# Multiscale Modelling of Graphene's Mechanical Properties

*Robert C. Sinclair*

A dissertation submitted in partial fulfillment
of the requirements for the degree of
**Doctor of Engineering**
of
**University College London**.

Department of Chemistry
University College London

April 28, 2020

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

# Abstract

Graphene is a household name, but is not yet a household product. The science of 2D materials was sparked with graphene's isolation in 2004, theoretical predictions made since have heralded nothing short of revolution in the fields of composite materials, electronics, and energy storage. Those revolutions are yet to materialise, but the feverish interest amongst the scientific community continues, motivated by graphene's tantalising properties and our dogged desire to exploit them.

One exciting prospect for graphene is to exploit its mechanical properties as an effective reinforcing component in composite materials. For this to be possible one must overcome the difficulty in producing high quality graphene dispersions in large quantities and effectively transfer its properties to the bulk material.

I developed an experimentally and theoretically verified forcefield for molecular dynamics, which replicates graphene's non-bonded interactions. Using this, I was able resolve graphene's unusual behaviour whilst in a low friction state (known as superlubricity) and the micromechanical exfoliation of nanoflakes of graphene. I find that graphene's low bending energy results in a pealing mechanism requiring less work than simply shearing graphite. I give insights into the nanostructure of graphene oxide, and predict that graphene oxide's percolation threshold will arrive at carbon oxygen ratios below 6, an important result for use in electronic devices.

The work presented in this thesis is part of an ongoing effort to develop a multiscale simulation method that links finite element analysis with molecular dynamics, with the aim of predicting macroscale properties of materials from nanoscale structures. This method exploits the power of high performance computing and shows that single scale simulation of graphene nanocomposites is often insufficient.

# Impact Statement

The potential benefits of graphene, if its theoretical promise can be realised, are well documented. As a flexible conductor, it could lead to a completely novel approach to electronic devices. Its uniquely high surface area to weight ratio gives it relevance for catalysis, sensing and energy storage. Having the highest measured tensile strength of any material means that graphene is a clear candidate for composite materials, which hope to exploit its mechanical properties.

Commercialisation of graphene requires reliable ways of producing large quantities of high quality graphene. In this research, I give insight into graphene-graphene interactions and the molecular mechanisms at play in graphene's exfoliation, including actionable instruction for those who wish to synthesis graphene via micromechanical exfoliation. The forcefield GraFF provides researchers with a new tool to probe graphene's dynamics on an atomic scale; I show new dynamics of graphene in a superlubric state, which may provide a way to separate graphene sheets by size. Improving graphene synthesis and having a strong theoretical understanding of the mechanisms at play bring the application of graphene closer.

Within this work I develop a general method for producing graphene oxide structures that can be used by other scientists interested in the field. This brings with it new understanding of graphene oxide's percolation threshold and reaction dynamics. Graphene oxide is significantly easier to work with, but usually comes with a large degradation of graphene's conductive and mechanical properties. By keeping the oxidised phase below its percolation threshold, one could preserve graphene's properties whilst taking advantage of graphene oxide's more facile synthesis.

Multiscale simulations of graphene-epoxy systems have great promise to give

scientists a new predictive tool to design composite materials from their atomic structure. Taking advantage of high performance computing resources we can provide a virtual lab where time from conception to testing is much shorter than in practical experiments. Our hierarchical multiscale model allows capture of mechanisms that occur on multiple length and time scales.

As part of this research I have also contributed to multiple software projects which are all open source. These include the scientific endeavours listed above, for example GraFF, a graphene oxide builder, and the hierarchical multiscale model. I have also contributed to software projects aimed at automating and exploiting high performance computing as part of the European project VECMA, collaborating with multiple research groups, creating tools to aid computational researchers in potentially any field.

Tangible research outputs include publications in Refs [1, 2, 3, 4, 5], and code repositories in Refs [6, 7, 8, 9, 10]. These have been disseminated in conference presentations and posters.

# Acknowledgements

Completing a four year research project alone would be unimaginably daunting, tortuous, and downright sad. Above all else, the research would be terrible. So here I would like to thank all the people who have helped me along the way and have made the past four years, on balance, an enjoyable experience.

