trajectory of these simulations we observed a time delay for a cluster of -1 spins to nucleate on the spin lattice of $+1$ spins. After this cluster nucleated we saw a quick reversal of the remainder of the grid to -1 in all simulations. The length of the initial delay for a stable cluster to nucleate however, varied across the simulations. The total number of MC sweeps is no longer a good choice for system time because a certain number of insertion moves *must* occur in order for the magnetisation to reverse, and this is observed to happen quickly. When we introduce field exchange moves in the dynamic RFIM, we therefore consider the time to be the total number of attempted spin flip moves.

In Figure 4.9 we show the average magnetisation for the IM and RFIM. Increased ξ results in a decrease in magnetisation reversal time. Indeed, at $\xi = 2.0$ the reversal is almost immediate. We attribute this decrease in reversal time to a decrease in the time taken to nucleate a stable cluster of -1 spins. The random field creates regions of the

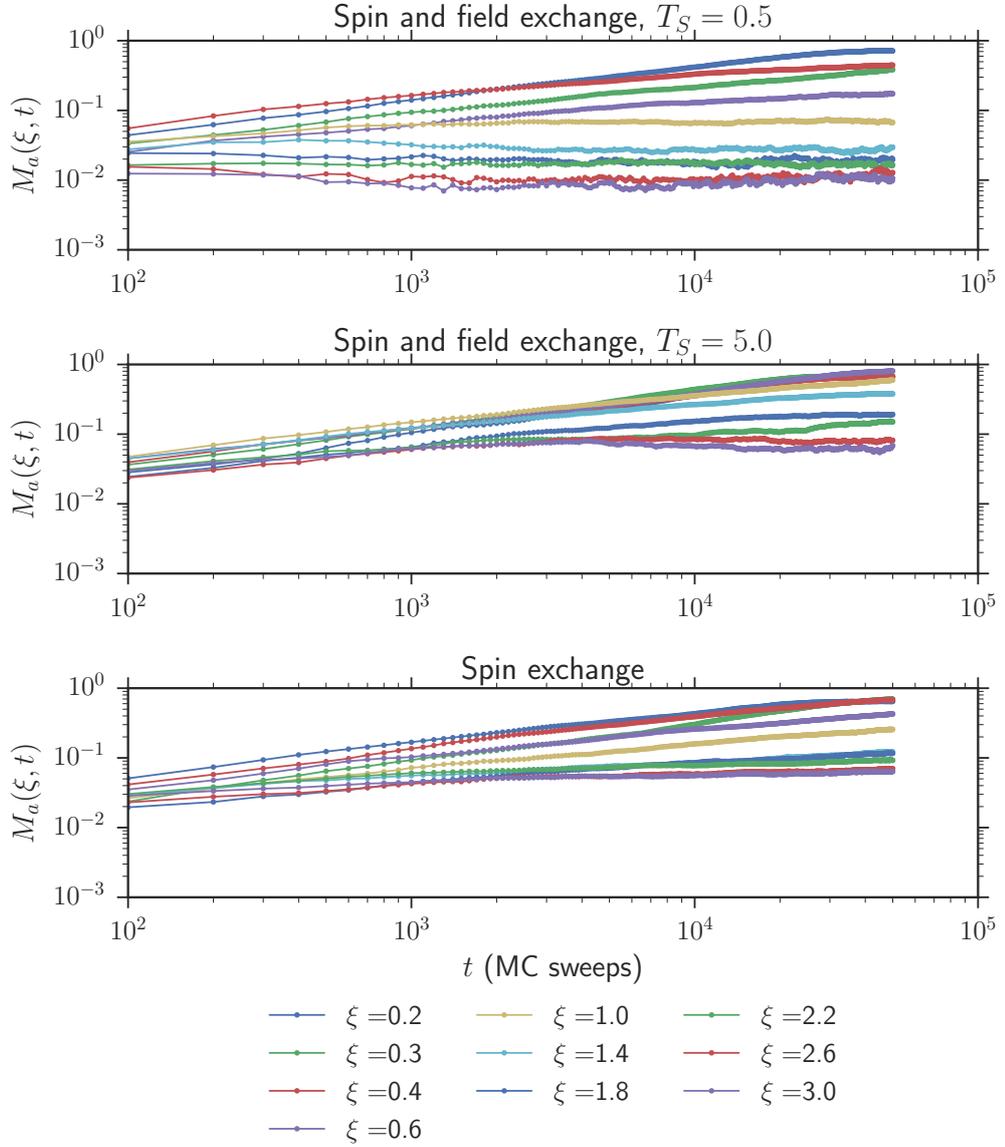


Figure 4.7 The average absolute magnetisation at the end of simulation $M_f(\xi)$ against different noise strength ξ , for the dynamic RFIM with spin exchange at two decoupled temperatures $T_S = 0.5$ and $T_S = 5.0$, and RFIM with spin exchange. These trajectories can be directly compared to Figure 4.2 and Figure 4.4. Each point is an average of 20 simulations, with errors on the order of magnitude of the point size.

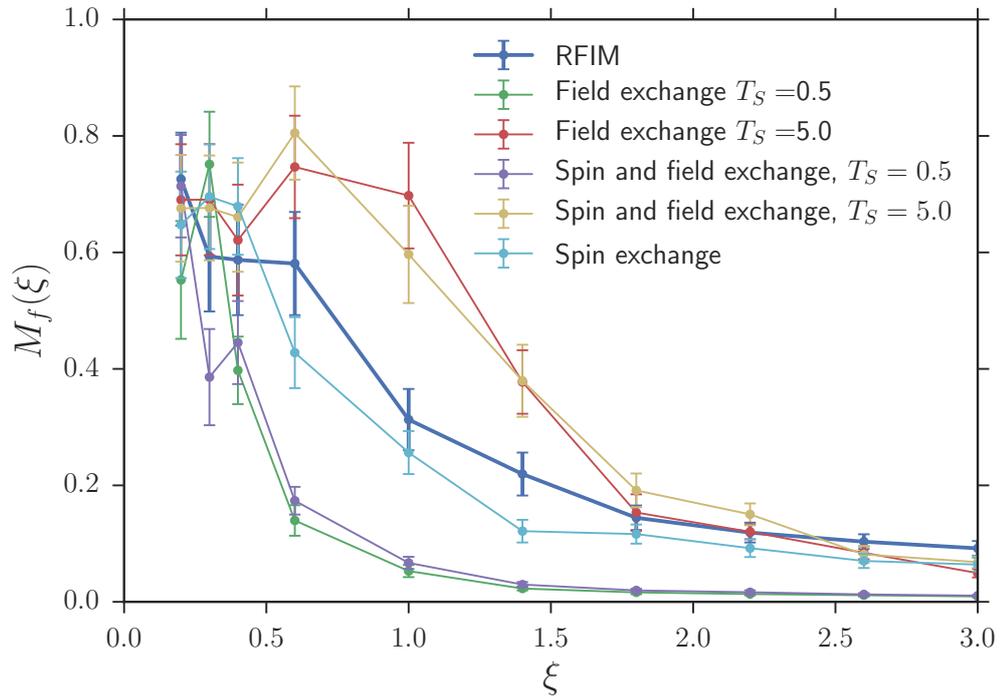


Figure 4.8 The average absolute magnetisation, at the end of simulation $M_f(\xi)$ against different noise strength ξ , for the dynamic RFIM with spin exchange at two decoupled temperatures $T_S = 0.5$ and $T_S = 5.0$, and RFIM with and without spin exchange. Increasing ξ leads to an increase in the domain pinning for the RFIM at low T_S , whilst higher T_S removes the pinning effect altogether. Note that the inclusion of spin exchange results in no significant change to the final magnetisation in each of the three cases. Each point is an average of 20 simulations, with standard error as shown.

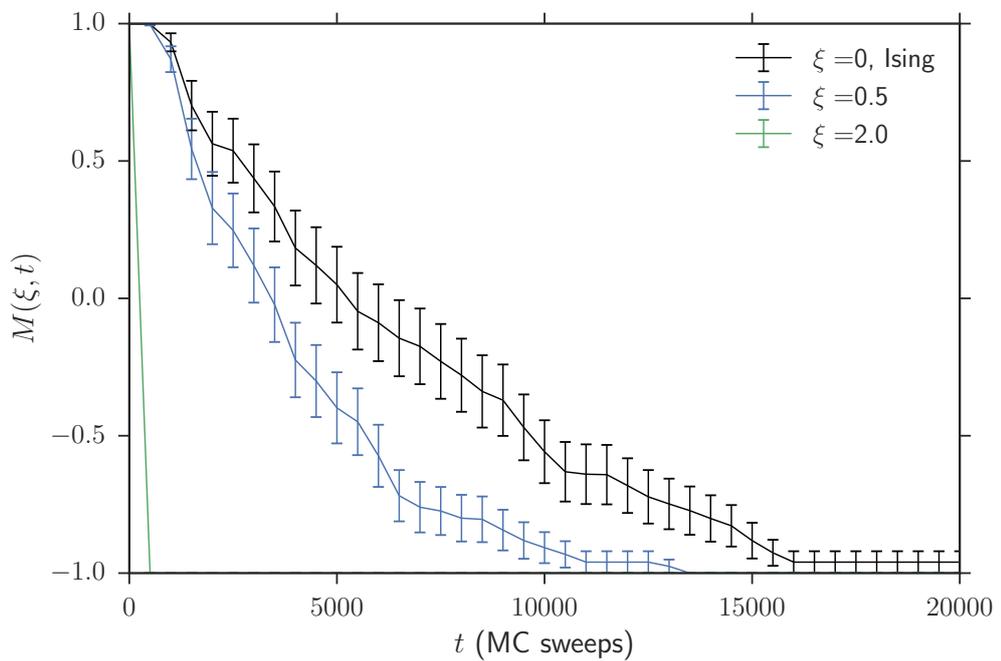


Figure 4.9 The time evolution of the average magnetisation for simulations of magnetisation reversal in the RFIM. Three cases are shown: $\xi = 0$, $\xi = 0.5$ and $\xi = 2.0$. The $\xi = 0$ case corresponds to the standard IM. Larger ξ reduces the magnetisation reversal time. Points are the average of 50 simulations.

field which are more favourable for nucleation, and creates a lower barrier to initial cluster formation.

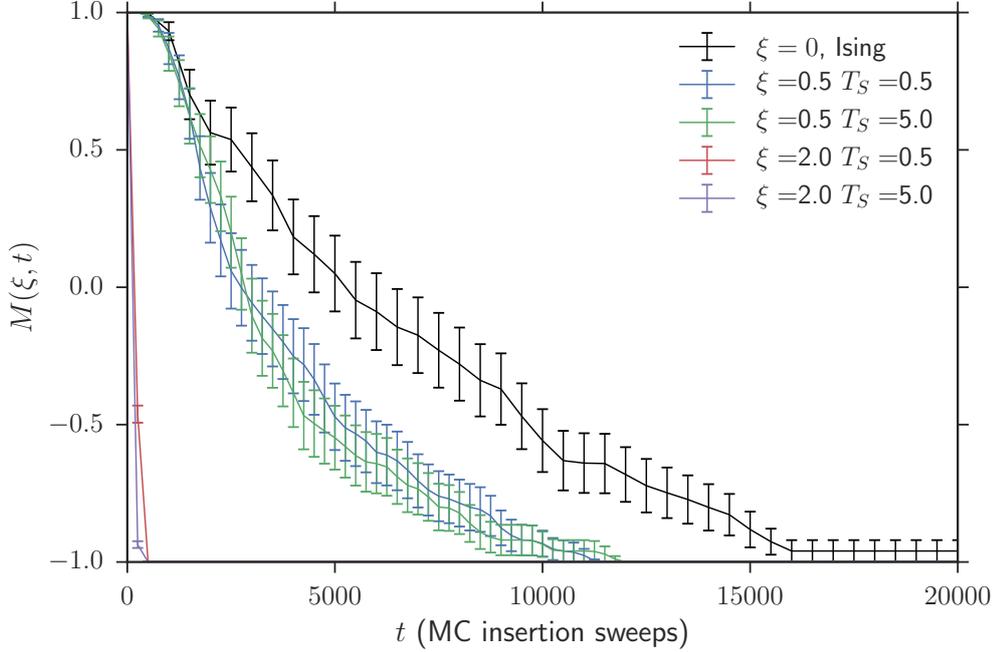


Figure 4.10 The time evolution of the average magnetisation for simulations of magnetisation reversal in the dynamic RFIM. Cases are shown for random field strength $\xi = 0$, $\xi = 0.5$ and $\xi = 2.0$ and decoupled field temperature $T_S = 0.5$ and $T_S = 5.0$. The $\xi=0$ case corresponds to the standard IM. The time in this plot is rescaled to be the number of insertion moves attempted, rather than the total number of MC moves. Points are the average of 50 simulations.

In Figure 4.10 we show the average magnetisation for the dynamic RFIM, with a timescale defined as the number of attempted insertion moves. We observe no change in behaviour with the change from RFIM to dynamic RFIM, with the evolution of average magnetisation for the dynamic RFIM falling within standard error of the RFIM.

4.4 Conclusion

In this chapter we have generalised our observations of island size enhancement in simulations of graphene growth to a standard formulation of the RFIM. We explored the parameter space of the RFIM, and reproduced previously published information about the magnetisation [115]. We then introduced field exchange with a decoupled temperature in an exploration of the dynamic RFIM. We observed that at low decoupled temperature T_S , the field exchange moves ordered the field underneath the domains, leading to a decrease in order parameter, but an increase in domain smoothness. At high T_S we observed the do-

mains pinned by the RFIM were unpinned, leading to the system becoming fully ordered. Whilst we did not provide a full description of the phase diagram in terms of T , ξ and T_S , we did provide a detailed study of the effects of the dynamic field at particularly interesting parameter choices. We then added Kawasaki style spin exchange dynamics into the model, and observed no substantial behavioural change.

Finally, we changed the growth regime under investigation to magnetisation reversal, using an external magnetic field to drive reversal of a fully ordered lattice to an opposite ordered lattice. This explored a different type of growth, with a time delay before nucleation, followed by a quick domain reversal. We observed that the introduction of a random field to this model reduced the time for reversal, due to a decrease in nucleation time which we attributed to an increase in local minima for islands to nucleate on.

The quenched and annealed limits of the RFIM have been studied extensively, and in some cases have exact analytical results [1, 87, 14]. In the quenched limit, the RFIM undergoes no phase transition and does not exhibit long range ordering. This is equivalent to the static RFIM initially introduced in this chapter, where our simulations showed a breakdown in long range ordering. In the annealed limit however, the RFIM does exhibit long range ordering. The principal difference between the simulation results presented here and the annealed RFIM extensively studied is the decoupled temperature of the spin and field lattices. We observe a behaviour dependence upon this decoupled substrate temperature T_S , observing long range order at high T_S and an absence of order at low T_S . This suggests the behaviour of the dynamic RFIM is similar to that of the annealed RFIM.

In the graphene growth model there were a number of notable differences to the model described in this chapter. The graphene model used a modified energetics scheme which favoured C–C bonds and heavily reduced the strength of H–H bonds, hexagonal geometry, atomic diffusion up to third near neighbours and a low temperature regime which enforced irreversible adatom attachment to the substrate. Including a random field in this model led to enhancement of island size, with a predicted critical roughness at which the island size enhancement was optimised.

In this chapter we saw an unpinning of domain boundaries through the introduction of a dynamic random field, which suggests that the observations in the previous chapter are a more general property of including a dynamic random field, rather than an artefact of the specific energetic regime and geometry. However, we also saw that at low decoupled substrate temperature T_S , the field lattice ordered under the spin domains, which led to formation of stable domains with smooth domain boundaries. In the graphene growth model the smooth domain boundaries were not detected due to the long range carbon diffusion moves already providing a mechanism for island morphology change. The spatial extent of the stable domains was lower in this chapter, but the constant low flux of monomers in

the graphene model, combined with high diffusion rates, made this harder to detect in the previous chapter. The critical roughness energy we discussed in the previous chapter is not as clearly defined in this chapter as here we investigated the transition to full rather than partial coverages.

The concept of using the RFIM as a model for surface growth on substrates with static defects has been previously proposed [27, 155], and the results presented in this chapter do support this idea. In simulations of magnetisation reversal in the static RFIM we demonstrated that increasing the random field strength decreased the nucleation time, and time to reversal. This could be incorporated into a more detailed model of graphene growth investigating nucleation on defects or steps [137, 135], where an increased random field would correspond to a more defective substrate. The introduction of a dynamic rather than a static field was shown to have no effect on magnetisation reversal. Combining observations from substrate temperature variation leading to island size enhancement, and random field variation increasing nucleation rate, a hybrid model could be constructed which could provide a more realistic 2D growth model.

The strength in using the RFIM to model surface growth is in the creation of two lattices for the substrate and the growing layer with a defined feedback interaction between them. Events in each layer can be decoupled in order to investigate different material properties. This concept allows the geometry of the growing layer to be strictly defined based on the geometry of that layer rather than on favourable substrate sites. Refinement of this method could make it applicable to a wide range of experimental systems.

Chapter 5

Off-lattice molecular dynamics simulations of graphene growth on an effective model substrate

5.1 Introduction

In Chapters 3 and 4 we presented results from lattice-based MC models of growth. Lattice MC modelling has its advantages in simplicity and speed, and we introduced a model for surface roughness, demonstrating that a possible route to island enhancement was through a dynamically rough substrate. In Chapter 3 we selected a honeycomb lattice to represent the atomic positions of carbon in the growing graphene layer, rather than that of the substrate. In reality the substrate layer will have a geometry that is not identical to the growing graphene layer, but may well have some matching symmetry. For example Cu(111) is a triangular lattice with good lattice matching to a graphene lattice, leading to the formation of graphene islands with a single orientation, and a graphene layer in registry with the substrate [145]. In this chapter we introduce a molecular dynamics (MD) model of graphene growth that is still strictly 2D, but which uses a realistic C–C bond potential [153] on an effective substrate potential defined from Cu positions in different crystallographic facets. This allows a direct investigation of the effects of the substrate geometry on the growth rate and orientation of growing graphene islands.

Graphene growth via CVD has been hailed as a promising route to synthesis due to the availability of cheap copper foils and relatively low cost of production [110]. These cheap copper foils are typically polycrystalline with a relatively high proportion of the copper substrate as Cu(100) [145, 141, 49]. Despite weak epitaxy with the copper surface [38], graphene islands grown on Cu facets predominantly grow with orientation determined

by the geometry of the substrate [141, 145]. Different orientations of graphene islands have been observed on the different facets of copper, with graphene on Cu(111) reported to be mostly aligned with the substrate, or rotated by 7° [37]. On Cu(100) graphene is mostly reported to grow in registry with the substrate [143], but also to form four lobed islands resulting from graphene islands growing with different orientations along the four fold symmetry of Cu(100) [143, 71]. Growth of graphene on Cu(111) has been reported as diffusion limited [90], whilst growth on Cu(100) as attachment limited [143]. Different growth rates have been observed on different facets [145], with observations from *in situ* CVD of graphene on Pt foils reporting a gap of approximately 2500 s between the formation of the first island on different facets [137]. Explanations for the different growth rates are attributed to grain orientation dependent precursor dissociation rates, nucleation barriers, and are likely to be affected by the density of nucleation sites such as step edges in the particular region.

Modelling of these growth processes tends to focus on early stages of growth and has identified the relative stability of carbon dimers as opposed to monomers, as well as reporting similar or higher rates of dimer diffusion [36, 147]. Graphene is reported to nucleate on copper either as small, domed islands or on step edges or terraces in the substrate [20]. These results can be integrated into kMC and mean field models [147], but incorporating realistic potentials for growing large carbon structures on a copper foil remains out of reach of most DFT methods. Whilst not all complications of real graphene growth can be incorporated into a strictly two dimensional model, in this chapter we explore the effects of substrate geometry on realistic growing graphene structures, and offer insight into the experimental findings.

5.2 Methods

This model of growth is a strictly 2D MD simulation, made using the popular molecular dynamics package LAMMPS [99]. The extensive and powerful LAMMPS code implements fast spatial decomposition algorithms in MD simulations. It includes default potential functions that allow for customisation, leading to simulation of a wide range of physical systems. These pre-written routines gave us a good base from which we construct a fast, well implemented MD code, without worrying about the computational challenges involved in this implementation. We also gain the ability to customise thermostats and inter-atomic potentials against an array of empirically fitted examples [97, 94] to suit the abstract model we want to explore.

As well as the advantages in speed and pre-written code, there are some downsides to using a pre-packaged simulation code such as LAMMPS. When you are not the author

of the code, there is a danger of being unaware of some nuance of the implementation. It is easy to use a potential which has not been optimised for the parameter space you are using for example, which would skew your results accordingly. The scope of the original code could be wider or narrower than your application of it, introducing complications which could otherwise be avoided. For example the model described in this chapter is entirely 2D, but potentials are generally fitted for 3D. As graphene is a well studied 2D material, we were able to find a suitable potential that considered 2D scenarios, but this could not always be the case.

In order to enforce a 2D simulation, the simulation box is defined to be very narrow in the z dimension. The z components of atomic velocities and forces are then zeroed at the end of every timestep, ensuring all atoms remain in the 2D plane. The two atomic species used, copper and carbon, reside inside this plane. The carbon atoms are free to diffuse in the plane, attach to each other and form islands with no pre-determined orientation, according to C–C potential. The copper atoms, arranged according to the geometry of crystallographic facets, remain static during the simulation, and have no interaction defined between them. Copper–carbon interactions are then defined through an effective potential, allowing us to model the substrate geometry.

It was important in constructing this model to select an appropriate C–C potential. As we outlined in Chapter 2, potentials vary in complexity, parametrisation and computational cost. We had the additional complication of needing a potential which not only models growth, but is also optimised for two dimensional structures. There was a danger of selecting a potential which was optimised for neither. Pair potentials, such as Lennard-Jones [54], give a generic description of inter-atomic forces and are generally well suited to large scale simulations which prize computational efficiency. Empirical many-body potentials such as the Tersoff [119] and Brenner [15] potentials are parametrised by fitting to either a set of experimental measurements or to quantum mechanical calculations, and are used when many body interactions are important. These potentials often fail to capture systems outside the exact ones for which they were parametrised, which led to the development of potentials derived directly from quantum mechanical principles [85].

The reactive bond order potential (REBO) [16], and its extensions [94] are good examples of empirically fitted many-body potentials suitable for modelling carbon structures. These potentials have proved successful for studying thermodynamic properties of bulk [30] and nano-scale [84] structures. Such empirically defined potentials have not seen as wide an application in growth modelling, as the wide array of states sampled during growth are unlikely to have been tested during potential development. In order to effectively model growth of carbon structures, Zhou et. al. [153] recently derived an analytic version of the empirical bond order potential (BOP) [97], parametrised towards growing

carbon structures. They report significant advantages over the REBO style potentials in capturing crystalline growth characteristics of graphite and graphene structures, satisfying our condition for a potential optimised for 2D growth. As this potential has been distributed within LAMMPS, it is an excellent choice for the abstract modelling of graphene growth we wish to investigate. This potential also captures defect energy trends from DFT.

Generally in MD simulations of CVD on copper [153], the bulk metal is simulated atomistically in the MD code, with copper atoms free to move and interact with other atoms through inter-atomic potentials. In these simulations we keep the copper atoms static, disable all Cu-Cu interactions and describe C-Cu interactions through Gaussian wells. This abstract model does not try to represent the real atomic structures of low-index Cu surfaces. A single potential well per Cu surface unit mesh represents some optimum adsorption position in each unit mesh and does reflect the overall symmetries and unit mesh sizes of the different low index faces. The well energies are given by

$$E(r) = -D \exp(-Wr^2), \quad (5.1)$$

where W is the inverse width of the well, and D the depth.

We arrange the copper atoms in the 2D plane according to four crystallographic facets, Cu(100), Cu(111), Cu(110), Cu(210) with lattice spacing given in Table 5.1, generating the effective surface potentials shown in Figure 5.1. Cheaply available copper foils with a Cu(100) structure, and the close resemblance of the geometry of Cu(111) to graphene makes these facets obvious choices for study. We selected Cu(110) and Cu(210) to compare island orientations to observations of faceting of Cu(100) to Cu(n10) observed in CVD graphene growth [141]. Here the mismatch epitaxy was demonstrated on single crystal (110), with two orientations observed from two coincident LEED spots. Graphene growth via CVD has been reported with different growth rates and orientation on polycrystalline foils [141] and DFT simulations on polycrystalline foils [83] have indicated the possibility of facet insensitive growth. This potential selection gives us a high degree of control over the form of the C-Cu interactions, and will allow us to investigate these findings.

The well width W was selected such that each well does not interact with that of neighbouring copper atoms, creating an eggbox-like potential of isolated Gaussian wells. Due to the two-dimensional nature of this study, the well depth D does not have an exact experimental counterpart, but can be thought of as a diffusion barrier to monomers—a monomer located in an isolated well has to overcome the barrier in order to escape the well and diffuse on the potential. Following this interpretation of D we chose to parametrise it from estimates of the monomer diffusion barriers between octahedral subsurface sites on Cu(111)—calculated to be 0.55 eV in one DFT study [20], and 0.5 eV from another—rather

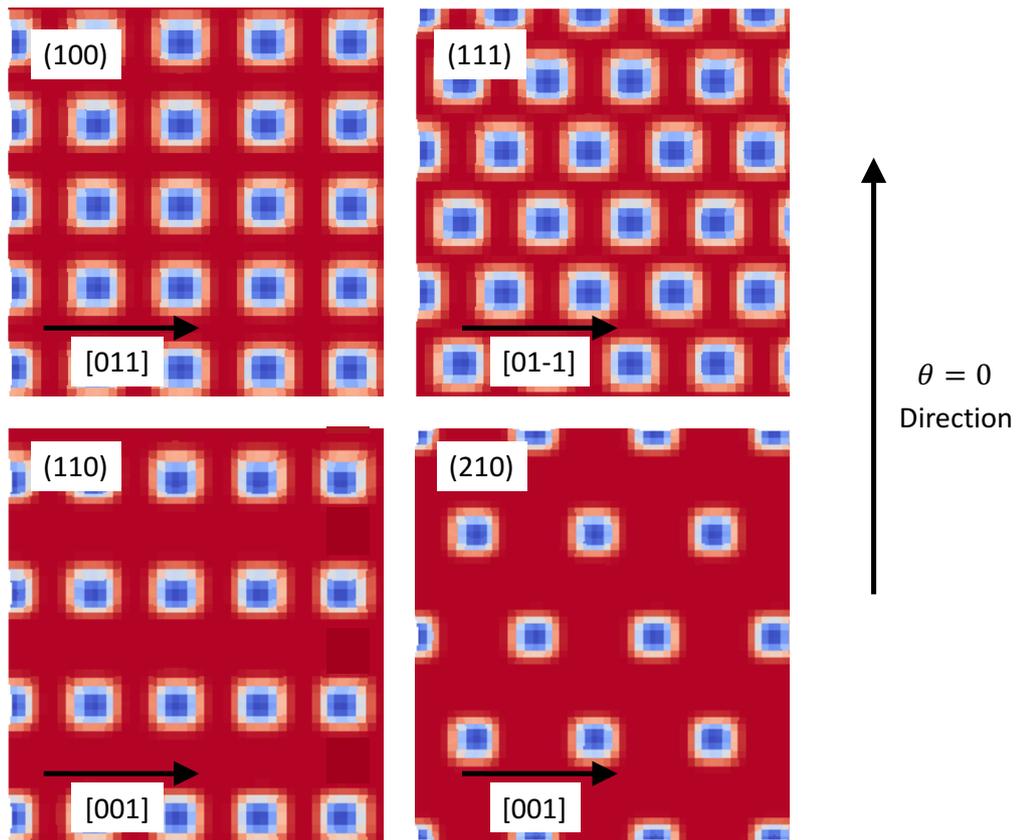


Figure 5.1 The four effective substrate potentials representing crystallographic facets of a copper lattice used in simulations presented in this chapter. Each image is a colour map of potential across the plane, with blue areas indicating most negative potential and the labels indicating crystallographic direction. This is not an image of the full plane, but a snapshot of the surface to show the geometry. The arrow indicates the direction of 0 degree orientation on all surfaces, used in the orientation calculations presented below. Horizontal crystallographic directions are marked on each substrate.

than from carbon adsorption energies on copper. Based on this argument, we chose a value of $D = 0.5 \text{ eV}$.

The lattice parameter taken was initially that of copper at room temperature, and is shown in Table 5.1. As simulations used periodic boundary conditions and the unit cell for each of these facets was different, the exact size of the box for each facet was slightly different, but sizes were chosen to give as similar to each other as possible, and are shown in Table 5.1.

Parameter	Description	Value
D	Copper well depth	0.5 eV
H	Copper well width	1.0 Å
a	Copper lattice parameter (300 K)	3.610 Å
T	Temperature	3000 K
t	timestep	1 fs
t_{\max}	runtime/number of steps	3.5×10^7
	Final fractional coverage	0.1
N	Grid size (100)	516.19 Å × 516.19 Å*
	Grid size (111)	516.19 Å × 514.09 Å*
	Grid size (110)	510.52 Å × 512.62 Å*
	Grid size (210)	514.18 Å × 518.30 Å*

Table 5.1 Default thermodynamic and kinetic parameters used in this model. Different substrates had slightly different N in order to enforce periodic boundary conditions, which are all listed. Other parameters are the same for all substrate potentials.

The evolution of a MD system occurs in timesteps which correspond to a real time, in our case 1 fs steps. At the end of each timestep, conditions are applied upon the system in the form of a thermostat, which ensures that particle velocities are selected from a certain ensemble. As our simulations have a constant temperature, rather than a constant energy, we select from a NVT ensemble—that is an ensemble which keeps particle number, volume and temperature constant. This is equivalent to keeping our system in contact with a heat bath. The thermostat we use is the Langevin thermostat, which subjects all particles to a random force, and lowers their velocities with a constant friction at the end of each timestep. The average magnitude of the friction and applied forces guarantee that the ensemble obeys the fluctuation–dissipation theorem and thus NVT statistics. Note that temperature here is in a statistical sense—the average system temperature is inferred from the total kinetic energy of the system, and the thermostat ensures that this temperature is on average maintained, with fluctuations appropriate to the size of the system. For more details of the implementation of thermostats in LAMMPS we refer the reader to the resources provided by the creators [99], or to one of many reference texts [66].

The protocol for simulations takes the same form as that proposed in Chapter 3,

and similar to that of real CVD [145]. There is an initial growth phase in which an atom is inserted every 100 fs until 12 500 atoms have been inserted. This corresponds to growth to a fractional coverage of 0.1. Atoms are inserted a distance of 5 Å away from any other atom, in order to minimise change in system energy. When the fixed number of atoms have been inserted, growth is halted and there is an annealing phase in which only atomic diffusion takes place.

An example snapshot of islands grown on the (111) potential are shown in Figure 5.3, at 1200 K and 3000 K. Islands grown at the lower temperature are dendritic in shape, whilst islands grown at higher temperature are compact. Observations of graphene islands grown at different temperatures are seen to undergo a dendritic to compact transition of shape at approximately 700 K–900 K on Cu(111) [91] and 892 K on Cu(100) [147]. This suggests that the temperature in this abstract model does not correspond to real world temperatures, but the behaviour exhibited by the graphene islands does. As we want to explore growth in the compact regime, we run simulations at the higher temperature of 3000 K.

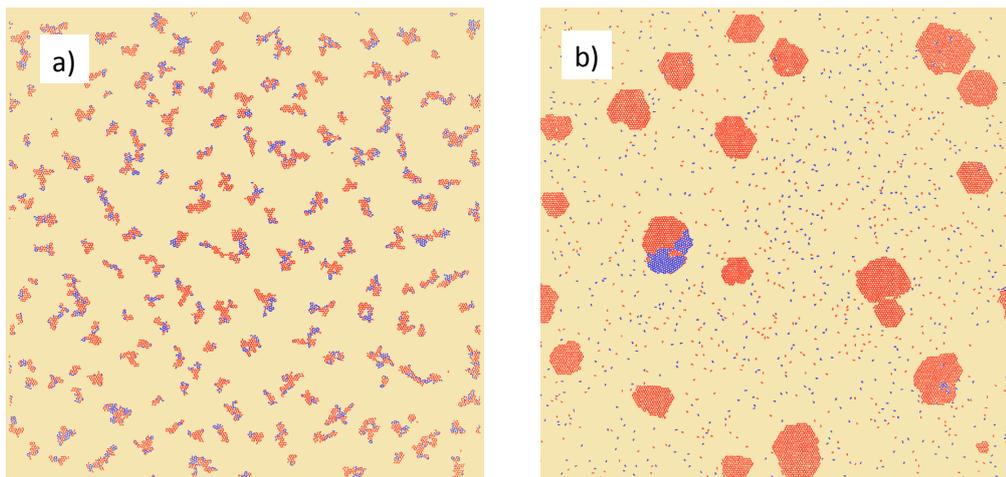


Figure 5.2 Snapshots at the end of simulation for simulations run at two different temperatures, 1200 K (a) and 3000 K (b). Simulations at the higher temperature show compact islands whilst simulations at the lower temperature show dendritic islands. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. With the exception of temperature, simulations were performed according to the parameters described in Table 5.1.

In order to measure some statistics of growth reported in this chapter, we use a recursive algorithm to detect clusters from a list of atomic positions reported at a given timestep t . This is slightly more complex than the similar methods reported in previous simulations of on-lattice growth as atoms are no longer rigidly bound to specific sites. Atoms are defined as neighbouring if they are within 2 Å of each other. Neighbouring atoms are recursively added to a list of all atoms in any given cluster, with analysis then performed on this cluster.

We define two quantities to observe during growth:

1. $F(t)$, the fraction of atoms in carbon islands containing more than 50 atoms
2. $N(t)$, the number of carbon islands containing more than 50 atoms, per simulation.

5.3 Results and discussion

Altering the symmetry of the effective substrate potential has a substantial effect on island growth rates and epitaxy, but the qualitative features of growth are similar for all potentials. For typical examples of growth on each potential we refer the reader to the supplementary videos for this chapter described in Appendix A, which plot the C–C bonds as the simulation evolves. We observed two principle orientations of island relative to the $\theta = 0$ direction in Figure 5.1, with C–C bonds aligned along $(-\pi/3, 0, \pi/3)$ or $(-\pi/2, -\pi/6, \pi/6)$. We discuss these orientations in greater detail later in the chapter, but mention this observation here to explain the colour scheme used in the videos and snapshots presented in this chapter. Bonds are coloured according to which of these orientation they are most closely aligned with: if they are aligned within $\pi/12$ of the peaks in the first orientation—or are in the range $[-5\pi/12, -3\pi/12]$, $[-\pi/12, \pi/12]$ or $[3\pi/12, 5\pi/12]$ —they are coloured red; if they are closer to the second orientation, they are coloured blue.

Snapshots of the graphene islands at the end of the simulation are shown in Figure 5.3. Islands are roughly circular or hexagonal in shape, with graphene islands favouring a zigzag termination and a preference for 120° corners. Graphene islands grown on copper have been observed with four-lobed, dendritic or hexagonal morphologies [110], with different morphologies linked to a variation of experimental growth conditions [128]. Graphene islands which nucleate at separate locations and merge to form larger islands deviate more strongly from this circular morphology, and have less regular shape, such as the one displayed in Figure 5.4 and Video 5.5. This is similar to the islands formed from simulations without roughness in Chapter 3, and is qualitatively similar to STM images of islands grown via CVD on polished copper [76]. We also observe a line of 5–7 defective rings along the grain boundary, which is similar to experimentally observed grain boundaries [49]. Very regular hexagonal islands have been observed in CVD of graphene on liquid Cu [39], and the absence of possible healing mechanisms such as a surface roughness as reported in Chapter 3 or in more detailed simulations [68] likely account for this difference.

We observe three stages of island formation in these simulations:

1. Saturation or induction. Carbon atoms are inserted into the system until they reach a critical saturation.

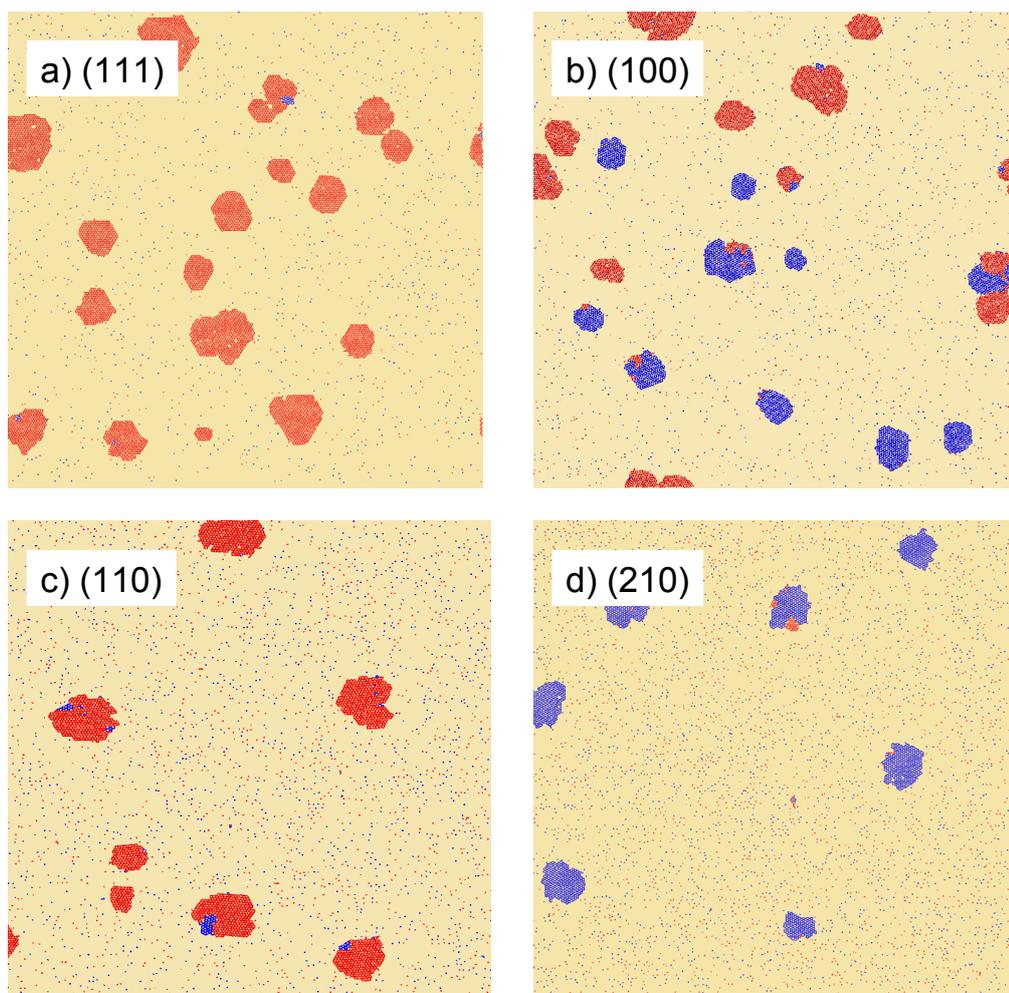


Figure 5.3 Snapshots of C–C bonds at the end of growth and annealing for the four effective substrate potentials, as labelled. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Simulations were performed according to the parameters described in Table 5.1.

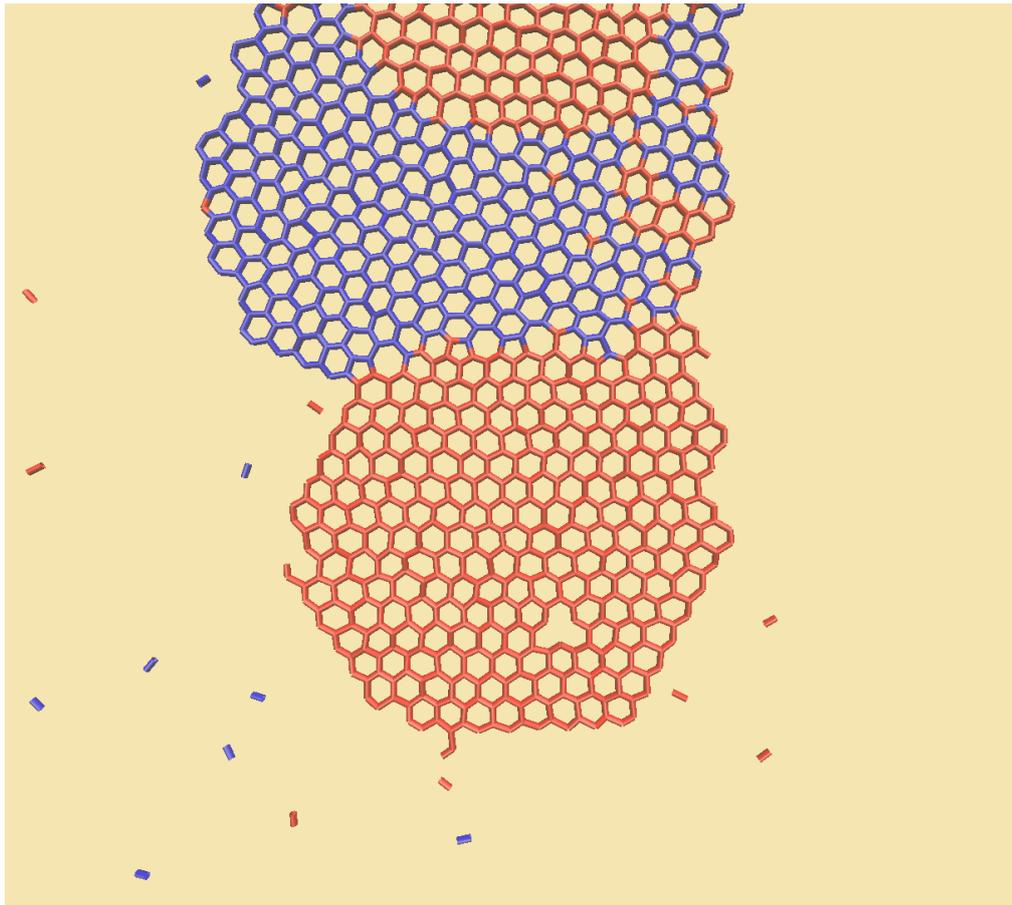


Figure 5.4 A boundary between two clusters which meet during growth, from a simulation on the (100) effective potential. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Simulations were performed according to the parameters described in Table 5.1.

2. Nucleation. When a certain concentration of carbon atoms is reached (which is different on each potential), islands begin to form.
3. Island growth. After the initial islands have nucleated, they slowly absorb the remaining small carbon islands in the system, forming larger islands.

These stages are clearly visible in the videos of growth described in Appendix A, and also in the time dependent population of island sizes in Figure 5.5. This figure shows an initial growth of small islands between one and five carbon atoms on the (100) effective potential, which drops off when nucleation begins, at around 0.5 ns. The larger islands then nucleate and grow until the end of nucleation, at around 1.5 ns, after which the number of large islands remains the same, but the population of other species slowly diminishes as they are absorbed into the larger clusters.

As seen in the videos, clusters are clearly highly mobile during the simulation, with clusters of substantial size diffusing and rotating during growth. Larger islands appear to move less, and to be fixed to a single orientation, indicating a higher degree of epitaxy with the surface. Whilst these larger islands contribute to carbon transport, the most mobile elements of the surface are the monomers and small clusters. The population of dimers is two orders of magnitude higher than any other island size in Figure 5.5, with monomers joining other islands almost immediately after being added to the system.

In DFT simulations of early stage graphene growth, dimers have been reported as more stable than monomers [36], and as having a far greater rate of formation and migration compared to other carbon fragments [104]. In essence, this demonstrates that the cost of breaking a carbon dimer is higher than the cost of migration. This is commensurate with STM imaging of low-temperature CVD of early stage graphene growth, which indicates the large dimer population [92]. Carbon dimers have also been identified as the dominant feedstock to graphene growth [147], due to their strong bond strength and faster mobility. We expand this discussion of the importance of dimers in these simulations with the calculation of diffusion rates later in this chapter.

The scale and scope of the simulations we present here are different to previous MD simulations of graphene growth. MD simulations have been used to identify active surface species in early stage island formation for graphene grown on copper [152] and nickel [81]. Studies such as these simulate the substrate as a bulk of material rather than a 2D material with a geometric potential. This tends to limit them to exploring systems containing 100s of carbon atoms on a ps timescale, whilst the approach presented here simulates 10000s of carbon atoms on a ns timescale.