Firstly, I am indebted to Peter Coveney, under whose excellent supervision I carried out the work presented in this thesis. Peter's mentorship, guidance, and immense knowledge helped at all times during my EngD, giving me thoughtful advice and the space to explore my own ideas. I am also unendingly grateful to James Suter, who has walked me through the trials and tribulations of academia, been the font of all molecular dynamics knowledge, and always held my research to the highest standard.

I would like to thank all the members of the centre for computational science, who have each passed on their individual brand of wisdom. Particularly David Wright and Robin Richardson, for our stimulating and lengthy discussions about *circles*, a testament to our shared psychosis. I would like to express my sincere gratitude to the scientists equally willing to join me in the pub as well as the lab — Sebastian, Emily, Matt, Alex, Katya, Adrian, Qingfen — I'll see you in the Housman. Special thanks go to Maxime, who never let me speak French.

From undergrad to thesis writing, my research has been exquisitely soundtracked by Lauren Laverne, Shaun Keaveny and the BBC 6 music team. Thanks for always keeping it groovy.

Finally, I want to thank my long suffering family and friends, whose love and support has kept me going, and somehow, kept me sane.

# Contents

# List of Figures

# List of Tables

# Abbreviations

| | |
|---|---|
| AMBER | Assisted Model Building with Energy Refinement |
| CG | Coarse Grain |
| C/O | Carbon to Oxygen ratio in graphene oxide |
| COMPASS | Condensed-phase Optimised Molecular Potentials for Atomistic Simulation Studies |
| cpuh | computer processing unit hours |
| CVD | Chemical Vapour Deposition |
| DFT | Density Functional Theory |
| DFA+MBD | Density Functional Approximation Many Body Dispersion |
| FEA | Finite Element Analysis |
| FEM | Finite Element Method |
| GAFF | General AMBER Force Field |
| GGA | Generalised Gradient Approximation |
| GO | Graphene Oxide |
| GNP | Graphene Nano Platelets |
| GraFF | Graphene ForceField |
| h-BN | Hexagonal Boron Nitride |
| HMM | Heterogeneous Multiscale Model |
| LAMMPS | Large-scale Atomic/Molecular Massively Parallel Simulator |

| | |
|---|---|
| MD | Molecular Dynamics |
| MLP | Multi-Layer Perceptrons |
| NPT | Isobaric-Isothermal Ensemble |
| NVE | Microcanonical Ensemble |
| NVT | Canonical Ensemble |
| OPLS | Optimised Potentials for Liquid Simulations |
| PBE | Perdew–Burke-Ernzerhof |
| PEG | Polyethylene Glycol |
| PDE | Partial Differential Equation |
| PES | Potential Energy Surface |
| PMMA | Polymethylmethacrylate |
| PDMS | Polydimethylsiloxane |
| QM | Quantum Mechanics |
| rGO | reduced Graphene Oxide |
| RF | Random Forest |
| STM | Scanning Tunneling Microscopy |
| vdW | van der Walls |
| vdW-DF2 | van der Walls Density Functional |
| vdW-optPBE | van der Walls optimised Perdew–Burke-Ernzerhof |
| VECMA | Verified Exascale Computing for Multiscale Applications |
| WHAM | Weighted Histogram Analysis Method |
| VVUQ | Verification Validation and Uncertainty Quantification |

## Chapter 1

# Introduction

Despite the far reaching attention graphene has garnered since its recent discovery, difficulties in its synthesis and processing have limited progress in virtually all its possible applications. In this chapter, I will discuss graphene's rise in importance, the current state of the graphene-polymer nanocomposite field, and the role that simulation can play in material design.

Graphene was first isolated in 2004 by the now famous and remarkably simple experiment of Geim *et al.* [11], where individual graphite layers were exfoliated using everyday Scotch tape. In response to this synthesis, which had previously been thought impossible [12], there was great speculation [13] about the applications which would take advantage of graphene's unique and superlative properties [14, 15, 16, 17, 18].