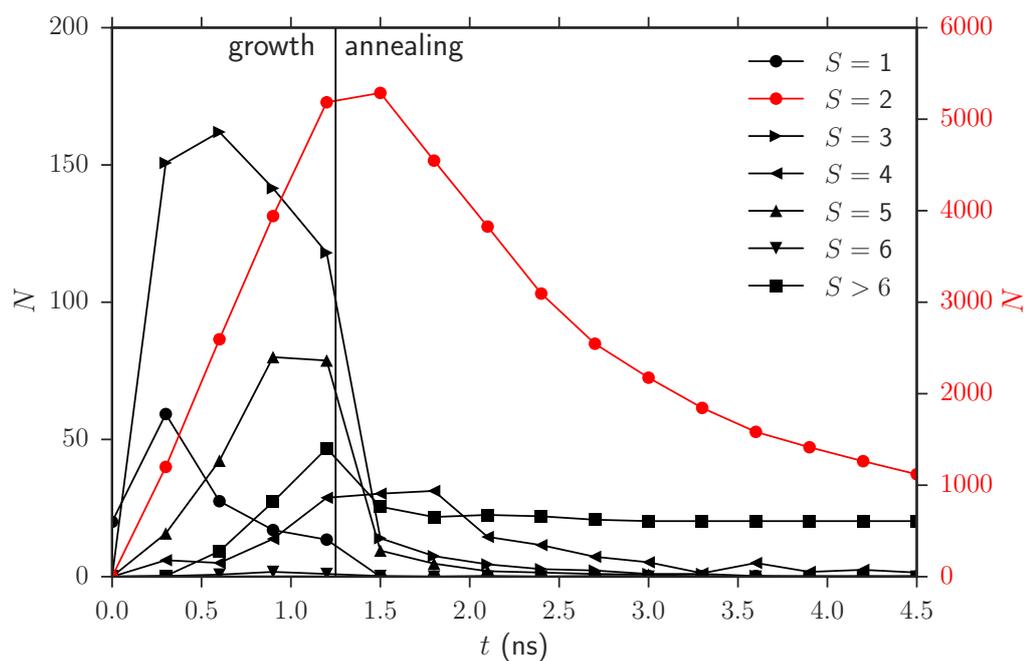


Figure 5.5 Counts of different island sizes for simulations on (100), over the entire simulation timescale. Note the dimer population count is substantially higher than the other island sizes, and is plotted on a separate scale. The growth phase finishes at 1.25 ns, after which no more carbon atoms are added to the system and it is left to anneal. Behaviour was similar in simulations performed on the other effective surface potentials, exhibiting the same characteristic features. Statistics were obtained from 5 simulations; the lines are guides to the eye and error bars are omitted for ease of visualisation.

5.3.1 Island orientation

One of the main aims of introducing a geometric potential is to investigate the orientation of clusters on different facets. We define the angle from the y axis, θ , and use this to measure all angles in the simulations. This direction is marked in Figure 5.1. The distribution of θ for the four different effective substrate potentials is shown in Figure 5.6. As for some of the substrates this contains a high level of background noise due to the relatively high population of dimers, the distribution of angles for bonds in an island containing more than 50 atoms is also shown. The threefold symmetry of graphene gives the bond angle distribution three peaks for a single orientation.

We observe a single orientation on (111) and (210), and two orientations separated by $\pi/3$ on (100). The distribution on (110) is centred around a single wide peak, with a wider and less regular form than that of the other three potentials. On (111), (100) and (210), the peaks can be fitted with a Gaussian curve with full width half maximum (FWHM) 0.181(1), 0.260(4) and 0.186(2) respectively. A Gaussian curve does not fit the orientation distribution from (110) well, and in order to determine the reasons behind this broader peak we employ the following method.

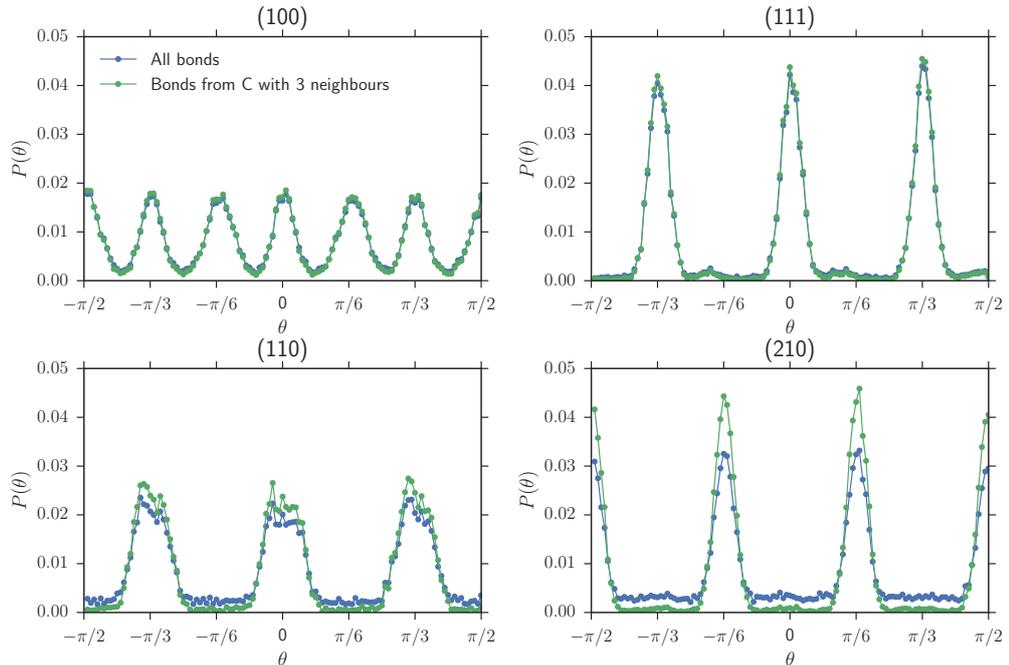


Figure 5.6 Angular distributions of bonds on the four different facets taken at the end of simulation. The blue lines show the angular distribution of all bonds and the green lines show the angular distribution of bonds connected to carbon atoms which have three bonds, thus excluding edges. On (111), (100) and (210), the peaks can be fitted with a Gaussian curve with full width half maximum (FWHM) 0.181(1), 0.260(4) and 0.186(2) respectively.

Firstly, we removed the threefold symmetry of θ by mapping the regions of θ space $[-\pi/2, -\pi/6]$, $[-\pi/6, \pi/6]$ and $[\pi/6, \pi/2]$ onto a variable ψ with domain $[0, \pi/3]$. The wide peak in the bond angle distribution observed three times in Figure 5.6 is then reduced to a single wide peak in ψ space. Secondly, we recalculate bond angles per cluster, giving us a set of bond angle distributions per cluster $P_i(\psi)$. Thirdly, we find the mean of each of these distributions, and divide the set of $P_i(\psi)$ into two sets, one with mean $> \pi/3$ and one with mean $< \pi/3$, and find their average distributions. We plot these two distributions in Figure 5.7, and observe that the wide curve seen in Figure 5.6 is actually two curves with the same FWHM $0.128(2)$, similar to that observed on (111) and (210). The peaks of these two orientations are separated by 10.8° . Whilst this is more complicated than simply computing the average of each cluster's orientation, it allows for excellent statistics on the resulting distributions of bond angle in Figure 5.7.

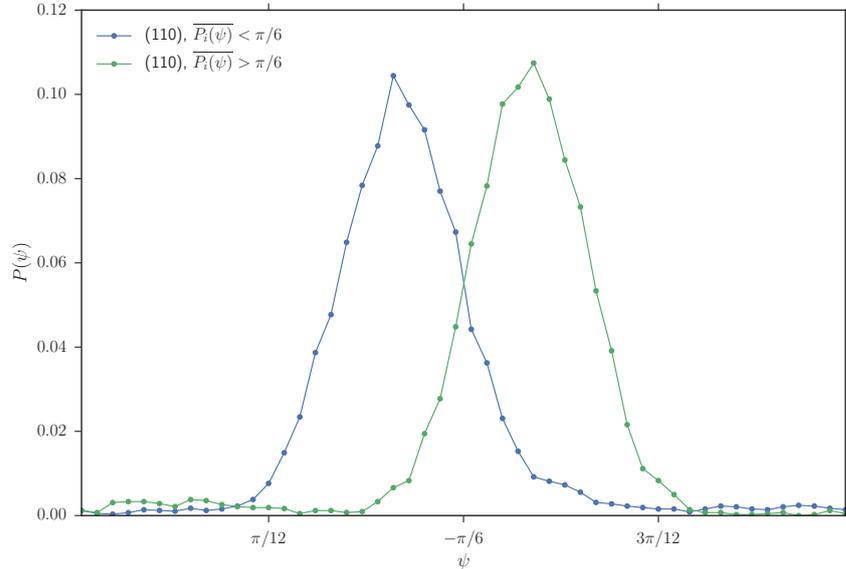


Figure 5.7 Angular distributions of C–C bonds from simulations on the (110) effective substrate potential. The angular distributions in θ seen in Figure 5.6 are reduced onto a domain ψ covering $[0, \pi/3]$, effectively removing the threefold symmetry of graphene. Plotted are the probability distribution of all clusters with mean $< \pi/3$ and $> \pi/3$. The separation between the two peaks is 10.8° . Simulations are performed using the parameters in Table 5.1, and averaged over 5 simulations, for statistics on between 50 and 100 islands.

These results make sense geometrically; (111) has a good lattice match with graphene, as does (210) and (100), with the two peaks seen in the (100) simulations from the symmetry of (100). On (110), there is a mismatch between the graphene and the substrate lattice, leading to a misalignment of the graphene with the substrate potential. Experimentally, graphene has been widely observed to grow with a single orientation on Cu(111) [145, 141, 50]. Growth reported on Cu(100) is more complicated, with observation of a

feedback effect restructuring the copper lattice underneath the graphene islands to Cu(110) [141]. In our simulations we observe a broader peak in the distribution on (100) which could indicate a more complex behaviour, but as the single orientation islands observed in videos and snapshots of growth show two main orientations, this broad peak is most likely due to the increase in multi-orientation islands, and the higher density of defects that comes with it.

The two orientations expected from a common periodicity which we observe on (110) are also observed experimentally [141] from LEED patterns of graphene grown on single crystal Cu(110), with the same deviation of 10.8° . As this model has no mechanism for the copper surface to change during a simulation, there is no way it could capture the feedback effect. Effective comparison to Cu(210) is challenging as a heated copper single crystal of this orientation would most likely restructure to a lower index facet at high temperature. Aside from this complication with (210), and the complex behaviour for (100), the growth on each crystallographic facet is accurately represented. Representative examples of defect free regions of graphene islands occupying the principle orientations discussed in this section are shown in Figure 5.8. Note that the bond distortion is visible, unlike in DFT simulations [20], and that the graphene island grown on (110) has a slight tilt.

5.3.2 Growth statistics

One of the main advantages of this style of modelling is the ability to observe growth statistics on the different effective substrate potentials over a long timescale, and with a greater level of detail than would be possible to observe experimentally. Already in this chapter we have discussed the population of small islands in Figure 5.5, and we now discuss the evolution of the fractional area $F(t)$ and island number $N(t)$, in Figure 5.9. Both of these quantities follow the same growth stages described above: a saturation stage where no large islands form, a period of rapid growth where islands nucleate and grow, and a slowing down of growth in the latter stages of the simulation. This three stage growth process leads to a sigmoidal 'S' shape growth function.

A growth model with a sigmoidal growth curve recently applied to data from graphene growth on copper [18] is the Gompertz curve, which has a history of application in generic growth models [142] and biological models of cell growth [73]. In this context, the island area enlargement rate is assumed to be proportional to the island area A , and an exponential decay

$$\frac{dA}{dt} \propto A e^{-kt}. \quad (5.2)$$

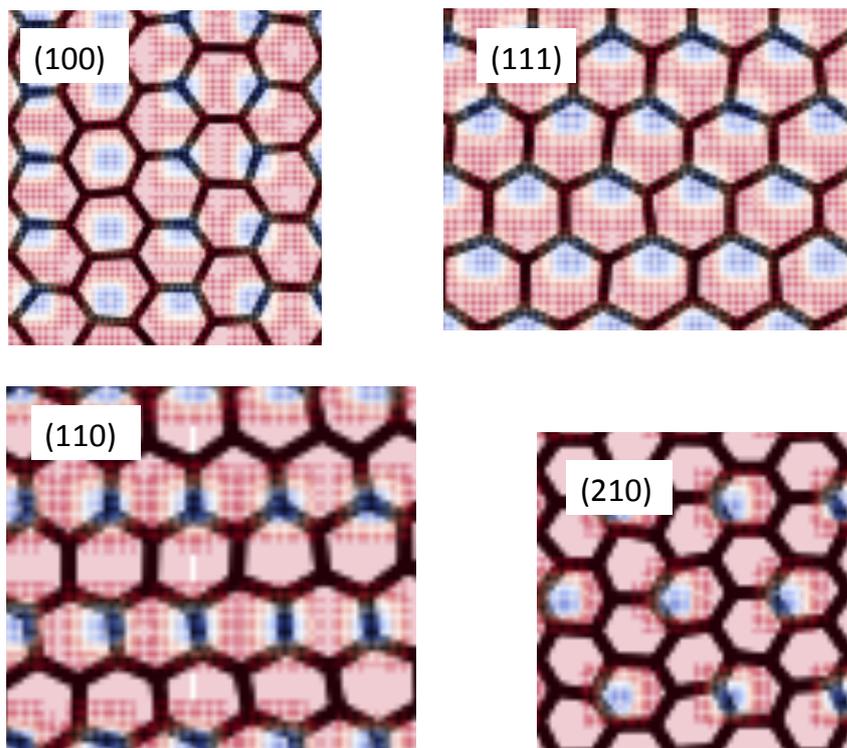


Figure 5.8 Examples of graphene island orientation and atomic position on top of the surface potential. Shown are: a) (100), b) (111), c) (110) and d) (210). The potential underneath the graphene follows the same convention as in Figure 5.1. Note the slight tilt on (110), and observable bond distortion.

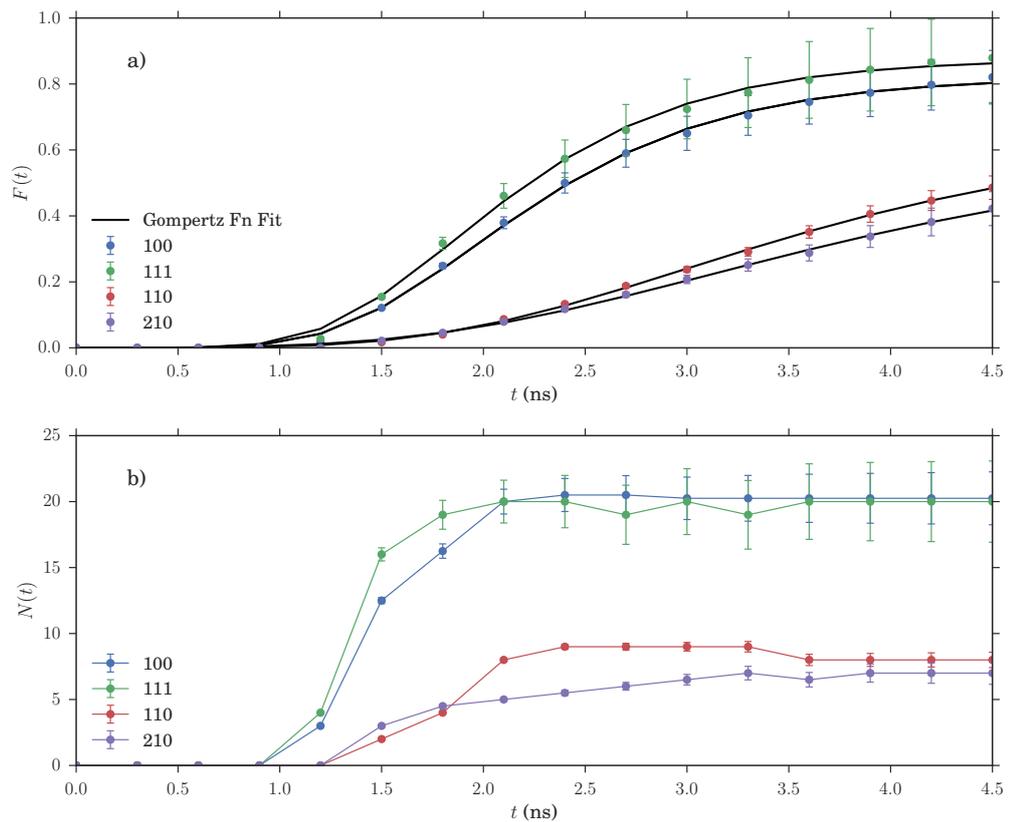


Figure 5.9 Comparative statistics on the growth of graphene clusters on the different surfaces. Pictured are a) the fraction of carbon atoms in clusters greater than 50, $F(t)$, including a fit from a Gompertz curve with parameters in Table 5.2, and b) the total number of clusters greater than 50 atoms per simulation, $N(t)$. Monomer insertion ceases after 1.25 ns. All statistics were obtained from 5 simulations with parameters in Table 5.1

This leads to a Gompertz type curve for the area of an island

$$A(t) = A_{\max} \exp \left\{ - \exp \left[\frac{e\mu_m}{A_{\max}} (t - \lambda) + 1 \right] \right\}, \quad (5.3)$$

where e is the base of the natural logarithm, and the fitting parameters have interpretation with A_{\max} as the maximum possible flake size at infinite time, λ the time lag and μ_m the maximum growth rate. This fit is applied to the data we obtained for the fractional area $F(t)$; we observed excellent agreement with the fitting parameters given in Table 5.2, and a $R^2 > 0.97$ throughout. The agreement of our data with this model indicates the assumptions about island area are correct: the two factors determining the rate of growth of islands are the size of island and the amount of feedstock remaining, which decreases exponentially over time as the carbon is absorbed by islands. As this has been used to model observations of experimental island size of graphene on copper [18], our model has agreement with experimental island growth. What is remarkable, however, is the difference in scale. The island sizes involved in the use of this growth curve in an experimental context are grown in minutes on the scale of μm , whilst the islands here are grown in ns on the scale of nm. This difference in scale points towards behavioural similarities of islands of different sizes.

We observe a different evolution of $A(t)$ and $N(t)$ on the different effective substrate potentials, but this difference is most pronounced between the isotropic potentials, (100) and (111) and the anisotropic ones (110) and (210). Nucleation begins at an earlier time on the isotropic potentials than the anisotropic, and the eventual number of islands which nucleate is higher. The sigmoidal shape of $A(t)$ is observed on all four potentials, with a Gompertz type trajectory fitting the data with different fitting parameters given in Table 5.2. This suggests that despite the differences in nucleation density, the growth follows the same characteristic features on all four potentials. We also report $A(t)$ is marginally higher for all t on the (111) potential than the (100).

Substrate	A_{\max}	λ (ns)	μ_m (ns ⁻¹)
(100)	0.88(1)	0.50(2)	1.20(3)
(111)	0.82(1)	0.44(1)	1.25(2)
(110)	0.65(2)	0.20(1)	1.77(2)
(210)	0.60(3)	0.16(1)	1.72(3)

Table 5.2 Fitting parameters for the four Gompertz type curves in Figure 5.9. Note A_{\max} is the fractional area and is thus dimensionless.

Different nucleation rates on polycrystalline foils have been reported for graphene grown on platinum [137], with incubation times before islands nucleate between 90 s and 2500 s. The authors link these findings to facet-dependent precursor dissociation rates, and

differences in carbon diffusion. In this model we do not account for precursor dissociation rates; in reality, carbon added to a system via CVD is added via some hydrocarbon precursor, which needs to be broken down in some way in order to add carbon to the surface. As we still observe different incubation times on different effective substrate potentials, we can confirm the importance of facet-dependent carbon diffusion rates.

On copper, growth on Cu(111) has been reported as faster than growth on other facets [141, 50, 145] which is in agreement with the observation that on the (111) effective substrate potential growth is indeed faster, although only marginally in this study. Higher index copper substrates are linked to growth faster than on Cu(100), and with more compact island morphologies [145] which are not replicated in this modelling. In the remainder of this chapter we discuss some reasons for these different growth rates in our model, and potential reasons for the discrepancy between ours and experimental findings.

5.3.3 Diffusion

Monomer diffusion rates were calculated from the trajectories of single particles inserted at a random location on each effective substrate potential. The particle displacement in x and y directions is recorded from the starting location over time, and the diffusion coefficients are obtained from the gradient of the variation of this position with time, averaged over 1000 simulations. The diffusion matrix

$$D = \begin{pmatrix} D_{xx} & D_{xy} \\ D_{yx} & D_{yy} \end{pmatrix} \quad (5.4)$$

was thus determined, with the coefficients D_{xx} , D_{yy} and D_{xy} calculated from the correlations between relevant positions. The total diffusion rate D is the trace of this matrix. Dimer diffusion rates were calculated using the same method with the following exceptions. When initialised, the dimers had a random orientation as well as random position, separated by the C–C bond length of 154 pm . The distance from this starting location was determined from the dimer centre of mass. Due to the speed of these simulations, the statistics on the diffusion coefficients calculated are good, with deviations typically only in the third significant figures. This allows us to estimate the diffusion coefficients of carbon atoms on the effective substrate potentials to a high degree of accuracy.

The diffusion coefficients for monomers and dimers on the four effective substrate potentials are shown in table 5.3. Monomer diffusion coefficients are similar across all effective substrate potentials, with the (111) surface having a slightly higher monomer diffusion rate. A difference in D_{xx} and D_{yy} indicates diffusion which is anisotropic, with particles preferring to diffuse in a certain direction. Monomer diffusion coefficients are slightly

anisotropic on (100), (111) and (210), and more significantly so on (110).

Substrate	Monomer			
	D_{xx}	D_{xy}	D_{yy}	$\text{Tr}(D)$
(100)	2.74(1)	$1.9(1) \times 10^{-4}$	2.69(1)	5.44(2)
(111)	3.12(2)	$7.2(2) \times 10^{-2}$	2.70(2)	5.85(4)
(110)	3.13(2)	-0.18(1)	2.36(1)	5.50(2)
(210)	2.59(1)	-0.39(1)	2.83(1)	5.42(2)
	Dimer			
	D_{xx}	D_{xy}	D_{yy}	$\text{Tr}(D)$
(100)	1.52(1)	0.12(1)	1.56(1)	3.07(1)
(111)	1.52(1)	0.12(1)	1.67(1)	3.20(2)
(110)	1.81(1)	0.13(1)	0.85(1)	2.66(1)
(210)	1.13(1)	-0.16(1)	1.02(1)	2.15(1)

Table 5.3 Monomer and dimer diffusion coefficients on the four different substrate potentials, calculated from the gradient of correlations between x and y positions as described in the text. All values have units $10^{-10}\text{m}^2\text{s}^{-1}$, and errors were estimated from the residuals of the numeric fit of the gradient.

The dimer diffusion coefficients were all lower than the monomer diffusion coefficients reported on each substrate. They were also more varied across the different effective substrate potentials than the monomer values, with dimer diffusion coefficients on (110) and (210) significantly lower than on (111) and (100). The dimer diffusion coefficient for (111) was again reported to be the highest. Dimer diffusion coefficients on (100), (111) and (210) were close to isotropic, with only a slight preference for a given direction shown. Dimer diffusion on (110) was the most anisotropic of any of the coefficients, with the D_{xx} coefficient almost double the D_{yy} . The cross terms D_{xy} , which determine the linearity of diffusion, were at least an order of magnitude lower than D_{xx} and D_{yy} on all effective substrate potentials, indicating all diffusion we report is linear.

Dimer diffusion rates have been reported as comparable [147] in a comprehensive theoretical study of dimer kinetics on Cu(100) and Cu(111), or greater [104, 20] in DFT, to monomer diffusion rates. Whilst our effective substrate potentials produce lower dimer diffusion coefficients, we still saw that a substantial proportion of the carbon islands were dimers in Figure 5.5. In a rate equation model of graphene growth based on attachment and diffusion barriers [147], dimers were reportedly one or two orders of magnitude more frequent than trimers or monomers respectively, which is similar to the results shown in Figure 5.5.

The mechanism by which the monomers and dimers diffuse is further explored in a supplementary video of a monomer and a dimer diffusing on a (100) substrate, Video 5.6.

Atomic trajectories were obtained in independent simulations, but are plotted together; the monomer and dimer are diffusing completely independently in this video. The monomer gets trapped in individual wells, only moving across the substrate when it escapes the well in which it is trapped. When it does escape the well, it does not necessarily diffuse to near neighbour wells, and frequently hops to a well several lattice sites away from the one it just escaped. Monomers are reported as stable in octahedral sites on Cu(111) [147, 104], with mobility between sites high. The long range monomer diffusion exhibited in this model leads to a low late time monomer density.

The dimer diffusion generally moves between near neighbour wells, one atom always seeking to be at the bottom of the well. The bond flexes, and the atom not located in a well can relocate to a nearby well. There is far less long-range hopping than exhibited by the monomer, and this behaviour is qualitatively similar to that reported in DFT studies [104]. The lower dimer diffusion coefficients reported on (110) and (210) could be attributed to the lowered probability of the free atom of the dimer pair reaching another well due to the lower density of wells on these substrates.

The detailed first principles studies of small island behaviour we have discussed in this section [104, 147, 20] focus on Cu(111) and Cu(100), the former due to the structural geometry similar to graphene and the latter due to its easy availability in cheap copper foils. In this work we have also simulated effective surface potentials based on Cu(110) and Cu(210). These two surfaces, visualised in Figure 5.1 are different to Cu(111) and Cu(100) in two key ways: they have a lower density of copper wells, with a greater separation between wells and they are not isotropic. The anisotropy of the surface potentials is reflected in the increased anisotropy in the diffusion coefficients, although (110) sees this effect far more than (210). We also saw different statistics of growth on the isotropic surfaces to the anisotropic, as reported earlier in this chapter in Figures 5.9. The isotropic surfaces had more graphene islands nucleate than the anisotropic, and reported a correspondingly higher fraction of the carbon atoms in large islands, $A(t)$.

5.3.4 Flake strain energy

Our model has two inter-atomic potentials which contribute to the stability of graphene islands: a C–C potential and a C–Cu potential. Using the BOP C–C potential in the absence of the C–Cu potential, a graphene island would form a perfectly hexagonal planar structure [153]. In order to locate carbon atoms at the bottom of the wells in the C–Cu potential, this hexagonal structure will likely be distorted or strained in the course of a simulation on a substrate. In order to investigate this strain, we estimate these two contributions to the total energy of an island.

The LAMMPS code has a built-in quenching mechanism in the minimise com-

mand. This command perturbs atoms locally as the temperature is lowered, causing each atom to seek its local minimum. Each island consisting of more than 50 atoms at the end of the simulation time is isolated onto its own potential grid. For each isolated island, we then obtain three total systems corresponding to the three situations shown in Figure 5.10. The first energy, E_1 is the energy of an island after quenching is performed on the original substrate. The substrate potential is then deactivated, with the well depth D set to 0, after which the system energy is recalculated to give a second energy E_2 . The third and final system energy E_3 is then obtained by quenching this flake in the absence of a C–Cu potential.

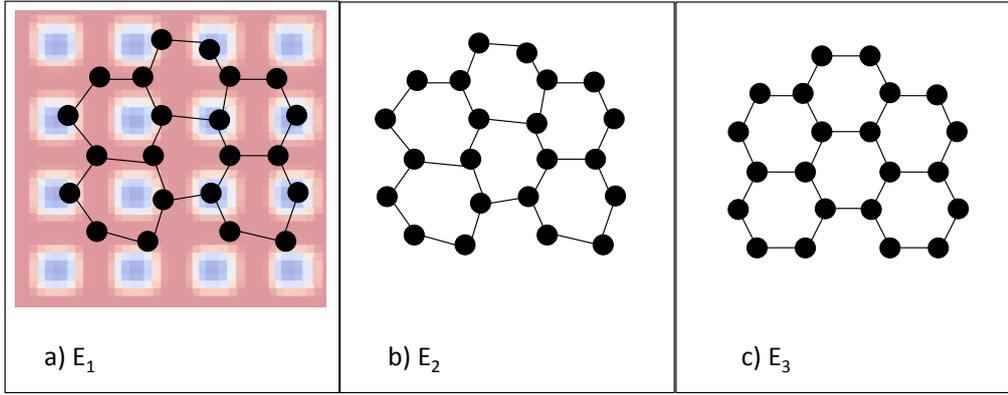


Figure 5.10 Diagram representing the state of a flake for calculations of the three energies used to calculate the strain energy E_P and the distortion energy E_D . Panel a) shows an island quenched after simulation on a substrate potential with system energy E_1 , b) shows this island after the potential is removed for the calculation of E_2 , and c) the flake after it is quenched without the potential for energy E_3 . Bond distortions are exaggerated for effect and do not represent real atomic positions.

For an island of size S , we calculate the substrate binding energy as the energy gain per atom from the effective potential E_P , and strain energy per atom E_S . The substrate binding energy E_P is defined as the energy gained by an island from the C–Cu interactions compared with a flake in the absence of potential

$$E_P = \frac{1}{S} (E_1 - E_2). \quad (5.5)$$

The strain energy E_S is defined as the energy difference between a quenched island without a potential, and the energy from the atomic positions also without a potential,

$$E_S = \frac{1}{S} (E_3 - E_2). \quad (5.6)$$

The two energies E_P and E_S allow us to estimate the extent to which a carbon island strains its C–C bonds, and the gain in energy it receives for favourably orienting on the effective substrate potential for the four different facets we investigate.

Substrate	E_S (eV)	E_P (eV)
(100)	-0.012(1)	0.40(3)
(111)	-0.012(1)	0.45(5)
(110)	-0.011(1)	0.28(3)
(210)	-0.012(2)	0.17(1)

Table 5.4 The strain and substrate binding energies E_S and E_P for the four surfaces, given per atom. Standard error on the final decimal place in brackets.

The values of E_P and E_S are shown for the four effective substrate potentials in Table 5.4. The strain energies of islands on different substrates differ by less than their standard errors and are much smaller than their substrate binding energies. The substrate binding energies on (111) and (100) are of similar magnitude and are higher than the binding energies of (110) and (210). It is surprising that the binding energy of (110) is lower given the epitaxial match reported in at least one direction for graphene on Cu(110) [141].

5.4 Limitations

The 2D model of graphene growth we have introduced in this chapter has replicated aspects of graphene growth on polycrystalline films. Notably we have reproduced the form of graphene area evolution with a Gompertz trajectory [18], and offered insight into the connection between dimer diffusion and nucleation rates on different substrates. The reproduction of realistic graphene structures and orientations must in large part be attributed to the modern potential we have used in this model, which has successfully modelled a system far outside its original conception and parametrisation. The simplistic effective potential we have reproduced has oriented graphene islands in a fashion consistent with most experimental observations. However, this 2D model of growth does not aim to—and cannot—capture all aspects of a real system.

Real substrates are never exactly planar, with atomic steps and small terraces [141]. The presence of step edges and terraces as nucleation sites for graphene has been linked to higher nucleation rates [137, 135]. Whilst it would be possible to incorporate some step modification to our effective substrate potential, at present this model does not capture this behaviour. On cooling after growth, graphene may wrinkle due to the thermal contraction of the copper lattice and graphene lattice occurring at different rates [69]. Wrinkling or rumpling effects are by nature three dimensional, and cannot be captured in this model. Small islands have also been reported to form domed structures [36, 133] whilst bonding to the substrate along the edges, indicating that there are some aspects of nucleation which this model cannot capture.

In Chapter 3 we noted that substrate roughness was a possible route to island cleaving, defect healing and eventually further growth through destabilisation of islands due to local roughness. No such mechanism exists in this model, and the lack of this and similar mechanisms prevents the islands in these simulations from growing to larger sizes. The static potential does not capture the nearly molten copper surface, which sublimates during a typical CVD growth run.

5.5 Conclusion and outlook

In this chapter we have introduced a 2D MD model of growth on effective substrate potentials, using a realistic bond order potential optimised for growing carbon structures. Motivated by observations of different orientations and growth rates on different crystallographic facets [141, 137], we have performed simulations on four different crystallographic facets, with geometry taken from copper lattices, Cu(111), Cu(100), Cu(110) and Cu(210). We have reported orientations on each facet consistent with that expected by lattice matching and experimental observations [141]. We have investigated growth on each facet, reporting a higher growth rate on Cu(111), and significantly lower growth rates on the anisotropic (110) and (210). We have discussed the reasons behind these reduced growth rates, in particular investigating monomer and dimer diffusion, and flake strain energy.

Whilst this style of modelling cannot capture all realistic CVD growth conditions, it has reported excellent agreement with experimental observations of graphene orientation and growth on different crystallographic facets [141, 50, 145]. This style of modelling, using realistic carbon potentials with manually configured substrate potentials, has allowed us to investigate a system on a far greater scale than is typically seen in DFT simulations, and with a far higher level of atomic detail than is typically seen in kMC or more coarse grained methods—and reproduce the strengths of both.

Looking forward, the potential for this novel modelling technique is significant. More realistic substrate potentials could be explored, incorporating lower than first layer substrate geometry. A mechanism for feedback between the two layers could be introduced such as static or dynamic surface roughness, similar to that in Chapter 3. The effects of temperature, carbon insertion rate, C–Cu interaction strength and fractional coverage could be further investigated, which would allow for better mapping to different experimental protocol. Epitaxial mismatch could be further investigated by changing the lattice separation in the effective potentials.

Perhaps the most exciting extension is the ability to stitch together each of the different facets discussed in this chapter, with some manually defined interface between different regions of the substrate potential. The introduction of such systematic defects, combined

with the realistic and experimentally justified graphene formation dynamics discussed in this chapter, would for the first time allow realistic simulation of graphene growth on a polycrystalline foil.

Chapter 6

Cluster moves in lattice Monte Carlo simulations

6.1 Introduction

So far in this thesis, we have discussed two different approaches to modelling graphene growth which have focused on different aspects of the process. The MC code in Chapter 3 modelled graphene growth on a fixed sheet of pre-defined possible atomic positions, with substrate–graphene interactions via a random site by site roughness energy. By introducing near neighbour moves in this roughness energy, we modelled a dynamic substrate which we argued introduced a temperature–induced mechanism of island size enhancement.

By contrast, in Chapter 5 we presented an off–lattice molecular dynamics (MD) model graphene growth utilising a recently developed inter-atomic carbon bond order potential optimised for growing carbon structures, with substrate interactions defined by a geometric effective substrate potential. Unlike the fixed lattice modelling, this allowed for free formation of clusters in 2D space, and we presented results on island orientation and growth rate on the different facets of the effective substrate potential. We observed kinetic behaviour qualitatively different to that in the MC code. Island nucleation was less frequent, and a significant fraction of the carbon was mobilised in dimers moving across the substrate potential. The videos of growth showed that clusters had a high level of mobility, diffusing and rotating in 2D space. We also observed that the largest clusters tended to stay in a fixed orientation, with less rotation and diffusion.

The importance of cluster kinetics was shown in modelling of gold nanoparticle formation on graphene [10] where islands were modelled as circles, allowing new atoms on the surface to adhere to existing circular islands. These islands were allowed to diffuse up to a maximum size, after which they formed hemispherical islands. These cluster kinetics

were found to reproduce experimental double peaked ISDs [96]. Low temperature scanning tunnelling microscopy performed *in situ* for graphene grown on Cu(111) with CVD reported the importance of small clusters as intermediates to graphene islands [133], and as possible precursors for growth defects [92]. Modelling of early stage CVD growth with DFT [36] demonstrates the importance of different cluster morphology in early stages of graphene growth, and indicates that understanding the mechanisms of these early stage clusters is key to understanding the growth process itself.

MC modelling is appealing due to its low computational demands. The MC code used in Chapters 3 and 4 was parallelised using the standard and portable Message passing system (MPI). A single simulation from the model in Chapter 3 would take around 10 minutes real time on 8 processors, whilst the more intensive work involved in the MD modelling made an equivalent simulation from Chapter 5 take around 2 days, also on 8 processors. Incorporating realistic cluster kinetics into an MC model could lead to more accurate studies, and a better understanding of the importance of such mechanisms to surface growth. Incorporating physically meaningful cluster movement into such code presents an additional challenge we explore in this chapter.

Virtual move Monte Carlo (VMMC) is a cluster moving algorithm developed for systems of strongly attracting colloidal particles [139]. This algorithm has been hailed as powerful due to its inclusion of cluster cleaving—and thus the preservation of detailed balance. It has been refined [107, 106] for translation and rotational moves, as well as pairwise interactions, and seen implementation in models of self assembly and growth [138]. The original authors of the VMMC anticipated such an algorithm being particularly appropriate for systems of strongly attracting particles, such as lattice based modelling, even whilst the algorithm itself was designed for off-lattice systems. One such on-lattice implementation of VMMC has been used to study phase separation in a lattice gas [43]. This implementation considers the importance of parametrising VMMC for lattice simulations with a view to preserving ergodicity and time reversibility. Note that when we discuss cluster moves, we refer to the diffusion of a cluster rather than the better known spin-cluster updates using an algorithm such as that proposed by Swendsen and Wang [134].

This chapter describes the parallel implementation of domain decomposed lattice gas simulations with diffusion and cluster diffusion described by the VMMC algorithm in code written in C with MPI. Firstly, in the methods section we discuss construction of a domain decomposed Ising-like lattice MC model of growth. We then extend this model to incorporate near neighbour diffusion moves, and finally introduce the full cluster moving algorithm. In Chapters 3 and 4, we described Ising-like models of growth with local diffusion and an additional roughness and field lattice respectively. The SL algorithm with diffusion, Algorithm 2 was the basis how the code for these chapters was constructed, and

we discuss alterations to the discuss algorithm which were necessary for these models. We test the performance of this code in the efficiency and speedup section presenting performance information across different system sizes and processor counts.

Finally, we implement the cluster moving algorithm in simulations of graphene growth similar to those in Chapter 3. We observe the effect of increasing the proportion of cluster moves, and identify the typical cluster sizes moved by the algorithm. We use insights obtained from this modelling to discuss the importance of cluster movement in simulations of graphene growth.

6.2 Methods

6.2.1 Domain decomposition in Monte Carlo simulations

Domain decomposition is a computational method for modelling physical processes over space. In problems with a spatial element, the principle is to split the entire simulation region into *domains*, and pass those domains to individual processors. Each processor stores a local copy of the information in the outermost layer of the neighbouring domain which bounds its own domain on each side. We refer to this local copy as the halo region. Operations are performed independently on each domain, with processors periodically updating the halos. During the simulation, and at its completion, data can be recalled onto the master processor for processing and output. A large spatial problem can thus be transformed into several smaller problems, which enables exploration of larger system sizes than would be possible in serial computing.

In Chapter 2, we discussed the concepts of detailed balance and ergodicity, and their importance in MC simulations. To recap, we require our algorithm to be ergodic—that the algorithm is designed in such a way that all points in the phase space would be sampled in infinite time. Whilst it is impossible to prove a simulation is ergodic without running it for infinite time, simulations which fulfil detailed balance—that is that at equilibrium each process is equilibrated by a reverse process—can safely be assumed to be ergodic. This gives us a requirement that each MC move in the cluster movement algorithm have a possible reverse process at equilibrium. Historically, there has been discussion of selecting an algorithm which preserved detailed balance as well as retaining a high level of computational efficiency [75, 26]. Such algorithms have been used to simulate Ising-like systems for a long time [7]. The domain decomposition protocol we used in this model is the synchronous sublattice (SL) algorithm [114], outlined in Algorithm 1. Domains are quartered into subdomains, which are then updated simultaneously across all processors. Provided the halos are updated before the next subdomain update, this removes the possibility of the same spin being updated by two different processors, or of a processor using an old spin for

updates.

In this and subsequent algorithms described in this chapter, we utilise three MPI routines: a gather routine (`MPI_Gather`), a scatter routine (`MPI_Scatter`) and a send receive routine (`MPI_Sendrecv`). The gather routine is an all-to-one operation where all processors send a chunk of information to a root, or master processor. Its inverse is the scatter operation, which breaks up a large piece of information into several smaller pieces, which are distributed across the network of processors. These routines distribute the information across the processors using a tree-like hierarchy rather than having a direct communication from each processor to the root. This reduces the number of communications—a typical bottleneck in speed. The send receive routine is a processor-to-processor routine, in which every processor in the grid posts a message for a certain neighbour (say the right neighbour), then waits to receive the same information from its neighbour which just sent it (the left neighbour). For more information on these routines, or any other aspect of the MPI system, we refer the reader to the MPI standard [44].

Algorithm 1 Synchronous Sublattice (SL) algorithm

```
1: Initialisation:
2:   Send domains to processors MPI_Scatter
3:   Divide domains into subdomains
4:   Get halos from neighbours MPI_Sendrecv
5:
6: for  $t < T$  do
7:   for Each subdomain do
8:     Perform MC moves
9:     Exchange halos MPI_Sendrecv
10:  if  $t\%100 == 0$  then
11:    Gather to master processor MPI_Gather
12:    Output statistics ( $t/100$ )
```

The coloured subdomains in Figure 6.1 shows how the SL algorithm operates. The blue subdomains on each domain are the subdomains which are currently being operated on. The red squares represent the neighbouring subdomain of the active subdomain in domain 1, in periodic boundary conditions. None of the red subdomains are operated on at the same time as the blue subdomain in domain 1, preventing the same region being updated at the same time. Halos can be exchanged before the active subdomain changes.

In this first implementation of a domain decomposed system, we consider an Ising-like model of growth; there is no atomic diffusion, only particle insertion and removal. As we take only first near neighbour contributions to energy calculations, the halo which needs to be stored on each processor is the outermost row of the neighbouring domains only. We update this halo after one sweep of the subdomain—i.e. after we have sampled a number

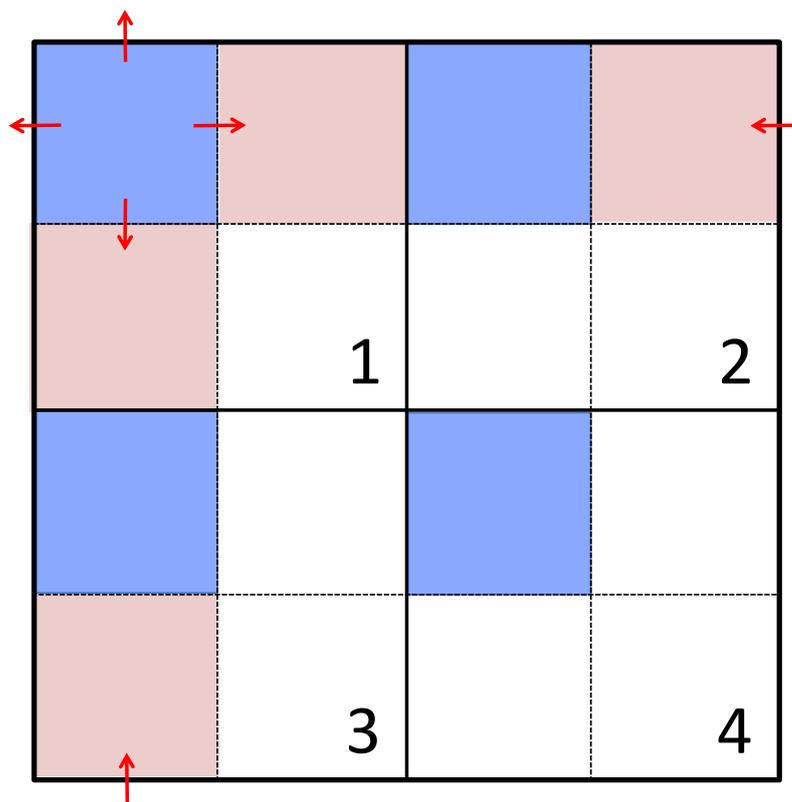


Figure 6.1 Diagram demonstrating the operation of the SL algorithm. The blue squares on each processor are updated simultaneously. The red squares which provide the halo for each blue square are static during the updates of the blue square. This ensures adjacent squares are never updated at the same time, preventing halo information from becoming out of date.

of sites equal to the total number of sites in the subdomain.

6.2.2 Incorporating diffusion

In this section we describe the inclusion of Kawasaki style spin exchange [57] into the SL algorithm described in the previous section. In the growth model we ultimately explore, this represents particle diffusion. This presents a challenge to our previous algorithm because there is now transport across the system which has to traverse the domain boundaries; the halo can no longer be completely static during simulation.

With serial code the diffusion move is simple: a random site is selected, along with a neighbouring target site. If the atomic species are different on each site, a trial swap can be made and the Metropolis acceptance probability for the move calculated. The move is then accepted or rejected based upon this acceptance probability.

To allow for transport in parallel, we extended the halo to two rows of atoms thickness, allowing diffusion into the first row of the halo. Only sites inside the subdomain could be selected for movement, and therefore only the first row of the halo could possibly change. The second row of the halo remained static as before for energy calculations. The corner element from diagonal neighbour processors becomes required for energy calculations, and had to be updated with the remainder of the halo. The SL algorithm with diffusion and exchange of corner elements is shown in Algorithm 2.