Amongst those exciting properties, that set graphene apart from other carbon nanomaterials (see figure 1.1), are: tensile strength of 1.0 TPa [14]; thermal conductivity of 5 KW $mK^{-1}$ [15]; 98% optical transmittance [16]; and high electron mobility (250,000 $cm^2V^{-1}$ s [19]). Perhaps graphene's most novel characteristic is its aspect ratio: a graphene sheet's dimensions can be in the order of metres [20] and still only one atom thick, possessing the highest surface area to weight ratio of any material: 2630 $m^2g^{-1}$ [21]. These properties give graphene great promise for micro-electronics, photovoltaics, energy storage, catalysis, biological engineering, and composite materials. In spite of the intense and heavily funded research, its theoretical promise and the commercial interest surrounding graphene, few appli-

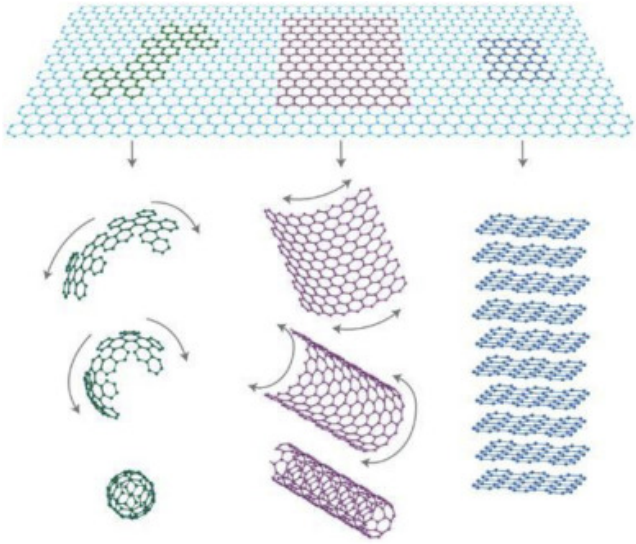

**Figure 1.1:** Schematic showing how 2D graphene is the basic building block for other carbon allotropes. It can be wrapped into 0D buckyballs, rolled into 1D nanotubes or stacked into 3D graphite. Taken from Ref.[13].

cations have found a place in industry.

Commercial successes for graphene in nanocomposites have been disappointing. Its most notable market seems to be luxury sports equipment, included in bicycle tyres, tennis rackets, and fishing rods. Evidence is rarely given to show that graphene is actively improving these products, although golf balls containing graphene nanocomposites are at least now used on the professional golf tour [22]. Time will tell if this is more than a passing fad. This lack of commercial interest is perhaps to be expected, since it is typical for a new material to take 20 years to reach the market after its initial conception [23, 24].

In other sectors, graphene is currently used in an ink for flexible circuitry [25] and as part of the cooling system in Huawei's new smartphone [26]. Exploiting graphene's electronic properties is not the primary goal of this research. However, its use in conductive inks may yet prove to be its most successful application.

Composite materials aim to combine the properties of two or more constituent components to create a material with new and enhanced physical or chemical properties. They are typically highly tuneable and diverse, common examples include concrete or carbon fibre plastics. Using graphene as a component in such a material could bring its uniquely high Young's modulus to practical scales. Graphene's ther-

mal and electrical conductivity are also attractive for such materials; Stankovich *et al.* [27] showed that the percolation threshold can be reached with as little as 0.1 vol% loading for room temperature electrical conductivity. Graphene's high aspect ratio has also shown promise for increasing a material's fatigue lifetime [28] and suppressing gas permeability [29].

Attempts to realise graphene's potential in composite materials have so far largely failed. Graphite is made of stacked graphene sheets held together by non-covalent forces which may intuitively seem easy to overcome. However, preparing a uniform distribution of exfoliated graphene throughout a polymer matrix has proved anything but trivial and remains an intense area of research [30, 31, 32]. Such is the difficulty in first isolating and then processing graphene in large quantities, computational experiments are being employed to predict the best materials to use in anticipation of an adequate production method. The challenge is comparable to clay nano-composites, a field that has existed for many years [33, 34], where large, mesoscale simulations have proven useful in design and understanding of these systems and in the future may also enlighten the field of graphene nanocomposites.

In this chapter, I will first discuss the synthesis of graphene and some of its derivatives. I will then outline the role of graphene in composite materials and the use of simulation to study those materials.

## 1.1 Graphene

Graphene can be thought of as a monolayer of graphite, but as with many nanomaterials, its properties vary greatly from the bulk. It consists of a honeycomb network of $sp^2$ hybridised carbon atoms, most of its exceptional mechanical and electronic properties derive from this extended delocalised electronic system.

### 1.1.1 Synthesis

#### 1.1.1.1 Mechanical Exfoliation

The first ever reported isolation of graphene was achieved by mechanical exfoliation of graphite by repetitive cleavage using Scotch tape [11]. The layers in graphite are held together by much weaker forces than the covalent bonds between atoms in the

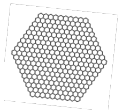

same layer, as such the stripping processes produces surprisingly good quality and large graphene sheets. This method is useful for analysis but there is little scope for scaling up to larger quantities. It is also difficult to remove the graphene sheet from the substrate [35], so it is of little use to composite material production.

#### 1.1.1.2 Chemical Vapour Deposition

A full description of the chemistry involved in graphene chemical vapour deposition (CVD) is beyond the scope of this review; for more information readers are directed to Ref. [36]. Simply put, a hydrocarbon precursor gas is fed into a reactor, the gas decomposes on a substrate surface and forms a coating. The coating's composition is sensitive to reaction temperature, pressure, precursor composition and surface chemistry. Graphene was first made using this method on a nickel substrate [37]; now, a copper substrate is more commonly used [38, 39]. Problems involved with CVD include control over the number of layers, film quality and exfoliation from the substrate. Recently, roll-to-roll methods have improved on the latter of these problems and show promise for large graphene coating applications [38] and in the future provide a method for scaling up to the quantities needed for composite materials.

#### 1.1.1.3 Liquid Phase Exfoliation

Unlike the previous synthesis routes, liquid phase exfoliation has clear potential for the large scale production of graphene necessary for composite materials. The process ideally involves exfoliating graphene directly from graphite in a solvent, usually stimulated by sonication or by inducing shear rates above $10^4$ $s^{-1}$ [40]. The success of this synthesis is somewhat disputed but control over the number of layers and dispersion stability is not good [41, 42], the synthesis usually producing graphene with multiple layers. The solvent is chosen to match the surface tension of graphene so that a dispersion is thermodynamically allowed, dimethylformamide and N-methyl-2-pyrrolidone are widely used for this purpose. Exceedingly long sonication times are needed and only dilute suspensions are stable: in one study graphite was sonicated in N-methyl-2-pyrrolidone for 460 hours; in the resulting suspension 20% of all graphite platelets, or 4% by weight, were monolayers (i.e.

graphene), and the suspension was stable at 1.2 mg $L^{-1}$ [43]. Long sonication times are needed to exfoliate the graphene sheets but also degrade the sheet size (although it has been suggested that they would not decrease below 1 $\mu$m [43]). Using shear forces to exfoliate graphene can result in better quality graphene flakes, averaging less than 10 layers and 300-800 nm in size [40], so this seems likely to replace sonication as the exfoliation method since shearing large volumes of solution is easier than sonication. However, the toxic solvents involved and poor control of size and layers, still represent significant drawbacks to these methods.

Dispersions of graphene can be assisted by inorganic salts, surfactants [44] and functionalisation [45], which often affect the properties of graphene and has been used to varying effect. For more information see Refs. [41, 42].

#### 1.1.1.4 Thermal Exfoliation

Thermal exfoliation requires graphite that has been intercalated by some compound as a starting material. The intercalated graphite is then rapidly heated, as the molecules between layers heat up they can be thought of as a gas, building up pressure between the graphite layers, forcing the graphene sheets to fully exfoliate. Temperatures of around 1000 °C are used and, as one might expect, the chemical composition of the graphene produced can become compromised. The process is generally reductive so this effect can be surprisingly small and the yield and size of the graphene sheets can also be good [42].