A diagram showing the operation of the SL algorithm with diffusion is shown in Figure 6.2. The blue region in this image represents an active subdomain, and the red region the halo for that subdomain. Examples of site diffusion are shown in this image, with any move possible except where it would be a move into the second layer of the halo. This implementation had outputs indistinguishable to serial simulations, with transport across domain boundaries. The algorithm obeys detailed balance because of the requirement to pick sites rather than atoms— whilst an atom exchanged into the halo could not be selected to move back into the subdomain, the site from which it moved could be selected (along with its halo neighbour), generating the reverse move.

One practical consideration of this algorithm is to ensure that atoms are not able to diffuse across boundaries faster than the halos are updated. If the halos are updated too infrequently, atomic diffusion across the boundaries would be different from that of a serial simulation. There is necessarily however, a frequency of halo updating which does leave the simulations indistinguishable from serial simulations, even if this is far more regularly than is practical. The ability to construct an algorithm which only uses local processor-to-processor communications is potentially very time saving in massively parallel simulations.

Altered versions of this algorithm were used in Chapters 3 and 4 to describe rough simulations of graphene growth and a dynamic random field Ising model respectively.

Both of these models called for an additional roughness or field lattice to be used. The implementation of this was to simply duplicate the MPI routines described in Algorithm 2 for the additional field lattice. The model in Chapter 3 also called for a honeycomb lattice and diffusion moves up to the third near neighbour. Incorporating third near neighbour diffusion in a honeycomb lattice required extending the halo to two rows for atomic diffusion and one for energy calculations for a total of three atom thickness. This also involved passing the corner elements in the first two rows of the halo, rather than just the corner element from the first row as described above.

Algorithm 2 Synchronous Sublattice algorithm with diffusion

```

1: Initialisation:
2:   Send domains to processors MPI_Scatter
3:   Divide domains into subdomains
4:   Get halos MPI_Sendrcv
5:   Get corner information MPI_Sendrcv
6:
7: for  $t < T$  do
8:   for Each subdomain do
9:     Perform MC moves with diffusion
10:    Exchange halos MPI_Sendrcv
11:    Exchange corners MPI_Sendrcv
12:   if  $t \% 100 == 0$  then
13:     Gather to master processor MPI_Gather
14:     Output statistics ( $t/100$ )

```

6.2.3 Incorporating cluster diffusion

In this section we introduce a method for implementing the SL algorithm with cluster movement. The implementation of operations on a large cluster in a domain decomposed simulation presents a challenge to the previously described SL algorithm, because of the consideration of how to move a cluster which spans multiple subdomains. Each processor has no knowledge of the full spatial extent of a cluster, and so cannot *a priori* be moved in the same way as in serial code.

The cluster moving algorithm we base our cluster moves on is the virtual move Monte Carlo (VMMC), or ‘cleaving’ algorithm [139], originally developed for systems of strongly interacting pairwise particles. It approximated collective dynamics in colloidal systems and avoided sampling of unphysical kinetic traps. It has since been tested and optimised [107, 106], and seen application in models of lattice gas self-assembly [43].

The VMMC algorithm is implemented as follows. Firstly, a cluster is formed by scanning the neighbour list of a randomly selected occupied site. Atoms are recursively

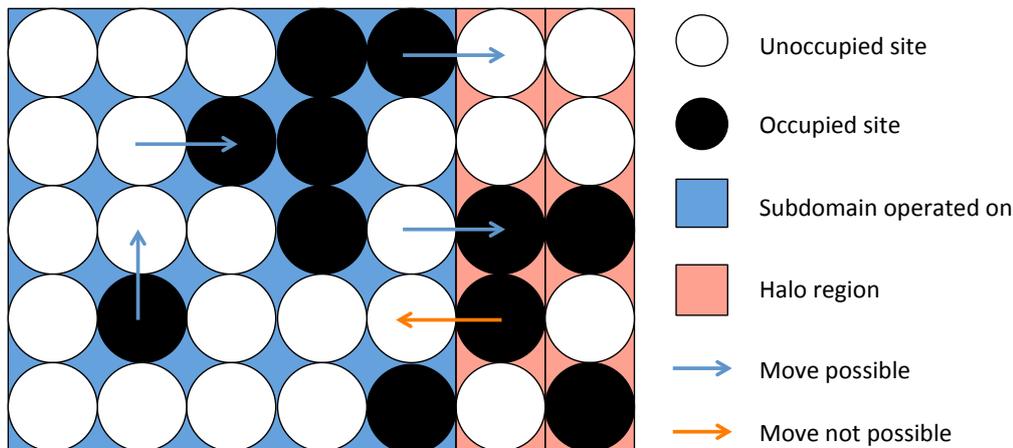


Figure 6.2 Diagram demonstrating implementation of SL algorithm with diffusion. Atoms in the blue region are being operated on in this cycle of MC moves, and can be selected to diffuse. Atoms can diffuse into the halo, but cannot be selected when inside the halo as they are no longer in the blue region. Atoms could diffuse out of the halo if an unoccupied site was selected which neighboured the halo, which means the moves fulfil detailed balance conditions of reversibility. Atoms in the second row of the halo are only used to obtain neighbouring energies for the moves trialled into the first row of the halo.

added to this list with a probability less than unity, $1 - \exp(-\lambda\varepsilon/T)$, where λ is a parameter which determines the likelihood of cluster selection and ε is the bond strength, set to 1 in this model. Each bond between atoms is only tested once. Atoms are added up to a maximum cluster size n_{\max} . A lattice direction is chosen at random, and the cluster is moved with the Metropolis acceptance probability according to the energy difference associated with its change in location. Cluster moves which intercept other atoms are rejected, and all alternative directions are randomly trialled until a direction which does not cause an intercept is selected, or all directions have been trialled. The parameter λ can be chosen between 0 and 1 whilst preserving detailed balance: a move which merges two clusters can be reversed as not all connected atoms in an island are necessarily selected as a cluster to be moved. We set it to 0.9 in this model, as a value closer to unity better mimics diffusive behaviour.

In the original VMCM algorithm, after a cluster is generated, the maximum cluster size n_{\max} is dynamically generated from a distribution such that $P(n_{\max} > n) = 1/n^2$. This gives a diffusion rate proportional to the inverse of the size of the cluster, which is consistent with Brownian motion. As in practice this means that the largest cluster that will be moved is of size around $n = 20$, we select $n_{\max} = 100$ to represent the observation in Chapter 5 that larger clusters were mobile, as well as previous observations of cluster movement in surface growth [10, 96, 36]. We justify this alteration with the comment that the movement of a large island on a substrate is likely to be non-Brownian; the size of the island gives it

inertia. This selection preserves detailed balance as it does not remove any possible cluster moves, merely altering the probability with which they are selected.

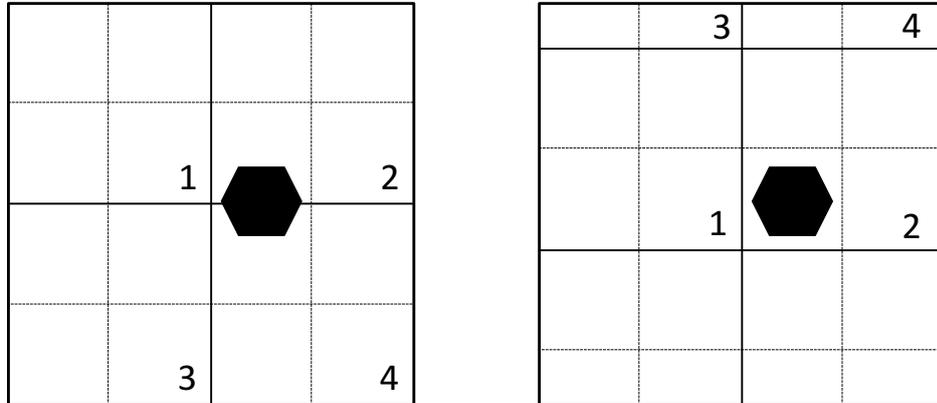


Figure 6.3 Diagram demonstrating the domain rescattering algorithm. The top panel shows a cluster which is unable to be moved by processor 2 or 4 because they possess incomplete information about the cluster. The bottom panel shows the processor grid after scattering; the cluster is now entirely within a single subdomain and can be moved.

Algorithm 3 Synchronous Sublattice algorithm with cluster moves

- 1: **Initialisation:**
 - 2: Send domains to processors `MPI_Scatter`
 - 3: Divide domains into subdomains
 - 4: Get halos `MPI_Sendrcv`
 - 5: Get corner information `MPI_Sendrcv`
 - 6:
 - 7: **for** $t < T$ **do**
 - 8: **for** Each subdomain **do**
 - 9: Perform MC moves with cluster moves
 - 10: Exchange halos `MPI_Sendrcv`
 - 11: Exchange corner information `MPI_Sendrcv`
 - 12: Gather to master processor `MPI_Gather`
 - 13: Reassign domains
 - 14: Send domains to processors `MPI_Scatter`
 - 15: **if** $t \% 100 == 0$ **then**
 - 16: Output statistics ($t/100$)
-

Methods which consider diffusion are more common in off-lattice, non spin system MC simulations. One method used in these off-lattice models to reduce the amount of local processor communications is to disallow diffusion out of domains, and to periodically

reassign the domains to the processors throughout the simulation [125, 108]. By operating on different regions of space after each reallocation, this method has transport indistinguishable from serial simulation, and has no known or apparent ergodicity problems. In the previous section, we discussed the importance of regular halo updates to ensure consistency of parallel simulations with serial ones. The method described here is another way of ensuring consistent parallel code for diffusion simulations; by regularly updating the domain boundaries we ensure there are no boundary effects across the grid.

In the left panel of Figure 6.3, there is a cluster which spans domains 2 and 4, and would be rejected for movement on the basis that it was not wholly contained on a single processor. Following the reallocation of domain boundaries, as in the right hand panel, the cluster is now wholly contained within a single subdomain and is free to be moved. Regular reallocation of the domains in this manner ensures all clusters are uniformly sampled for movement.

The SL algorithm with cluster moves is described in Algorithm 3. Note that we still have a limitation that clusters bigger than a subdomain can never be moved, and that clusters of size similar to a subdomain will be sampled less frequently than substantially smaller clusters as it is less likely to select domain boundaries which completely enclose the larger cluster. However, such clusters are not relevant here firstly due to the imposition of a maximum cluster size of 100, with typical subdomain sizes of over 4000 sites, and secondly as we do not simulate for long enough for clusters to reach sizes similar to a subdomain.

6.3 Testing and performance

In order to determine if parallel simulations performed using the cluster diffusion algorithm described in this chapter produce trajectories indistinguishable from their serial counterparts we ran simulations of graphene growth with cluster moves and monomer diffusion in different proportions at different grid sizes. From these simulations we computed the mean island size for each simulation run, $\bar{S}_i(t)$, the ensemble average island size as a function of time, $\bar{S}(t)$ and the autocorrelation function of \bar{S} , $R(\tau)$ for time gap τ

$$R(\tau) = \frac{\text{E} [(\bar{S}_i(t) - \bar{S}(t)) ((\bar{S}_i(t + \tau) - \bar{S}(t + \tau)))]}{\sigma_t \sigma_{t+\tau}}. \quad (6.1)$$

We define the probabilities of selecting an insertion, diffusion or cluster move as p_I , p_D and p_C respectively. Noting that these probabilities sum to 1 and that in simulations of growth we expect there to be high levels of diffusion for each particle inserted and lower diffusion probabilities for large clusters, we set the monomer diffusion attempt prob-

ability $p_D = 1 - p_I - p_C$. In Figure 6.4 is plotted the mean island size in the early stages of simulations of growth with $p_I = 10^{-5}$ and $p_C = 0.01$ across different grid lengths and processor numbers. We observe $\bar{S}(t)$ which appear indistinguishable between serial and parallel simulations except on the smallest grid length simulated, $N = 64^2$, where the parallel simulation deviates from its serial counterpart. At this value of N the subdomains would only have size 16^2 . Clusters grown in these simulations are moved up to size $n_{\max} = 100$, which corresponds to moving a cluster roughly half the total size of the subdomain. We attribute this deviation of parallel simulations on this gridsize to this similarity in size of clusters to the subdomain, and note that when used for real growth care must be taken to ensure $n_{\max} \ll N/p^2$.

We also compute the autocorrelation function $R(\tau)$ of $\bar{S}(t)$ for simulations in serial and parallel at $N = 128$, shown in Figure 6.5. As the time gap increases the quality of the statistics get worse, but at low time gap τ we see very little deviation between the two autocorrelation functions. This is a good indication of the indistinguishability of parallel and serial simulations and is often used as such [58].

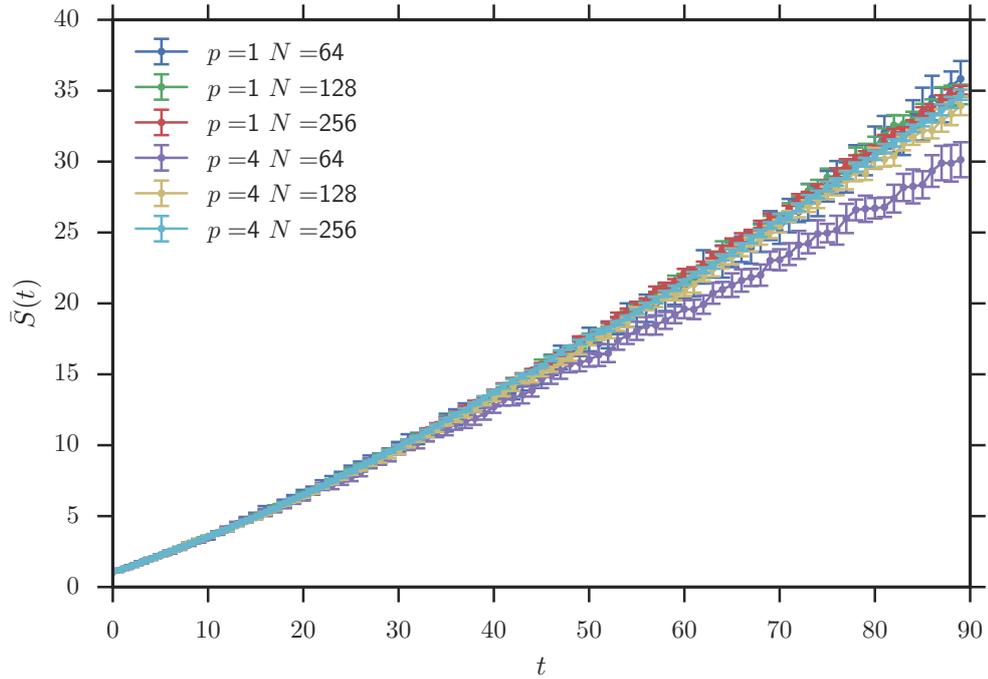


Figure 6.4 Mean island size evolution for different processor number p and grid length N (total particle number N^2). Simulations were performed with insertion probability $p_I = 10^{-5}$ and cluster move probability $p_C = 0.01$. Time is in 100 MC sweeps. The lower grid size of $N = 64$ sees a deviation of parallel and serial trajectories, as the cluster move algorithm is now selecting clusters close to the size of a subdomain. Otherwise parallel results are indistinguishable from serial simulations.

In order to test the speed and performance of this algorithm, we present results from timed simulations under two regimes across a range of grid length N (so total particle number is N^2) and processor number p . Firstly, using an implementation of Algorithm 1 with only particle insertion, equivalent to the standard Ising model. Secondly, with the above ratios of diffusion and cluster moves. These timings, per system sweep, are presented in Figures 6.6 and 6.7 respectively. In the first case we observe good performance with increasing p , with continuous decreases in time taken. However, we observe substantially longer simulations with the introduction of diffusion and cluster moves. Whilst there is an initial decrease in simulation time with increasing p , as p increases further there is an increase in simulation time.

This was far from ideal, and in order to investigate the reasons behind the effect we isolated the portion of the code performing MC moves (discussed in the test below as *moves*) and the overheads incurred from processor communication (*overheads*). Figure 6.7 shows the time taken per sweep of the system, for the moves and the overheads. The overheads increase with increasing p , as more processor to processor communications are required. The moves decrease with increasing p , as the subdomain size diminishes. The

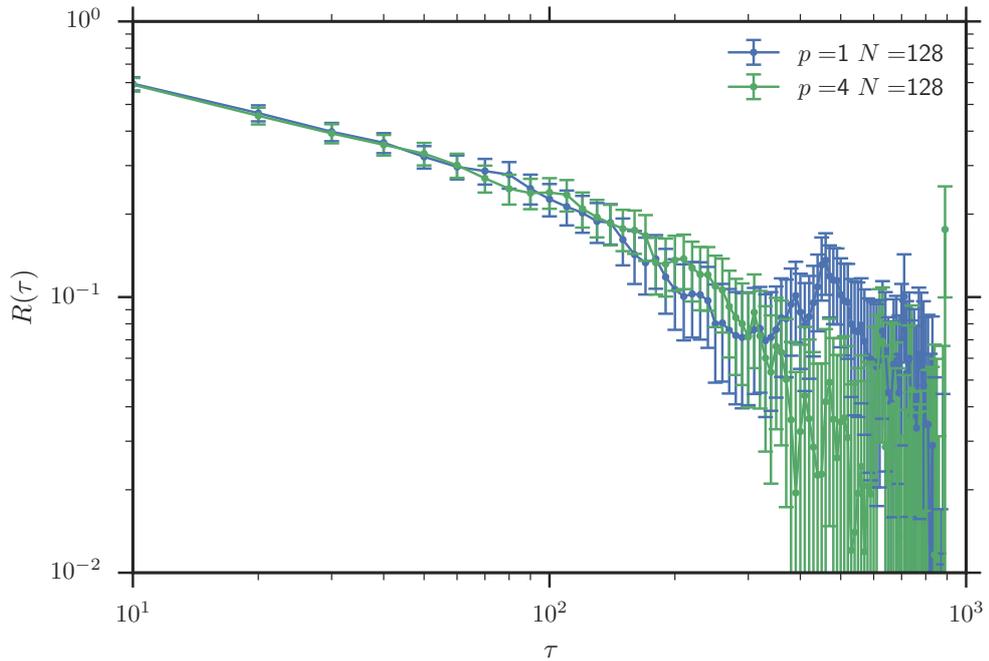


Figure 6.5 Autocorrelation function of $\bar{S}(t)$, $R(\tau)$ for time lag τ . Simulations were performed with insertion probability $p_I = 10^{-5}$, cluster move probability $p_C = 0.01$ and grid length $N = 128$ for one and four processors. The time lag τ is in MC sweeps. There is good agreement between the two functions at low τ , but at higher τ there is more uncertainty in the data, and there is a divergence between the two.

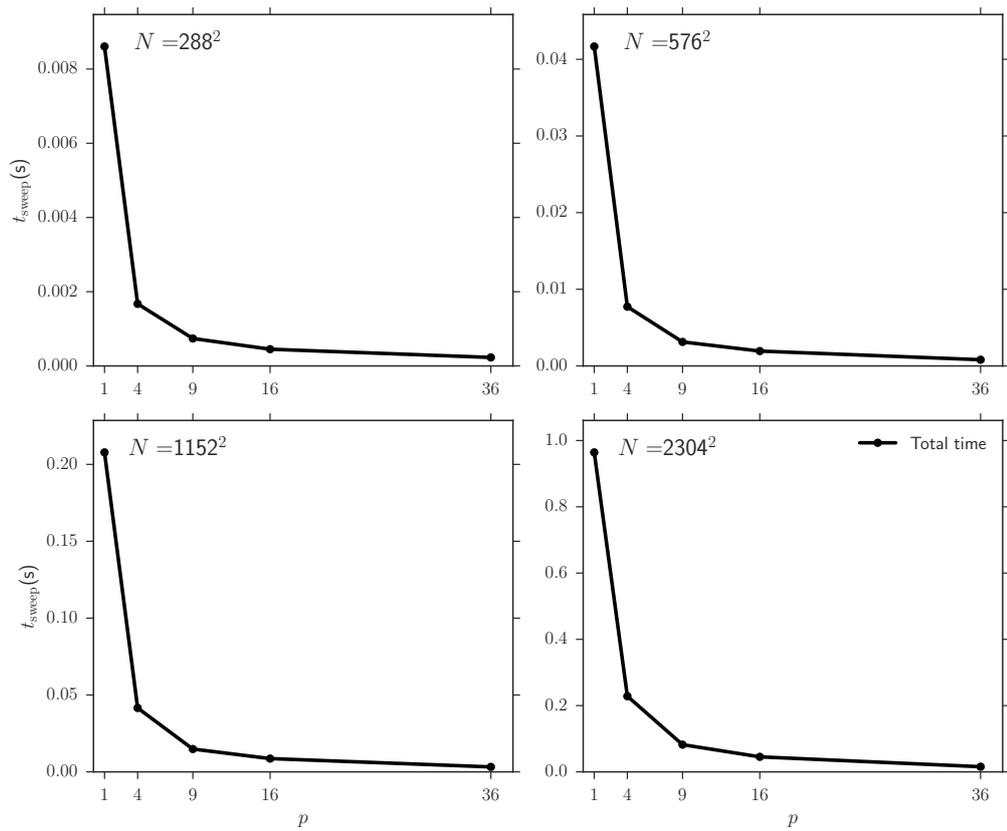


Figure 6.6 Timings for different gridsize and processor number for a simulation of the Ising model using the domain decomposition scheme described in Algorithm 1. Timings are obtained for different grid sizes and processor number, as labelled. Error bars are of the order of the point size.