A successful experiment is described in Ref. [46], where thermal exfoliation is combined with liquid-phase exfoliation. Graphite was first expanded by rapid heating, then intercalated by oleum and tetrabutylammonium in DMF, resulting in a 90% yield of graphene sheets.

A detailed review of other graphene exfoliation methods and combinations of the methods described above is given in Refs [41, 42, 18].

### 1.1.2 Electronic Properties

This research focuses on synthesis of graphene and exploiting graphene's potential as a mechanical reinforcement, but it would be remiss not to mention graphene's

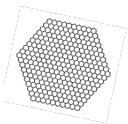

electronic properties. Graphene's high electron mobility has been known since its discovery [11, 17]. Developments in fabrication have shown that graphene nanoribbons can have lower resistances than theoretically predicted, showing great potential as components in nanoelectronics [47].

Recently, the discovery of superconductivity in bilayer graphene has sparked intense interest in graphene as a unique electronic material. By rotating two layers of graphene by an angle of 1.1°, Jarillo-Herrero *et al.* have shown that graphene is superconducting at temperatures below 1.5 kelvin [48, 49]. As of January 2019 only two labs have synthesised superconducting bilayer graphene [50], with Young *et al.* showing that its electronic properties can also be controlled by applying pressure to the system, influencing the layer spacing between the two layers of graphene [51]. Theory had predicted that unusual electronic properties would occur in twisted bilayer graphene [52]; however no consensus exists on the explanation for bilayer graphene's superconductivity. The unique ability to switch between superconducting and insulating, by means other than controlling the temperature, is particularly exciting.

## 1.2 Graphene Oxide

Until this point, all data and synthesis routes have referred to graphene and have intentionally avoided the mention of graphene oxide (GO). GO is by no means a new compound. It was first synthesised in 1859 [53] and has very different properties to the pristine graphene referred to above, but because of its similar dimensions and relative ease of synthesis and processing, it is often used in place of graphene. In a worryingly wide spectrum of the current scientific literature, GO is misleadingly referred to as simply graphene, so great care should be taken when reading around this subject to inspect the synthesis routes and characterisation data described.

GO retains the carbon skeleton of graphene with oxygen-containing functional groups on the surface and edges added during its preparation. The oxygen groups are thought to have an essentially random arrangement [55] with epoxy and alcohol groups on the surfaces and carboxyl groups at the edges. Less frequently, 5 or

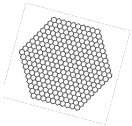

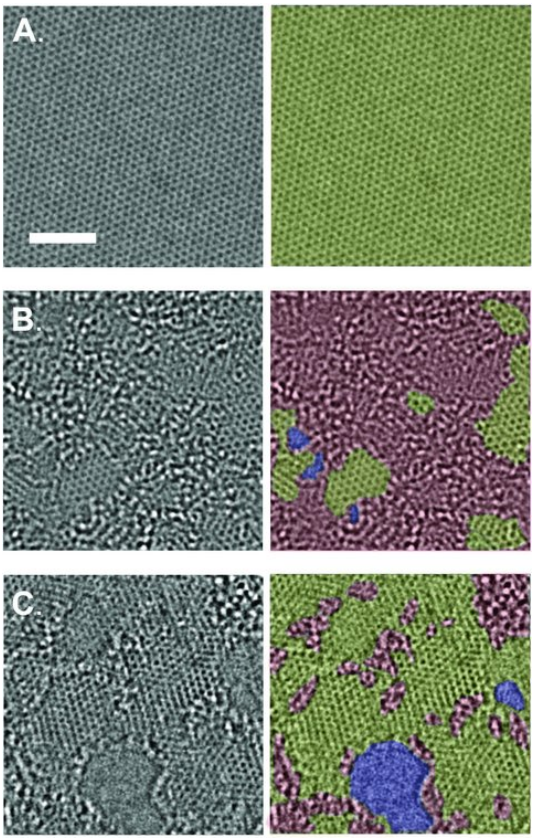


**Figure 1.2:** Ultra-high resolution TEM images showing the atomic structure of graphene, graphene oxide and reduced graphene oxide (from top to bottom). Images on the right are coloured by the author (green: $sp^2$ graphene; red: disordered regions indicating oxidisation; blue: holes). Scale bar: 2 nm. The graphene oxide was made via a modified Hummer's method and reduced by hydrazine. Taken from Ref. [54].

6-membered lactone rings will appear throughout the structure [56]. Their exact nature and arrangement is not known and whether oxidised "islands" occur is still a matter of debate [54, 57, 31, 58, 59]. The subject will be explored in chapter 6 of this thesis. The structure and density of the functional groups will depend on the synthesis routes taken but a C/O ratio of less than 4 is typical for GO [60]. GO is commonly synthesised by the Hummers' method [61]: in which graphite is reacted with sulfuric acid, sodium nitrate and potassium permanganate; this method typically achieves a C/O ratio of 2. After the oxidation takes place the layer spacing in graphite is increased from 0.34 nm to 0.6 - 1.0 nm [18]: the layers are extremely weakly held together and a gentle stir in water provides enough shear force to fully exfoliate the graphite oxide into graphene oxide. Compare this synthesis to the ones

described above for graphene, and one can see the attraction of using GO.

The transformation to graphene oxide should not be taken lightly as it results in significant degradation of its conductivity, stability and mechanical properties [18]. The loss of graphene's desirable properties is largely due to the disruption of the $sp^2$ carbon network (see figure 1.2), while the thickness of a sheet is effectively doubled during oxidation and holes or other defects can occur. While a Young's modulus of 250 GPa has been measured for GO [62], measurements vary and depend heavily on the preparation method; however, a stiffness of 40 GPa and strength of 120 MPa are typical [63, 18]. The electrical and thermal conductivity of GO is several orders of magnitude lower than graphene [27]. Samples will often begin to degrade above 100 °C, limiting polymer coprocessing and reducing GO's use as a composite material component. Its use in nanocomposites is promising though, as the functionalised surface can improve the polymer-nanocomposite interaction, improving stress transfer between matrix and sheet leading to a stronger material [30].

### 1.2.1 Reduced Graphene Oxide

Reducing graphene oxide can partially restore some of the properties of graphene. Typically, a strong reducing agent such as hydrazine can yield a C/O ratio of 10 [64]. A simple alkaline wash can reduce the oxygen content by removing oxidative debris left from Hummer's-like methods [65], using this method a C/O ratio of 6 is easily achieved. Reducing graphene oxide is mainly used to restore some of the delocalised $\pi$-network and improve the sample's conductivity, but the reductive process damages the carbon network and creates holes in the sheet which offset any improvement in the material's stiffness or strength [54] (see figure 1.2).

Novel routes toward synthesising graphene through GO have produced C/O ratios in excess of 250 [56]. This is no doubt remarkable and shows the promise in this field for a route towards mass production of graphene, but for now, the distinction between products remains necessary as significant defects and $sp^3$ carbons still occur in the final structure.