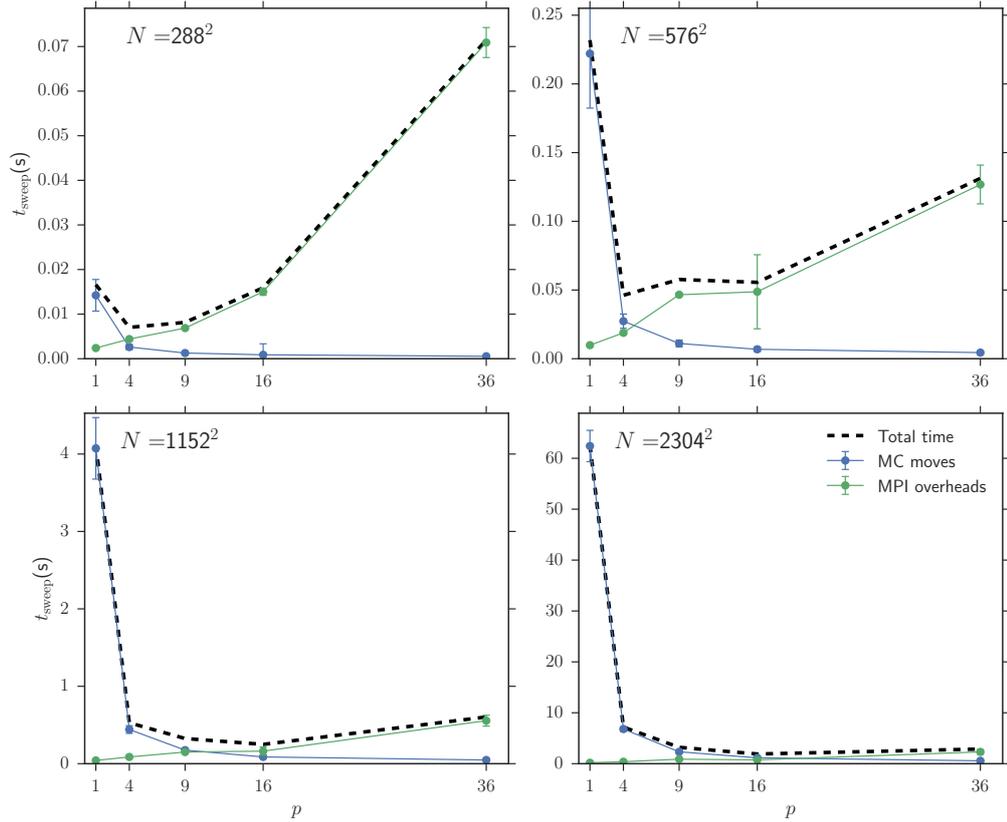


Figure 6.7 Timings for different gridsize and processor number for a simulation using insertion, diffusion and cluster moves with probabilities $p_I = 10^{-5}$, $p_C = 0.01$ and p_D obtained from subtracting these two probabilities from unity. Timings are obtained for different gridsizes and processor number, as labelled. The time simulations take for MPI moves and MC overheads are shown, as well as the total time per MC sweep. Different gridsizes have a different optimal number of processors, the minimum of the total time.

balance between these two processes leads to an optimal p at each N , which increases with N .

In processor-to-processor communications, the biggest contributor to efficiency and scalability is the number of communications, rather than the message size. The number of simultaneous communications in a halo swapping method is always 4 when exchanging for a single layer halo, or 8 when also exchanging corners from diagonal neighbours, regardless of p . A naive gather and scatter operation would scale with p , but a more intelligent tree-like model could reduce this to scaling with $\log p$. We observe scaling with p in this code, which is likely because the implementation of the scatter or gather routines is not optimised for small p .

The brief investigation of the performance and efficiency of the parallel cluster move algorithm in this section has shown that Algorithm 3 produces results indistinguishable from serial simulations, provided the size of clusters being moved does not approach the subdomain size. We have also observed substantial decreases in simulation time going from serial to small p processors, and note that as we increase N , the optimal number of processors for simulations also increases.

6.4 A lattice Monte Carlo graphene growth model with cluster moves

6.4.1 Simulation protocol

We now present results from graphene growth simulations using the cluster moving algorithm described in this chapter. This work is an extension of the model described in Chapter 3, with no substrate roughness, and the inclusion of cluster moves. The Hamiltonian

$$H_0 = \sum_{\langle i,j \rangle} E_{CC} s_i s_j + \sum_i \delta_{s_i,0} \min_j (E_{CH} \delta_{s_j,1}), \quad (6.2)$$

favours the stronger C–C bond over the C–H bond. The other key features of this model are as follows:

1. The model was on a honeycomb lattice, representing the sites of a perfect graphene sheet.
2. Each lattice site could be carbon or hydrogen, with hydrogen assumed to be in excess, with correspondingly lower C–H than C–C energies.
3. The possible MC moves were particle insertion and diffusion to first second and third near neighbours in near hexagons (1NN, 2NN and 3NN respectively).

4. A chemical potential difference ($\mu_C > \mu_H$) ensured that insertion of carbon was almost always accepted.
5. Growth via carbon insertion up to a fixed coverage, θ , followed by an extended annealing phase with no further insertion.

The simulations in this chapter introduce an additional parameter: the probability of selecting a cluster diffusion move p_C . We perform simulations at three values of p_C : $p_C = 0.001$, $p_C = 0.005$ and $p_C = 0.01$. The results from Chapter 3 without cluster moves ($p_C = 0$) are also included for comparison.

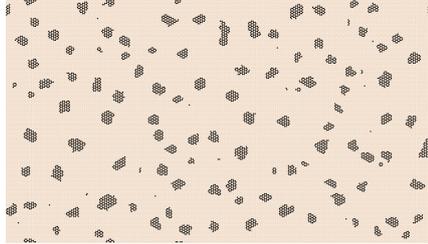
6.4.2 Results

Snapshots from different stages of growth are shown in Figure 6.8, for $p_C = 0$ and $p_C = 0.01$. There are also representative videos of growth at these move attempt frequencies included on the supplementary CD and described in Appendix A, where the $p_C = 0$ video is the case without surface roughness from Chapter 3, Video 3.1 and the $p_C = 0.01$ case Video 6.1. The mean island size evolution is plotted in Figure 6.9. In this figure the simulations which include cluster moves at any ratio have a higher mean island size, \bar{S} during growth. There is a period of island size increase at the start of annealing, where small clusters are highly mobile due to the lack of newly inserted monomers blocking cluster movement. There is also a progressively higher final \bar{S} for higher p_C . The inclusion of cluster moves then results in substantially larger islands, for example at $p_C = 0.01$, twice as large at the end of annealing.

In Chapter 3, we introduced a dynamic scaling relation obeyed by island size distributions (ISDs)

$$N_S = \frac{\theta}{\bar{S}^2} f(S/\bar{S}), \quad (6.3)$$

where N_S is the unscaled ISD, θ the coverage, \bar{S} the average island size, and f a dimensionless scaling function. This scaling distribution has been widely observed and investigated in both experimental and computational works [2, 98]. The top four panels of Figure 6.10 show the unscaled ISDs, N_S for the four values of p_C . Notably, the ISD for $\theta = 0.1$ increasingly deviates from the standard form of ISD curve with increasing p_C . The cluster moving algorithm can move islands up to size 100, and at the end of the short annealing phase used in these simulations there are still substantial quantities of smaller islands in the $\theta = 0.1$ ISD. Also shown is a second run of the $\theta = 0.1$, marked long, which is annealed for an additional 10,000 MC sweeps. After this longer anneal this ISD has the same form as the others. A video of growth at $\theta = 0.1$ is included on the supplementary CD, and described in Appendix A. The clusters move for longer than for $\theta = 0.3$, continuing to diffuse after



$p_C = 0, t = 100$



$p_C = 0.01, t = 100$



$p_C = 0, t = 200$



$p_C = 0.01, t = 200$



$p_C = 0, t = 350$



$p_C = 0.01, t = 350$



$p_C = 0, t = 400$



$p_C = 0.01, t = 400$

Figure 6.8 Snapshots from simulations of growth at 4 timesteps (time in 100 MC sweeps), for cluster diffusion to atomic diffusion ratios $p_C = 0$ and $p_C = 0.01$. The final two snapshots shown are from the end of the growth and annealing stage. In the $p_C = 0.01$ snapshots there are fewer, larger islands at all times.

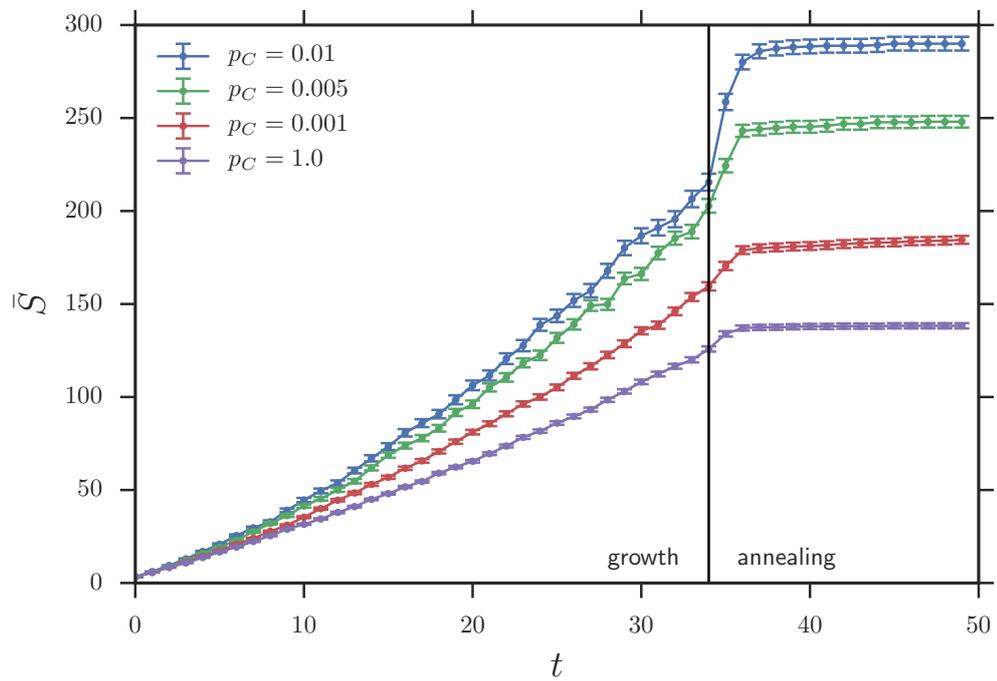


Figure 6.9 Time evolution of average island size, for four different values of the cluster diffusion to atomic diffusion attempt ratio p_C . Time is in 100 MC sweeps, and values are averaged over 20 runs. When $p_C > 0$ there is a period of rapid growth at the beginning of the annealing phase, and increasing p_C always leads to increased island size.

the end of growth.

The scaled form of these ISDs are plotted in the lower four panels of Figure 6.10. The largest θ simulations notably diverge with increasing p_C , demonstrating that the introduction of the cluster algorithm removes the scaling behaviour in the system. Deviations from scaling behaviour suggest a process in the growth which is not independent of scale [52], and as the cluster behaviour introduces fixed parameters defining the size of cluster which can be moved, the breakdown of scaling is to be expected.

6.5 Clusters, dimers and graphene growth

A notable result from the MD modelling in Chapter 5 was that a surprisingly high proportion of carbon transport was through the movement of dimers. In multiscale modelling, dimer movement was shown to be the dominant feeding species in epitaxial graphene growth, due to strong bonds and comparable diffusion rates [147]. There are also likely to be large numbers of dimers in experiment from broken down feedstock [149], although the exact composition of early stage islands remains unknown.

The simulations of graphene growth in lattice MC simulations described in this thesis could have two possible interpretations of dimer movement. The first interpretation is of monomer diffusion to 2NN as a dimer move for monomers with neighbours, with the moving atom considered to be rotated about the fixed dimer element. As all other possible moves of a monomer in a dimer would be rejected due to the breakage of the C–C bond, the diffusion of monomers up to 3NN could thus be considered equivalent to a dimer move. In Chapter 3, we remarked that with only 1NN diffusion moves, only dendritic islands were formed. With the introduction of 2NN/3NN moves came regular circular or hexagonal islands. This change in island morphology can thus be connected to the mobility of dimers in real growth.

The second interpretation of a dimer move is the more literal one described in this chapter, a dimer could be moved with the cluster moving algorithm. This would in effect be a longer range dimer move, as both atoms would move to a further site. For an isolated dimer however, this would be a valid alternate way of moving. If a dimer in a cluster were selected, such as that in 6.11, the movement could lead to the formation of a vacancy within the cluster, as shown in the figure. This introduction of an unphysical process into the model must be considered carefully, but as snapshots and movies do not reveal many vacancies in graphene islands, we can conclude that there are sufficient mechanisms available to heal any such defects.

Figure 6.12 shows different cluster sizes present over the simulation timescale for $p_C = 0.01$. In contrast to the similar figure in Chapter 5 (Figure 5.5), which showed an

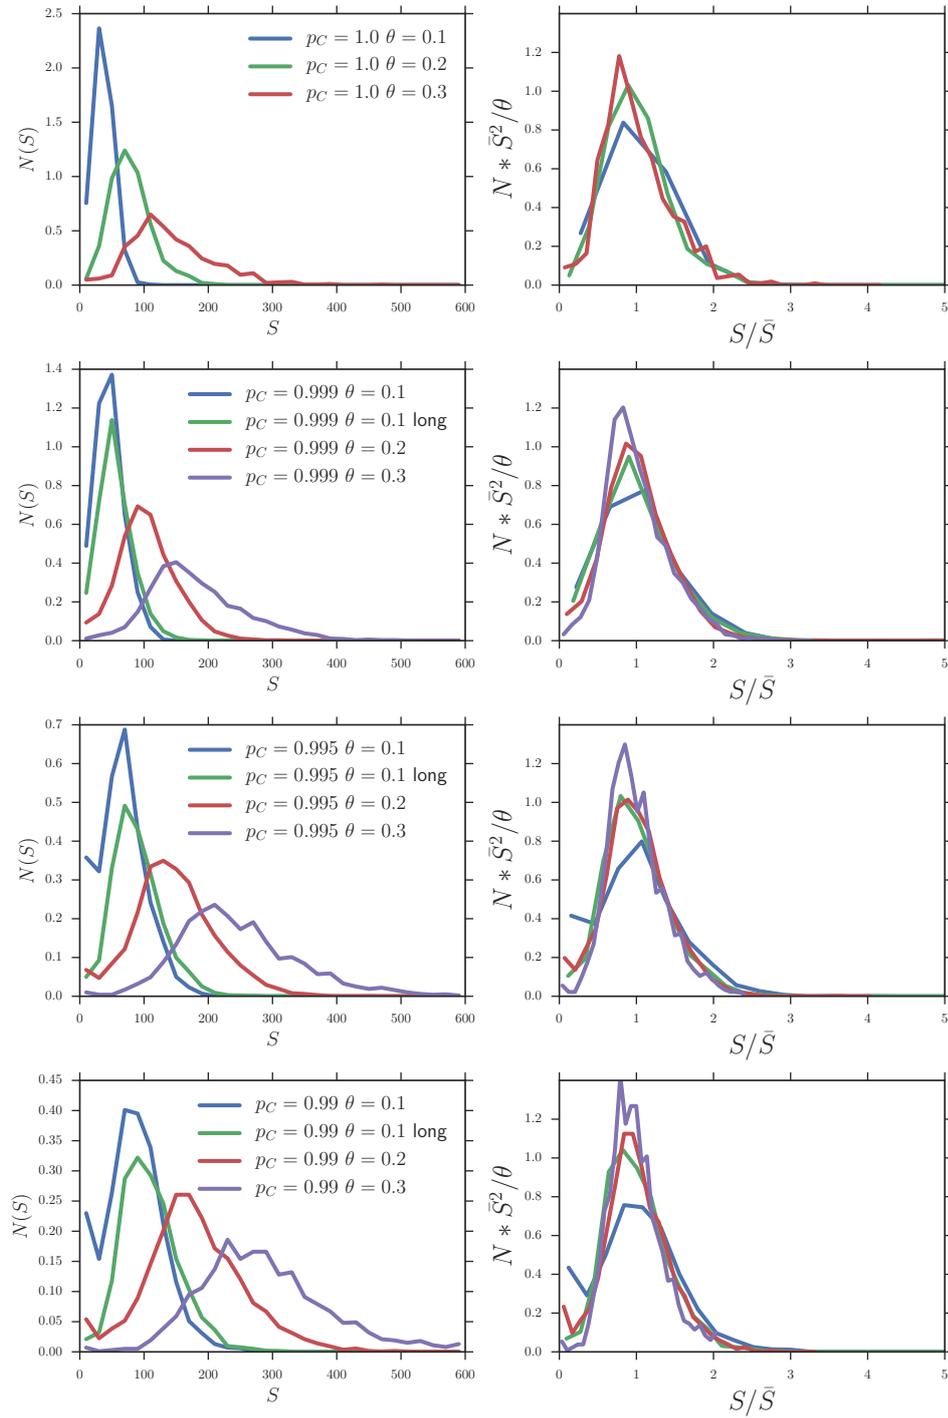


Figure 6.10 Island size distributions (ISDs), for four different cluster diffusion to atomic diffusion attempt ratios p_C , and three different coverages θ . From top to bottom, $p_C = 0, 0.001, 0.005, 0.01$. The leftmost four ISDs are the unscaled average of 50 runs, and the rightmost four rescaled according to the scaling relation described in the text. Also included when $p_C > 0$ is a run to coverage $\theta = 0.1$ with an extended annealing phase, marked long.

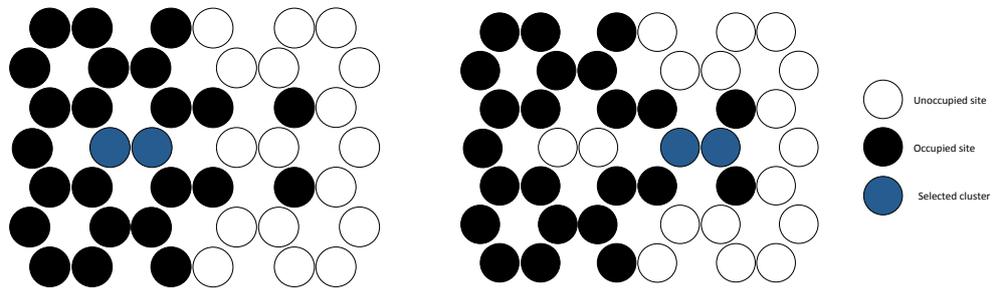


Figure 6.11 Diagram demonstrating the issue encountered with the cluster moving algorithm described when effecting a dimer move in a hexagonal geometry. The movement of the highlighted dimer would preserve number of occupied near neighbours, and would be accepted in our algorithm. It would however create a vacancy in the cluster, which is unphysical.

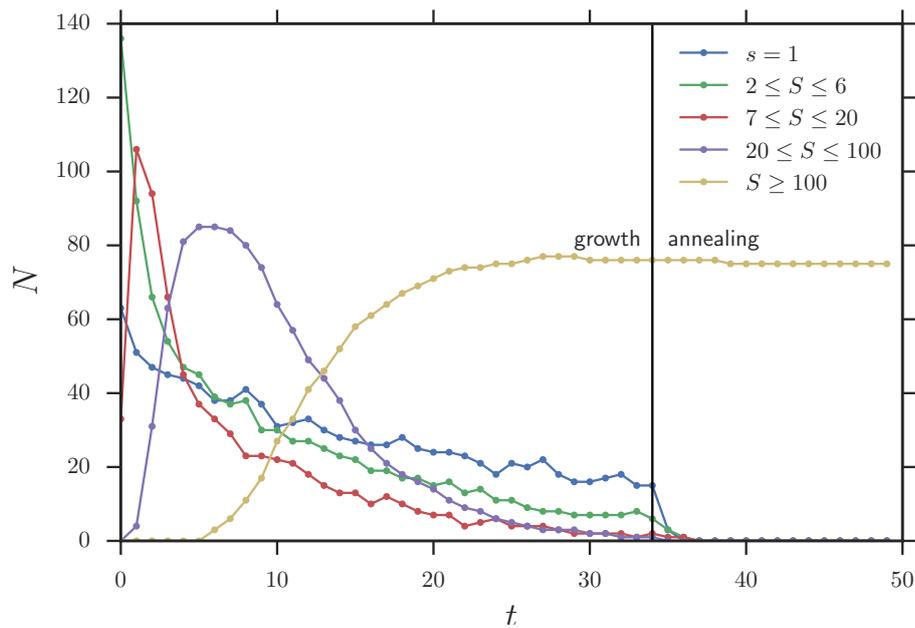


Figure 6.12 Cluster sizes present across simulation timescale, for simulations grown to coverage $\theta = 0.3$, with cluster diffusion to atomic diffusion ratio $p_C = 0.01$. Cluster sizes are grouped into appropriate categories, representing different scales of island. In early stages there is a mixture of island sizes; in the middling stages of growth there is a significant amount of the surface in islands between size 20 and 100; whilst by the end of the simulation all remaining islands are a size of at least 100, the minimum cluster size moved. The y axis is the number observed at each timestep, summed over 20 simulations.

overwhelming majority of clusters as dimers which eventually nucleated into islands, this simulation has a wide array of island sizes throughout. The energetics included in this lattice MC modelling do not capture the nucleation in the MD modelling.

In DFT simulations [36], and STM of low temperature CVD growth [92] the stability of dimers relative to monomers, existence of small clusters, and form of nucleation clusters to be important to understanding the nucleation process. Lattice based simulation can never truly capture the full mechanism behind such processes as only a carbon 6–ring can be formed. The observations of triangles, squares and pentagons of carbon atoms in the STM, and of deviations from the 6-rings in medium sized clusters in DFT indicate the importance of the incorporation of such non-hexagonal geometry. Our energetics also fail to capture the incubation period between adding carbon to a system and nucleation, observed in the MD simulations presented in chapter 5 and in real experimental procedures [18, 137]. Implementing dimer and cluster moves in the manner described in this chapter does nevertheless improve the accuracy and relevance of MC modelling to real graphene growth.

6.6 Conclusion

In this chapter we have introduced a parallel algorithm for the movement of clusters in 2D lattice MC simulations. We identified the problem involved with extending a serial cluster movement algorithm [139, 43] to parallel code, namely that clusters which span multiple processors cannot be sampled with the same probability as clusters which are entirely contained on a single processor. We proposed and implemented a solution adapted from off–lattice polymer simulations [125, 108] by regularly redefining the processor domains, allowing for uniform sampling of clusters.

We applied this algorithm to simulations of graphene growth, which demonstrated an increase in island size with increased frequency of cluster moves, and an extended period of growth for low coverage. We also demonstrated that introducing a scale–dependent event into the model broke down the scale–independent form of the island size distribution. Finally, we discussed the implementations of this model, and of MC modelling in general, to real graphene growth—noting the importance of dimers and long–range carbon moves to a regular island morphology, and the difficulties in capturing all experimental features in simulations of graphene growth. Nevertheless, this cluster move scheme represents state–of–the –art in lattice based MC modelling, and makes a significant contribution to understanding systems of surface growth.

Chapter 7

Conclusion

In this thesis we have presented results from abstract models of surface growth, aiming to capture different aspects of real growth systems with each model. We have focussed much of the discussion on graphene grown via CVD on a copper substrate, motivated by a breadth of interesting results which we discussed in Chapter 2. Better understanding of the processes governing growth is necessary for CVD-grown graphene to achieve the outstanding properties of exfoliated graphene, such as carrier mobility, and to allow large-area processing of graphene for device applications [110]. All of our results are predicated on abstract rather than system-specific models, meaning that they are broadly applicable to the increasingly important field of 2D and quasi-2D materials epitaxy [144, 40, 5].