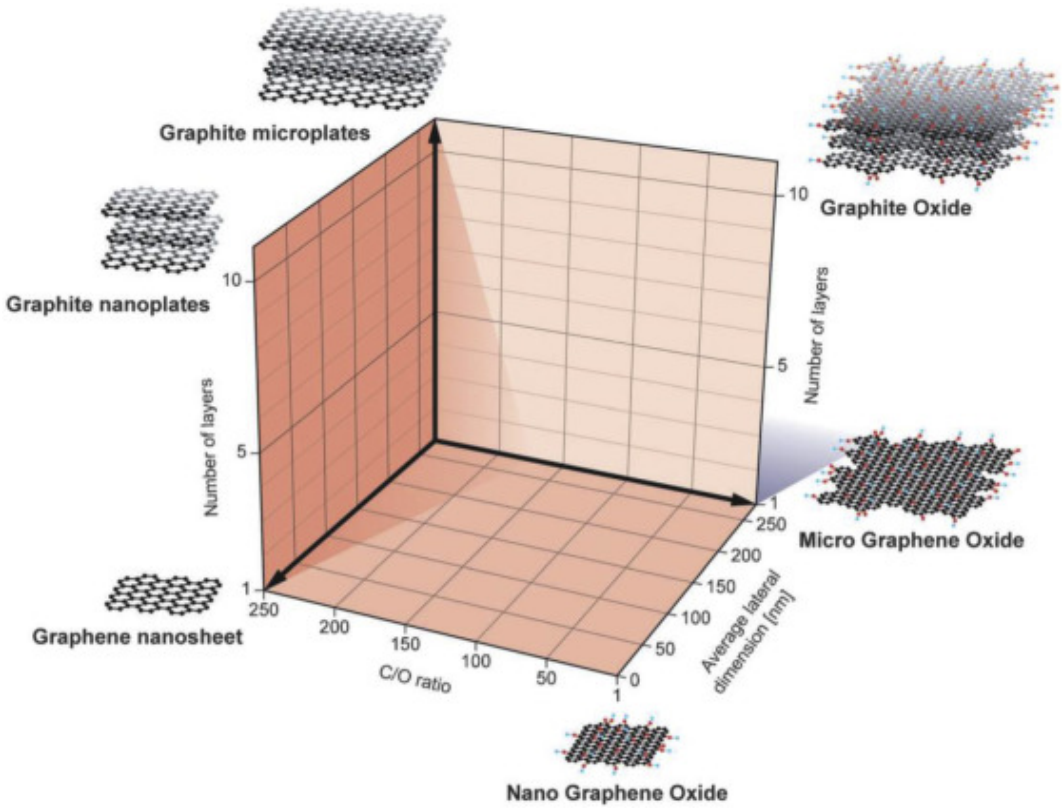


**Figure 1.3:** A guide to naming graphene derived compounds as proposed in Ref. [60].

## 1.3 Nomenclature

As has been alluded to already, the nomenclature in this subject is often confusing and at times, misleading. Articles frequently use their own terminology for established systems and make use of ambiguous definitions to exaggerate their results. The problem is significant enough that reviews have been published specifically calling for consistency and honesty from the scientific community, proposing a universal nomenclature to be followed; see figure 1.3 [60, 66]. This nomenclature system is starting to be followed as journals begin to demand consistency.

The headline from the aforementioned reviews is that the term 'graphene' should be kept for the single layer, pristine network of $sp^2$ hybridised carbons with dimensions greater than a micrometre. These are strict recommendations, but to avoid confusion any deviation from these criteria should be explicitly defined and characterised by the author. It is also worth pointing out again that GO has been known for over 150 years and much of the recent research on it has been published

on the coattails of graphene's rise to fame and rush to explore 2D material properties in general. Terms that have since become standard: 'few-layered graphene' refers to 2-5 layers graphene, and 'multi-layer graphene' refers to 5-10 layers of graphene [30].

It is understandable that in a highly competitive field, such as graphene, authors should want their research to stand out (over 35,000 papers referencing graphene were published in 2018 according to `webofknowledge.com`), but this means the need for rigorous exclusion of bias from scientific work becomes acutely relevant in this field [67].

## 1.4 Graphene-Polymer Nanocomposites

Graphene nanocomposites have made it to market in a small number of specialised applications but are yet to be taken up by large industries. This is simply because graphene reinforced polymers are yet to demonstrate improvement on existing materials at a reasonable cost. That is not to say there is no place for graphene if its potential is eventually realised. For a comprehensive review of the state of graphene composite materials, readers should look no further than the up-to-date review by Young and his colleagues [30].

Graphene nanocomposites could directly replace carbon fibre composites, an industry which demands 72,000 tonnes worldwide [68], providing lighter weight components and an feasier fabrication process. The aviation industry is a major target for graphene nanocomposite producers. For example, the 787 Dreamliner is 80% polymer composite by volume and each 787 aircraft contains approximately 35 metric tons of carbon fibre reinforced plastic. One study found that replacing the carbon fibre composite with carbon nanocomposites would reduce the airframe's weight by 17.32%, resulting in average fuel savings of over 10% [69].