In Chapter 3 we discussed results from MC simulations in the semi-grand canonical ensemble in which, prompted by graphene growth on copper substrates which were nearly molten[95] or liquid[39], we introduced a random site by site roughness energy. We kept this energy constant, or allowed it to diffuse locally in simulations of *static* and *dynamic* roughness. The introduction of static roughness was shown to decrease the graphene island size with increasing roughness energy. The dynamic roughness, with a high decoupled substrate temperature, led to an enhancement of graphene island size at critical roughness energy, and we speculated that this dynamic roughness gave graphene islands a new mechanism for healing. We generalised this result to the RFIM, which has a rich history of study [4, 115] in Chapter 4. By examining models of quenching we showed the generality of the dynamic roughness effect discovered in Chapter 3. We discussed the analogue between magnetisation reversal in the RFIM and nucleation on a substrate with static defects, showing that a more rough field or substrate resulted in faster nucleation.

We then introduced a 2D MD model in Chapter 5, which represented C–C bonds via a recently released bond order potential optimised for growing graphene structures far from typically parametrised geometries [153], on an effective copper substrate potential

determined by atomic positions in copper facets. In this model we captured the epitaxy between the substrate and the graphene layer well, reproducing experimental observations of different growth rates on different copper facets and orientations determined by the geometry of the substrate [145, 141]. We further investigated this system by calculating diffusion rates and flake strain energies, connecting the lower nucleation density on lower index substrates to the lower dimer diffusion rates and island energy gain from the effective substrate potential.

Finally, following observations of mass transport through dimer diffusion in Chapter 5, and observations of cluster diffusion and movement in the same chapter as well as in DFT simulations [36, 147] and low temperature CVD [92], we introduced cluster diffusion into the MC model outlined in Chapter 3 in Chapter 6. We discussed the implementation of such a move scheme in scalable parallel MC simulations. We observed an increase in island size and the breakdown of a dynamic scaling relation via the introduction of the scale dependent cluster moves. This state of the art technical extension to our MC modelling was a significant step towards more realistic simulation of the early stages of growth. We discussed the next steps for this work, of integrating a more complex Hamiltonian which could take into account a wider range of C–C bond energetics.

Throughout this thesis we have discussed some of the next steps which could be taken to improve the modelling techniques discussed here. Extensions to the model of substrate roughness in Chapter 3 could incorporate a spatially correlated rough mesh, perhaps corresponding to the potentials described in Chapter 5, with dynamically rough deviations. Whilst it would not be possible to investigate orientation on a completely static lattice, investigation of a substrate with a good lattice match to graphene could provide insights into growth. The cluster kinetics discussed in Chapter 6 could be introduced into the model of roughness, along with a Hamiltonian which better represented the energetics of C–C bonding. The potential application for the method outlined in Chapter 5 is significant, with the possibility of the first simulations on large scale polycrystalline foils with realistic graphene growth potentials.

Each of the models we have used in this thesis has been able to capture a different aspect of growth, which has been experimentally determined as important. It is currently impossible to have everything in any single model: bigger systems inevitably mean a trade off in computational efficiency or detail. By selecting a single aspect of an experimental system—roughness in Chapter 3 or geometric influence on growth in Chapter 5—we have been able to explore the specific effects of that aspect, and offer insight into related work. There are aspects of graphene growth which remain unaddressed by the work in this thesis and which are hard to model with strictly 2D models, such as wrinkling [69], structural feedback [141], copper terraces [103] and the importance of small domed islands in early

stages of growth [36]. Nevertheless, the results in this thesis provide insight into experimental findings, elucidate the mechanisms which lead to the formation of low defect graphene sheets and open up new avenues for computational modelling.

Appendix A

Supplementary videos

Videos are supplied alongside this thesis, on the provided CD. These videos will eventually be hosted in the Warwick Research Archive Portal (WRAP). Each of the four results chapters have accompanying movies of representative growth, and monomer and dimer diffusion for Chapter 5. Videos for each chapter are contained in the folders labelled on the CD, and are labelled by chapter below, along with a brief summary of their content. Video file formats are m4v and mp4, and should play on most modern video players such as QuickTime, VLC or Windows Media Player.

Video 3.1 Filename `3.1_no-noise.mp4`

Video of MC graphene growth on a substrate with no roughness. Simulation parameters are given in Table 3.1. Carbon atoms are added to the system until the target coverage has been reached, after which insertion moves are no longer attempted, and only diffusion is permitted. Note atoms are never removed from clusters, and clusters with no dangling bonds are stable. Once all atoms are added to the system, there is little change in island size during annealing.

Video 3.2 Filename `3.2_static.mp4`

Video of MC graphene growth on a substrate with static roughness, $\xi = 1.2$. Simulation parameters are given in Table 3.1. Carbon atoms are added to the system until the target coverage has been reached, after which insertion moves are no longer attempted, and only diffusion is permitted. The Voronoi cells surrounding each lattice site correspond to the strength of the roughness at that site, with darker shades corresponding to more negative roughness energies. Note atoms diffuse until reaching local minima and islands tend to be smaller. As the simulation progresses, diffusion is suppressed, and islands remain small.

Video 3.3 Filename `3.3_dynamic.mp4`

Video of MC graphene growth on a substrate with dynamic roughness, $\xi = 1.2$, $T_S = \infty$. Simulation parameters are given in Table 3.1. Carbon atoms are added to the system until the target coverage has been reached, after which insertion moves are no longer attempted, and only diffusion is permitted. The Voronoi cells surrounding each lattice site correspond to the strength of the roughness at that site, with darker shades corresponding to more negative roughness energies. Note islands are no longer stable if they have no dangling bonds, with a high degree of cluster movement and cleaving throughout the annealing phase, which leads to eventual larger island size.

Video 4.1 Filename 4.1_Ising.mp4

Video of an MC Ising model simulation of growth as quenching from above the Curie temperature. The blue pixels correspond to down spins (-1) and the yellow to up spins ($+1$). The simulation is initialised from a lattice of up and down spins selected with equal probability at each site.

Video 4.2 Filename 4.2_rfim_xi2.0.mp4

Video of a RFIM simulation of growth as quenching from above the Curie temperature with random field energy $\xi = 2.0$. The blue pixels correspond to down spins (-1) and the yellow to up spins ($+1$). The simulation is initialised from a lattice of up and down spins selected with equal probability at each site, and a random field at each site selected from a uniform probability distribution of width ξ . Note that in contrast to Video 4.1, domains are pinned at latter stages of the simulation and evolve very slowly as domain boundaries form in local minima of the field.

Video 4.3 Filename 4.3_dynamic_rfim_xi2.0_T0.5.mp4

Video of a RFIM simulation of growth as quenching from above the Curie temperature, with random field energy $\xi = 2.0$, substrate temperature $T_S = 0.5$. The video shows two lattices side by side, the left is the spin lattice, and the right is the field lattice. The blue pixels correspond to down spins (-1) and the yellow to up spins ($+1$) in the spin lattice, and blue spins to positive sites, and yellow to negative sites in the field lattice. The simulation is initialised from a lattice of up and down spins selected with equal probability at each site, and a random field at each site selected from a uniform probability distribution of width ξ . In this simulation, the field lattice and the spin lattice align together, increasing the pinning effect, but smoothing the domain boundaries in comparison to Video 4.2.

Video 4.4 Filename 4.4_dynamic_rfim_xi2.0_T10.0.mp4

Video of a RFIM simulation of growth as quenching from above the Curie temperature with random field energy $\xi = 2.0$, substrate temperature $T_S = 10.0$. The video

shows two lattices side by side, the left is the spin lattice, and the right is the field lattice. The blue pixels correspond to down spins (-1) and the yellow to up spins ($+1$) in the spin lattice, and blue spins to positive sites, and yellow to negative sites in the field lattice. The simulation is initialised from a lattice of up and down spins selected with equal probability at each site, and a random field at each site selected from a uniform probability distribution of width ξ . In this simulation, the field lattice and the spin lattice do not align together, with the field lattice constantly shifting throughout the simulation. This leads to an alignment with a single spin, and shrinking of domains.

Video 5.1 Filename 5.1_graphene_on_CU100.mp4

Video from a simulation of graphene growth on an effective substrate potential, with atomic positions according to copper positions in Cu(100). Copper atom positions are not shown. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Carbon atoms are inserted into the system up until a fixed coverage, after which only atomic diffusion is permitted.

Video 5.2 Filename 5.2_graphene_on_CU111.mp4

Video from a simulation of graphene growth on an effective substrate potential, with atomic positions according to copper positions in Cu(111). Copper atom positions are not shown. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Carbon atoms are inserted into the system up until a fixed coverage, after which only atomic diffusion is permitted.

Video 5.3 Filename 5.3_graphene_on_CU110.mp4

Video from a simulation of graphene growth on an effective substrate potential, with atomic positions according to copper positions in Cu(110). Copper atom positions are not shown. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Carbon atoms are inserted into the system up until a fixed coverage, after which only atomic diffusion is permitted.

Video 5.4 Filename 5.4_graphene_on_CU210.mp4

Video from a simulation of graphene growth on an effective substrate potential, with atomic positions according to copper positions in Cu(210). Copper atom positions are

not shown. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Carbon atoms are inserted into the system up until a fixed coverage, after which only atomic diffusion is permitted.

Video 5.5 Filename 5.5_cluster_joining_Cu100.mp4

Video of several small clusters nucleating near each other and joining in simulations of graphene growth on an effective Cu(100) potential. Bonds are coloured according to which of the two most commonly observed orientations they are closest to, for a more detailed explanation of this colouring scheme see the text. Note that boundaries of 57 defect rings form along the join between the different orientation clusters.

Video 5.6 Filename 5.6_diffusion.mpg

Video of a dimer and a monomer diffusing on a Cu(100) effective substrate potential. The black circles represent carbon atoms, and the period lattice of atoms behind shows the positions of the wells in the effective substrate potential. The trajectories for the carbon atoms are obtained from separate simulations, but are shown on the same lattice for side-by-side comparison. Note the monomer tends to be trapped in a particular well then undergo a long distance diffusion to a different well, whilst the dimer tends to hop between adjacent sites.

Video 6.1 Filename 6.1_cluster_moves_p0.01_theta0.3.mp4

Video of graphene growth in a lattice Monte Carlo model, with cluster diffusion moves selected with a probability of 0.01, grown up to a coverage of 30%. Carbon atoms are inserted into the system until it reaches a fixed coverage, at which point insertion moves are disabled and only diffusion is permitted. Clusters diffuse a lot in the early stages of growth, and lead to islands reshaping in later stages of growth.

Video 6.2 Filename 6.2_cluster_moves_p0.01_theta0.1.mp4

Video of graphene growth in a lattice Monte Carlo model, with cluster diffusion moves selected with a probability of 0.01, grown up to a coverage of 10%. Carbon atoms are inserted into the system until it reaches a fixed coverage, at which point insertion moves are disabled and only diffusion is permitted. Clusters diffuse a lot in the early stages of growth, and as the cluster size at the end of the growth stage is still lower than the maximum cluster size, continue to move during the annealing phase.

Bibliography

- [1] Aizenman, M. and J. Wehr
1989. Rounding of first-order phase transitions in systems with quenched disorder. *Physical review letters*, 62(21):2503.
- [2] Amar, J. G. and F. Family
1995. Critical cluster size: Island morphology and size distribution in submonolayer epitaxial growth. *Physical Review Letters*, 74(11):2066.
- [3] Anderson, S. R.
1987a. Growth and equilibration in the two-dimensional random-field Ising model. *Physical Review B*, 36:8435–8446.
- [4] Anderson, S. R.
1987b. Growth and equilibration in the two-dimensional random-field Ising model. *Physical Review B*, 36(16):8435.
- [5] Bakti Utama, M. I., Q. Zhang, J. Zhang, Y. Yuan, F. J. Belarre, J. Arbiol, and Q. Xiong
2013. Recent developments and future directions in the growth of nanostructures by van der waals epitaxy. *Nanoscale*, 5:3570–3588.
- [6] Banhart, F., J. Kotakoski, and A. V. Krasheninnikov
2011. Structural Defects in Graphene. *ACS Nano*, 5(1):26–41.
- [7] Barkema, G. and T. MacFarland
1994. Parallel simulation of the Ising model. *Physical Review E*, 50(2):1623.
- [8] Basham, M., F. Montalenti, and P. A. Mulheran
2006. Multiscale modeling of island nucleation and growth during Cu(100) homoepitaxy. *Physical Review B*, 73(4):045422.
- [9] Battaile, C. C. and D. J. Srolovitz
2002. Kinetic Monte Carlo simulation of chemical vapor deposition. *Annual Review of Material Research*, 32(1):297–319.

- [10] Bell, G. R., P. M. Dawson, P. A. Pandey, N. R. Wilson, and P. A. Mulheran
2014. Size-dependent mobility of gold nano-clusters during growth on chemically modified graphene. *APL materials*, 2(1):012109.
- [11] Bell, G. R., T. J. Krzyzewski, P. B. Joyce, and T. S. Jones
2000. Island size scaling for submonolayer growth of InAs on GaAs(001)-(2×4): Strain and surface reconstruction effects. *Physical Review B*, 61(16):10551.
- [12] Berry, J., N. Provatas, J. Rottler, and C. W. Sinclair
2012. Defect stability in phase-field crystal models: Stacking faults and partial dislocations. *Physical Review B*, 86(22):224112.
- [13] Bolkhovityanov, Y. B. and O. P. Pchelyakov
2008. GaAs epitaxy on Si substrates: modern status of research and engineering. *Physics-Uspokhi*, 51(5):437–456.
- [14] Bray, A. J.
2002. Theory of phase-ordering kinetics. *Advances in Physics*, 51(2):481–587.
- [15] Brenner, D. W.
1990. Empirical potential for hydrocarbons for use in simulating the chemical vapor deposition of diamond films. *Physical Review B*, 42(15):9458.
- [16] Brenner, D. W., O. A. Shenderova, J. A. Harrison, S. J. Stuart, B. Ni, and S. B. Sinnott
2002. A second-generation reactive empirical bond order (REBO) potential energy expression for hydrocarbons. *Journal of Physics: Condensed Matter*, 14(4):783.
- [17] Brune, H., M. Giovannini, K. Bromann, and K. Kern
1998. Self-organized growth of nanostructure arrays on strain-relief patterns. *Nature*, 394(6692):451–453.
- [18] Celebi, K., M. T. Cole, J. W. Choi, F. Wyczisk, P. Legagneux, N. Rupesinghe, J. Robertson, K. B. K. Teo, and H. G. Park
2013. Evolutionary Kinetics of Graphene Formation on Copper. *Nano Letters*, 13(3):967–974.
- [19] Chandler, D.
1987. Introduction to modern statistical mechanics. *Introduction to Modern Statistical Mechanics*, by David Chandler, pp. 288. Foreword by David Chandler. Oxford University Press, Sep 1987. ISBN-10: 0195042778. ISBN-13: 9780195042771, P. 288.

- [20] Chen, H., W. Zhu, and Z. Zhang
2010. Contrasting Behavior of Carbon Nucleation in the Initial Stages of Graphene Epitaxial Growth on Stepped Metal Surfaces. *Physical Review Letters*, 104:186101.
- [21] Chen, L.-Q.
2002. Phase-field models for microstructure evolution. *Annual Review of Materials Research*, 32(1):113–140.
- [22] Cushing, G. W., V. Johaneck, J. K. Navin, and I. Harrison
2015. Graphene growth on pt (111) by ethylene chemical vapor deposition at surface temperatures near 1000 k. *The Journal of Physical Chemistry C*, 119(9):4759–4768.
- [23] Ding, G., Y. Zhu, S. Wang, Q. Gong, L. Sun, T. Wu, X. Xie, and M. Jiang
2013. Chemical vapor deposition of graphene on liquid metal catalysts. *Carbon*, 53:321–326.
- [24] Doudevski, I. and D. K. Schwartz
1999. Dynamic scaling of the submonolayer island size distribution during self-assembled monolayer growth. *Physical Review B*, 60(1):14.
- [25] Dubrovskii, V. G. and N. V. Sibirev
2015. Analytic scaling function for island-size distributions. *Physical Review E*, 91:042408.
- [26] Eick, S. G., A. G. Greenberg, B. D. Lubachevsky, and A. Weiss
1993. Synchronous relaxation for parallel simulations with applications to circuit-switched networks. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 3(4):287–314.
- [27] Eng, P. J., P. W. Stephens, and T. Tse
1992. Anomalous power-law ordering kinetics of Pb on Ni(001). *Physical Review B*, 46:5024–5027.
- [28] Enstone, G., P. Brommer, D. Quigley, and G. R. Bell
2016. Enhancement of island size by dynamic substrate disorder in simulations of graphene growth. *Physical Chemistry Chemical Physics*, 18(22):15102–15109.
- [29] Esconjauregui, S., C. M. Whelan, and K. Maex
2009. The reasons why metals catalyze the nucleation and growth of carbon nanotubes and other carbon nanomorphologies. *Carbon*, 47(3):659–669.
- [30] Fasolino, A., J. Los, and M. I. Katsnelson
2007. Intrinsic ripples in graphene. *Nature Materials*, 6(11):858–861.

- [31] Fisher, D. S.
1986. Scaling and critical slowing down in random-field Ising systems. *Physical Review Letters*, 56:416–419.
- [32] Frenkel, D. and B. Smit
2002. Understanding molecular simulation: from algorithms to applications. *Computational Sciences Series*, 1:1–638.
- [33] Gaillard, P., J.-N. Aqua, and T. Frisch
2013. Kinetic Monte Carlo simulations of the growth of silicon germanium pyramids. *Physical Review B*, 87(12):125310.
- [34] Gaillard, P., T. Chanier, L. Henrard, P. Moskovkin, and S. Lucas
2015. Multiscale simulations of the early stages of the growth of graphene on copper. *Surface Science*, 637(0):11 – 18.
- [35] Gao, J., J. Yip, J. Zhao, B. I. Yakobson, and F. Ding
2011a. Graphene Nucleation on Transition Metal Surface: Structure Transformation and Role of the Metal Step Edge. *Journal of the American Chemical Society*, 133(13):5009–5015.
- [36] Gao, J., Q. Yuan, H. Hu, J. Zhao, and F. Ding
2011b. Formation of Carbon Clusters in the Initial Stage of Chemical Vapor Deposition Graphene Growth on Ni(111) Surface. *Journal of Physical Chemistry C*, 115(36):17695–17703.
- [37] Gao, L., J. R. Guest, and N. P. Guisinger
2010. Epitaxial Graphene on Cu(111). *Nano Letters*, 10(9):3512–3516.
- [38] Geim, A. K. and I. V. Grigorieva
2013. Van der Waals heterostructures. *Nature*, 499(7459):419–425.
- [39] Geng, D., B. Wu, Y. Guo, L. Huang, Y. Xue, J. Chen, G. Yu, L. Jiang, W. Hu, and Y. Liu
2012. Uniform hexagonal graphene flakes and films grown on liquid copper surface. *Proceedings of the National Academy of Science USA*, 109:7992–7996.
- [40] Gibb, A. L., N. Alem, J.-H. Chen, K. J. Erickson, J. Ciston, A. Gautam, M. Linck, and A. Zettl
2013. Atomic Resolution Imaging of Grain Boundary Defects in Monolayer Chemical Vapor Deposition-Grown Hexagonal Boron Nitride. *Journal American Chemical Society*, 135:6758–6761.

- [41] Giovannetti, G., P. A. Khomyakov, G. Brocks, P. J. Kelly, and J. van den Brink
2007. Substrate-induced band gap in graphene on hexagonal boron nitride: Ab initio density functional calculations. *Physical Review B*, 76:073103.
- [42] Gonçalves, L. and R. Stinchcombe
1986. Two-dimensional ising model in an annealed random field: Exact results. *Physical Review B*, 33(7):4762.
- [43] Grant, J. and R. L. Jack
2012. Quantifying reversibility in a phase-separating lattice gas: An analogy with self-assembly. *Physical Review E*, 85(2):021112.
- [44] Gropp, W., E. Lusk, N. Doss, and A. Skjellum
1996. A high-performance, portable implementation of the MPI message passing interface standard. *Parallel Computing*, 22(6):789–828.
- [45] Hellerstedt, J., M. T. Edmonds, J. H. Chen, W. G. Cullen, C. X. Zheng, and M. S. Fuhrer
2014. Thickness and growth-condition dependence of in-situ mobility and carrier density of epitaxial thin-film Bi₂Se₃. *Applied Physics Letters*, 105(17):173506.
- [46] Hohenberg, P. and W. Kohn
1964. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864.
- [47] Hopp, S. F. and A. Heuer
2010. Kinetic Monte Carlo study of nucleation processes on patterned surfaces. *Journal of Chemical Physics*, 133(20):204101.
- [48] Hoshen, J. and R. Kopelman
1976. Percolation and cluster distribution. I. Cluster multiple labeling technique and critical concentration algorithm. *Physical Review B*, 14(8):3438.
- [49] Huang, P. Y., C. S. Ruiz-Vargas, A. M. van der Zande, W. S. Whitney, M. P. Levendorf, J. W. Kevek, S. Garg, J. S. Alden, C. J. Hustedt, Y. Zhu, J. Park, P. L. McEuen, and D. A. Muller
2011. Grains and grain boundaries in single-layer graphene atomic patchwork quilts. *Nature*, 469(7330):389–392.
- [50] Ishihara, M., Y. Koga, J. Kim, K. Tsugawa, and M. Hasegawa
2011. Direct evidence of advantage of Cu(111) for graphene synthesis by using Raman mapping and electron backscatter diffraction. *Materials Letters*, 65(19):2864 – 2867.

- [51] Itoh, M., G. Bell, B. Joyce, and D. Vvedensky
2000. Transformation kinetics of homoepitaxial islands on GaAs (001). *Surface Science*, 464(2):200–210.
- [52] Itoh, M., G. R. Bell, A. R. Avery, T. S. Jones, B. A. Joyce, and D. D. Vvedensky
1998. Island Nucleation and Growth on Reconstructed GaAs(001) Surfaces. *Physical Review Letters*, 81:633–636.
- [53] Jeon, J. H., M. Song, H. Kim, W.-J. Jang, J.-Y. Park, S. Yoon, and S.-J. Kahng
2014. Quintuple layer Bi₂Se₃ thin films directly grown on insulating SiO₂ using molecular beam epitaxy. *Applied Surface Science*, 316:42 – 45.
- [54] Jones, J. E.
1924. On the determination of molecular fields. II. From the equation of state of a gas. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 106, Pp. 463–477. The Royal Society.
- [55] Joyce, B., D. Vvedensky, G. Bell, J. Belk, M. Itoh, and T. Jones
1999. Nucleation and growth mechanisms during MBE of IIIV compounds. *Materials Science Engineering B*, 67:7.
- [56] Kang, H. and W. Weinberg
1989. Dynamic Monte Carlo with a proper energy barrier: Surface diffusion and two-dimensional domain ordering. *The Journal of Chemical Physics*, 90(5):2824–2830.
- [57] Kawasaki, K.
1966. Diffusion constants near the critical point for time-dependent Ising models. I. *Physical Review*, 145(1):224.
- [58] Kelling, J., G. Ódor, and S. Gemming
2017. Dynamical universality classes of simple growth and lattice gas models. *arXiv preprint arXiv:1701.03638*.
- [59] Kidambi, P. R., B. C. Bayer, R. Blume, Z.-J. Wang, C. Baetz, R. S. Weatherup, M.-G. Willinger, R. Schloegl, and S. Hofmann
2013. Observing graphene grow: catalyst–graphene interactions during scalable graphene growth on polycrystalline copper. *Nano letters*, 13(10):4769–4778.
- [60] Kim, H., C. Mattevi, M. R. Calvo, J. C. Oberg, L. Artiglia, S. Agnoli, C. F. Hirjibehedin, M. Chhowalla, and E. Saiz
2012. Activation Energy Paths for Graphene Nucleation and Growth on Cu. *ACS Nano*, 6(4):3614–3623.

- [61] Kofke, D. A. and E. D. Glandt
1988. Monte Carlo simulation of multicomponent equilibria in a semigrand canonical ensemble. *Molecular Physics*, 64(6):1105–1131.
- [62] Kohn, W.
1999. Nobel Lecture: Electronic structure of matterwave functions and density functionals. *Reviews of Modern Physics*, 71(5):1253.
- [63] Kohn, W. and L. J. Sham
1965. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133.
- [64] Koma, A.
1992. Van der waals epitaxya new epitaxial growth method for a highly lattice-mismatched system. *Thin Solid Films*, 216(1):72–76.
- [65] Krzyzewski, T. J., P. B. Joyce, G. R. Bell, and T. S. Jones
2002. Scaling behavior in InAs/GaAs(001) quantum-dot formation. *Physical Review B*, 66:201302.
- [66] Leach, A. R.
2001. *Molecular modelling: principles and applications*. Pearson education.
- [67] Lee, J.-H., E. K. Lee, W.-J. Joo, Y. Jang, B.-S. Kim, J. Y. Lim, S.-H. Choi, S. J. Ahn, J. R. Ahn, M.-H. Park, C.-W. Yang, B. L. Choi, S.-W. Hwang, and D. Whang
2014. Wafer-Scale Growth of Single-Crystal Monolayer Graphene on Reusable Hydrogen-Terminated Germanium. *Science*, 344(6181):286–289.
- [68] Li, H.-B., A. J. Page, C. Hettich, B. Aradi, C. Kohler, T. Frauenheim, S. Irle, and K. Morokuma
2014. Graphene nucleation on a surface-molten copper catalyst: quantum chemical molecular dynamics simulations. *Chemical Science*, 5:3493–3500.
- [69] Li, X., W. Cai, J. An, S. Kim, J. Nah, D. Yang, R. Piner, A. Velamakanni, I. Jung, E. Tutuc, et al.
2009. Large-area synthesis of high-quality and uniform graphene films on copper foils. *Science*, 324(5932):1312–1314.
- [70] Li, X., L. Colombo, and R. S. Ruoff
2016. Synthesis of graphene films on copper foils by chemical vapor deposition. *Advanced Materials*, 28(29):6247–6252.

- [71] Li, X., C. W. Magnuson, A. Venugopal, J. An, J. W. Suk, B. Han, M. Borysiak, W. Cai, A. Velamakanni, Y. Zhu, et al.
2010. Graphene films with large domain size by a two-step chemical vapor deposition process. *Nano Letters*, 10(11):4328–4334.
- [72] Liu, F. and H. Metiu
1994. Stability and kinetics of step motion on crystal surfaces. *Physical Review E*, 49(4):2601.
- [73] Lloyd, H.
1974. Estimation of tumor cell kill from Gompertz growth curves. *Cancer Chemotherapy Reports. Part 1*, 59(2 Pt 1):267–277.
- [74] Losurdo, M., M. M. Giangregorio, P. Capezzuto, and G. Bruno
2011. Ellipsometry as a Real-Time Optical Tool for Monitoring and Understanding Graphene Growth on Metals. *Journal of Physical Chemistry C*, 115(44):21804–21812.
- [75] Lubachevsky, B.
1988. Complex Syst. 1, 1099 1987. *Journal of Computational Physics*, 75:103.
- [76] Luo, Z., S. Kim, N. Kawamoto, A. M. Rappe, and A. C. Johnson
2011. Growth mechanism of hexagonal-shape graphene flakes with zigzag edges. *ACS Nano*, 5(11):9154–9160.
- [77] Ma, J., D. Alfè, A. Michaelides, and E. Wang
2009. Stone-Wales defects in graphene and other planar sp^2 -bonded materials. *Physical Review B*, 80:033407.
- [78] Marsden, A. J., M.-C. Asensio, J. Avila, P. Dudin, A. Barinov, P. Moras, P. M. Sheverdyeva, T. W. White, I. Maskery, G. Costantini, N. R. Wilson, and G. R. Bell
2013a. Is graphene on copper doped? *Physica Status Solidi Rapid Research Letters*, 7(9):643–646.
- [79] Marsden, A. J., M. Phillips, and N. R. Wilson
2013b. Friction force microscopy: a simple technique for identifying graphene on rough substrates and mapping the orientation of graphene grains on copper. *Nanotechnology*, 24(25):255704.
- [80] Meca, E., J. Lowengrub, H. Kim, C. Mattevi, and V. B. Shenoy
2013. Epitaxial graphene growth and shape dynamics on copper: phase-field modeling and experiments. *Nano Letters*, 13(11):5692–5697.

- [81] Meng, L., Q. Sun, J. Wang, and F. Ding
2012. Molecular dynamics simulation of chemical vapor deposition graphene growth on ni (111) surface. *The Journal of Physical Chemistry C*, 116(10):6097–6102.
- [82] Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller
1953. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092.
- [83] Mi, X., V. Meunier, N. Koratkar, and Y. Shi
2012. Facet-insensitive graphene growth on copper. *Physical Review B*, 85:155436.
- [84] Mo, Y., K. T. Turner, and I. Szlufarska
2009. Friction laws at the nanoscale. *Nature*, 457(7233):1116–1119.
- [85] Moriarty, J. A.
1988. Density-functional formulation of the generalized pseudopotential theory. III. Transition-metal interatomic potentials. *Physical Review B*, 38(5):3199.
- [86] Mulheran, P. and J. Blackman
1997. Simulation and theory of island growth on substrate steps. *Surface Science*, 376(1-3):403410.
- [87] Nattermann, T.
1997. Theory of the random field ising model. *Spin glasses and random fields*, 12:277.
- [88] Newman, M. and G. Barkema
1996. Monte Carlo study of the random-field Ising model. *Physical Review E*, 53(1):393.
- [89] Newman, M. and G. Barkema
1999. *Monte Carlo Methods in Statistical Physics*. Oxford University Press: New York, USA.
- [90] Nie, S., J. M. Wofford, N. C. Bartelt, O. D. Dubon, and K. F. McCarty
2011a. Origin of the mosaicity in graphene grown on Cu (111). *Physical Review B*, 84(15):155425.
- [91] Nie, S., J. M. Wofford, N. C. Bartelt, O. D. Dubon, and K. F. McCarty
2011b. Origin of the mosaicity in graphene grown on Cu (111). *Physical Review B*, 84(15):155425.
- [92] Niu, T., M. Zhou, J. Zhang, Y. Feng, and W. Chen
2013. Growth intermediates for CVD graphene on Cu (111): carbon clusters and defective graphene. *Journal of The American Chemical Society*, 135(22):8409–8414.

- [93] Nowak, U.
1995. Magnetisation reversal and domain structure in thin magnetic films: theory and computer simulation. *Ieee Transactions on Magnetics*, 31(6):4169–4171.
- [94] OConnor, T. C., J. Andzelm, and M. O. Robbins
2015. AIREBO-M: A reactive model for hydrocarbons at extreme pressures. *The Journal of Chemical Physics*, 142(2):024903.
- [95] Page, A. J., Y. Wang, H.-B. Li, S. Irle, and K. Morokuma
2013. Nucleation of Graphene Precursors on Transition Metal Surfaces: Insights from Theoretical Simulations. *Journal of Physical Chemistry C*, 117(28):14858–14864.
- [96] Pandey, P. A., G. R. Bell, J. P. Rourke, A. M. Sanchez, M. D. Elkin, B. J. Hickey, and N. R. Wilson
2011. Physical Vapor Deposition of Metal Nanoparticles on Chemically Modified Graphene: Observations on Metal-Graphene Interactions. *Small*, 7(22):3202–3210.
- [97] Pettifor, D. G. and I. I. Oleinik
1999. Analytic bond-order potentials beyond Tersoff-Brenner. I. Theory. *Physical Review B*, 59:8487–8499.
- [98] Pimpinelli, A. and T. L. Einstein
2007. Capture-Zone Scaling in Island Nucleation: Universal Fluctuation Behavior. *Physical Review Letters*, 99:226102.
- [99] Plimpton, S.
1995. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics*, 117(1):1–19.
- [100] Press, W. H.
2007. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.
- [101] Pristovsek, M., R. Kremzow, and M. Kneissl
2013. Energetics of Quantum Dot Formation and Relaxation of InGaAs on GaAs(001). *Japanese Journal of Applied Physics*, 52:041201.
- [102] Pristovsek, M., B. Röhmer, M. Breusig, R. Kremzow, and W. Richter
2007. In situ scanning tunnelling microscopy during metal-organic vapour phase epitaxy. *Journal of Crystal Growth*, 298:8–11.

- [103] Rasool, H. I., E. B. Song, M. Mecklenburg, B. C. Regan, K. L. Wang, B. H. Weiller, and J. K. Gimzewski
2011. Atomic-Scale Characterization of Graphene Grown on Copper (100) Single Crystals. *Journal of The American Chemical Society*, 133(32):12536–12543.
- [104] Riikonen, S., A. Krasheninnikov, L. Halonen, and R. Nieminen
2012. The Role of Stable and mobile carbon adspecies in copper-promoted graphene Growth. *The Journal of Physical Chemistry C*, 116(9):5802–5809.
- [105] Robertson, A. W., B. Montanari, K. He, C. S. Allen, Y. A. Wu, N. M. Harrison, A. I. Kirkland, and J. H. Warner
2013. Structural Reconstruction of the Graphene Monovacancy. *ACS Nano*, 7(5):4495–4502.
- [106] Ružička, Š. and M. P. Allen
2014a. Collective translational and rotational Monte Carlo cluster move for general pairwise interaction. *Physical Review E*, 90(3):033302.
- [107] Ružička, Š. and M. P. Allen
2014b. Collective translational and rotational Monte Carlo moves for attractive particles. *Physical Review E*, 89(3):033307.
- [108] Schmid, F., D. Düchs, O. Lenz, and B. West
2007. A generic model for lipid monolayers, bilayers, and membranes. *Computer Physics Communications*, 177(1):168–171.
- [109] Schulze, T. P. and P. Smereka
2012. Kinetic Monte Carlo simulation of heteroepitaxial growth: Wetting layers, quantum dots, capping, and nanorings. *Physical Review B*, 86(23):235313.
- [110] Seah, C.-M., S.-P. Chai, and M. Abdul Rahman
2014. Mechanisms of graphene growth by chemical vapour deposition on transition metals. *Carbon*, 70:1 – 21.
- [111] Sekerka, R. F.
2004. Morphology: from sharp interface to phase field models. *Journal of Crystal Growth*, 264(4):530–540.
- [112] Seymour, M. and N. Provatas
2016. Structural phase field crystal approach for modeling graphene and other two-dimensional structures. *Physical Review B*, 93(3):035447.

- [113] Shi, Y., W. Zhou, A.-Y. Lu, W. Fang, Y.-H. Lee, A. L. Hsu, S. M. Kim, K. K. Kim, H. Y. Yang, L.-J. Li, J.-C. Idrobo, and J. Kong
2012. Van der Waals epitaxy of MoS₂ layers using graphene as growth templates. *Nano Letters*, 12:2784–2791.
- [114] Shim, Y. and J. G. Amar
2005. Semirigorous synchronous sublattice algorithm for parallel kinetic Monte Carlo simulations of thin film growth. *Physical Review B*, 71(12):125432.
- [115] Sinha, S. and P. K. Mandal
2013. Dynamical properties of random-field Ising model. *Physical Review E*, 87(2):022121.
- [116] Stanley, H. E. and N. Ostrowsky
2012. *Random fluctuations and pattern growth: experiments and models*, volume 157. Springer Science & Business Media.
- [117] Takahasi, M.
2013. X-ray Diffraction Study of Crystal Growth Dynamics during Molecular-Beam Epitaxy of III–V Semiconductors. *Journal of The Physical Society of Japan*, 82(2):021011.
- [118] Tarasenko, A. and L. Jastrabik
2011. Diffusion of particles over dynamically disordered lattice. *Physical chemistry chemical physics*, 13:2300–2306.
- [119] Tersoff, J.
1986. New empirical model for the structural properties of silicon. *Physical Review Letters*, 56(6):632.
- [120] Tetlow, H., J. P. de Boer, I. Ford, D. Vvedensky, J. Coraux, and L. Kantorovich
2014. Growth of epitaxial graphene: Theory and experiment. *Physical Reports*, 542:195–295.
- [121] Todorov, I. T., W. Smith, K. Trachenko, and M. T. Dove
2006. DL_POLY_3: new dimensions in molecular dynamics simulations via massive parallelism. *Journal of Materials Chemistry*, 16(20):1911–1918.
- [122] Torquato, S.
2013. *Random heterogeneous materials: microstructure and macroscopic properties*, volume 16. Springer Science & Business Media.

- [123] Tóth, G. I., T. Pusztai, G. Tegze, G. Tóth, and L. Gránásy
2011. Amorphous nucleation precursor in highly nonequilibrium fluids. *Physical Review Letters*, 107(17):175702.
- [124] Tsukamoto, S., T. Honma, G. Bell, A. Ishii, and Y. Arakawa
2006. Atomistic Insights for InAs Quantum Dot Formation on GaAs(001) using STM within a MBE Growth Chamber. *Small*, 2:386–389.
- [125] Uhlherr, A., S. J. Leak, N. E. Adam, P. E. Nyberg, M. Doxastakis, V. G. Mavrantzas, and D. N. Theodorou
2002. Large scale atomistic polymer simulations using Monte Carlo methods for parallel vector processors. *Computer Physics Communications*, 144(1):1–22.
- [126] Venables, J., G. Spiller, and M. Hanbucken
1984. Nucleation and growth of thin films. *Reports on Progress In Physics*, 47(4):399.
- [127] Vishwanath, S., X. Liu, S. Rouvimov, P. C. Mende, A. Azcatl, S. McDonnell, R. M. Wallace, R. M. Feenstra, J. K. Furdyna, D. Jena, and H. G. Xing
2015. Comprehensive structural and optical characterization of MBE grown MoSe₂ on graphite, CaF₂ and graphene. *2D Materials*, 2:024007.
- [128] Vlassiouk, I., M. Regmi, P. Fulvio, S. Dai, P. Datskos, G. Eres, and S. Smirnov
2011. Role of hydrogen in chemical vapor deposition growth of large single-crystal graphene. *ACS Nano*, 5(7):6069–6076.
- [129] Vlassiouk, I., S. Smirnov, M. Regmi, S. P. Surwade, N. Srivastava, R. Feenstra, G. Eres, C. Parish, N. Lavrik, P. Datskos, S. Dai, and P. Fulvio
2013. Graphene Nucleation Density on Copper: Fundamental Role of Background Pressure. *Journal physical chemistry C*, 117(37):18919–18926.
- [130] Voigtländer, B., T. Weber, P. Šmilauer, and D. E. Wolf
1997. Transition from Island Growth to Step-Flow Growth for Si/Si(100) Epitaxy. *Physical Review Letters*, 78(11):21642167.
- [131] Vvedensky, D.
2004. Multiscale modelling of nanostructures. *Journal of Physics: Condensed Matter*, 16(50):R1537.
- [132] Walter, A. L., S. Nie, A. Bostwick, K. S. Kim, L. Moreschini, Y. J. Chang, D. Innocenti, K. Horn, K. F. McCarty, and E. Rotenberg
2011. Electronic structure of graphene on single-crystal copper substrates. *Physical Review B*, 84:195443.

- [133] Wang, B., X. Ma, M. Caffio, R. Schaub, and W.-X. Li
2011a. Size-Selective Carbon Nanoclusters as Precursors to the Growth of Epitaxial Graphene. *Nano Letters*, 11(2):424–430.
- [134] Wang, J.-S. and R. H. Swendsen
1990. Cluster monte carlo algorithms. *Physica A: Statistical Mechanics and its Applications*, 167(3):565–579.
- [135] Wang, Y., A. J. Page, Y. Nishimoto, H.-J. Qian, K. Morokuma, and S. Irle
2011b. Template Effect in the Competition between Haeckelite and Graphene Growth on Ni(111): Quantum Chemical Molecular Dynamics Simulations. *Journal of The American Chemical Society*, 133(46):18837–18842.
- [136] Wang, Z.-J., G. Weinberg, Q. Zhang, T. Lunkenbein, A. Klein-Hoffmann, M. Kurnatowska, M. Plodinec, Q. Li, L. Chi, R. Schloegl, and M.-G. Willinger
2015. Direct observation of graphene growth and associated copper substrate dynamics by in situ scanning electron microscopy. *ACS Nano*, 9(2):1506–1519. PMID: 25584770.
- [137] Weatherup, R. S., A. J. Shahani, Z.-J. Wang, K. Mingard, A. J. Pollard, M.-G. Willinger, R. Schloegl, P. W. Voorhees, and S. Hofmann
2016. In Situ Graphene Growth Dynamics on Polycrystalline Catalyst Foils. *Nano Letters*, 16(10):6196–6206.
- [138] Whitlam, S., E. H. Feng, M. F. Hagan, and P. L. Geissler
2009. The role of collective motion in examples of coarsening and self-assembly. *Soft Matter*, 5(6):1251–1262.
- [139] Whitlam, S. and P. L. Geissler
2007. Avoiding unphysical kinetic traps in Monte Carlo simulations of strongly attractive particles. *The Journal of Chemical Physics*, 127(15):154101.
- [140] Whitlam, S. and R. Jack
2015. The statistical mechanics of dynamic pathways to self-assembly. *Annual Review of Physical Chemistry*, 66:143.
- [141] Wilson, N. R., A. J. Marsden, M. Saghir, C. J. Bromley, R. Schaub, G. Costantini, T. W. White, C. Partridge, A. Barinov, P. Dudin, A. M. Sanchez, J. J. Mudd, M. Walker, and G. R. Bell
2013. Weak mismatch epitaxy and structural Feedback in graphene growth on copper foil. *Nano Research*, 6(2):99–112.

- [142] Winsor, C. P.
1932. The Gompertz curve as a growth curve. *Proceedings of The National Academy of Sciences*, 18(1):1–8.
- [143] Wofford, J. M., S. Nie, K. F. McCarty, N. C. Bartelt, and O. D. Dubon
2010. Graphene islands on Cu foils: the interplay between shape, orientation, and defects. *Nano Letters*, 10(12):4890–4896.
- [144] Wood, G., A. Marsden, J. Mudd, M. Walker, M. Asensio, J. Avila, K. Chen, G. Bell, and N. Wilson
2015. Van der Waals epitaxy of monolayer hexagonal boron nitride on copper foil: growth, crystallography and electronic band structure. *2D Materials*, 2:025003.
- [145] Wood, J. D., S. W. Schmucker, A. S. Lyons, E. Pop, and J. W. Lyding
2011. Effects of Polycrystalline Cu Substrate on Graphene Growth by Chemical Vapor Deposition. *Nano Letters*, 11(11):4547–4554.
- [146] Wu, B., D. Geng, Y. Guo, L. Huang, Y. Xue, J. Zheng, J. Chen, G. Yu, Y. Liu, L. Jiang, et al.
2011. Equiangular Hexagon-Shape-Controlled Synthesis of Graphene on Copper Surface. *Advanced Materials*, 23(31):3522–3525.
- [147] Wu, P., Y. Zhang, P. Cui, Z. Li, J. Yang, and Z. Zhang
2015. Carbon dimers as the dominant feeding species in epitaxial growth and morphological phase transition of graphene on different Cu substrates. *Physical Review Letters*, 114(21):216102.
- [148] Yan, Z., J. Lin, Z. Peng, Z. Sun, Y. Zhu, L. Li, C. Xiang, E. L. Samuel, C. Kittrell, and J. M. Tour
2012. Toward the Synthesis of Wafer-Scale Single-Crystal Graphene on Copper Foils. *ACS Nano*, 0(0).
- [149] Yazyev, O. V. and A. Pasquarello
2008. Carbon diffusion in CVD growth of carbon nanotubes on metal nanoparticles. *Physica Status Solidi (B)*, 245(10):2185–2188.
- [150] Yoon, D., Y.-W. Son, and H. Cheong
2011. Negative thermal expansion coefficient of graphene measured by Raman spectroscopy. *Nano Letters*, 11(8):3227–3231. PMID: 21728349.

- [151] Yu, Q., L. A. Jauregui, W. Wu, R. Colby, J. Tian, Z. Su, H. Cao, Z. Liu, D. Pandey, D. Wei, et al.
2011. Control and characterization of individual grains and grain boundaries in graphene grown by chemical vapour deposition. *Nature Materials*, 10(6):443–449.
- [152] Zhang, W., P. Wu, Z. Li, and J. Yang
2011. First-Principles Thermodynamics of Graphene Growth on Cu Surfaces. *Journal of Physical Chemistry C*, 115(36):17782–17787.
- [153] Zhou, X., D. Ward, and M. E. Foster
2015. An analytical bond-order potential for carbon. *Journal of Computational Chemistry*, 36(23):1719–1735.
- [154] Zhu, Y. and T. Wang
2015. Kinetic Monte Carlo simulation of the initial growth of Ag thin films. *Applied Surface Science*, 324:831836.
- [155] Zuo, J.-K., G.-C. Wang, and T.-M. Lu
1988. Growth kinetics of a chemisorbed overlayer in the presence of impurities. *Physical Review Letters*, 60:1053–1056.