The key obstacles in creating useful graphene-polymer nanocomposite materials include synthesising large quantities of well defined graphene, achieving an even distribution of exfoliated graphene, and maintaining the dispersion through any processing required for the material (figure 1.4 shows the states of graphene

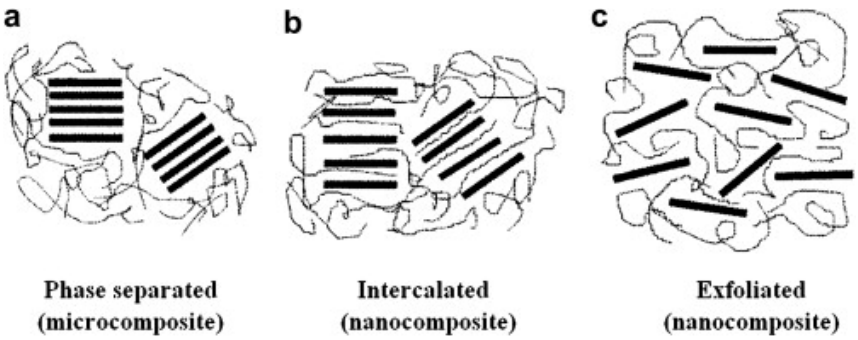


**Figure 1.4:** Schematic showing three morphological states for graphene-polymer nanocomposites. For most applications, a fully exfoliated composite is desired.

exfoliation in a polymer). The first of these challenges is addressed in the previous sections; the other two are largely dependent on the graphene-polymer interaction energy and the processing techniques used. The range of systems being studied is broad, only the key examples relevant to our research will be discussed below.

### 1.4.1 Model Nanocomposites

To study the ideal behaviour of graphene embedded in a polymer, researchers can deposit graphene onto a polymer substrate (either by mechanical exfoliation or CVD), then spin coat the sample with polymer to create a 'sandwich-like' configuration that can be analysed in detail.

An ideal composite material, with complete stress transfer between matrix and filler, will obey the rule of mixtures when measuring, for example, the composite's Young's modulus:

$$E_c = V_f E_f + (1 - V_f) E_m; \tag{1.1}$$

where $E$ is the elastic modulus; $V_f$ is the volume fraction of the filler; and subscripts $f$, $m$ and $c$ denote the filler, matrix and composite. The performance of a filler in a composite is measured by calculating its effective contribution. Once the elastic modulus of a composite is known, one can replace $E_f$ with $E_{\text{eff}}$ in equation 1.1 and solve for $E_{\text{eff}}$. $E_{\text{eff}}$ will usually be lower than the ideal stiffness of graphene (1.0 TPa [14]). Several factors such as non-ideal flake orientations, incomplete dispersion,

and ineffective stress-transfer can contribute to this. Vlassiouk *et al.* produced a model graphene-PMMA (polymethylmethacrylate) composite with effective modulus, $E_{\text{eff}}$, of 1.0 TPa [70]. This result is remarkable and shows the potential of graphene as a composite filler; achieving this behaviour in less labour-intensive synthesis routes will be the goal over the coming years.

From model nanocomposite studies other interesting characteristics can be gleaned. Stress-transfer to graphene is less effective after strains of approximately 0.5% strain, when slippage/creep mechanisms start to occur [30]. Graphene flakes need lateral dimensions over 30$\mu$m for effective reinforcement [71, 72]. Compared to monolayer graphene, bilayer graphene has similar reinforcement capabilities, but more layers degrade its reinforcement capabilities, for example 10-layer graphene has 40% the reinforcement of monolayer graphene [73].

### 1.4.2 Preparation of Bulk Nanocomposites

Ideally, graphene could be exfoliated directly into a polymer melt and then injection molded as needed, like any other plastic. This system would also be good to study by simulation, as it avoids solvents and limits the number of species that need to be considered. Despite the obvious practical benefits to this method, no system that could mediate graphite exfoliation down to graphene is known [32]. This is unsurprising given the numerous difficulties encountered when exfoliating graphite to graphene described above. Instead, a separate exfoliation step is usually followed by some method of dispersion into the polymer.

Perhaps the most common method used for dispersing is solution mixing [32]. Typically GO or some other functionalised graphene will be used to generate the suspension. The polymer can be added by simply dissolving it in the suspension or mixing with an already solvated polymer. Once the constituents have been mixed a non-solvent for the polymer will precipitate the plastic with a good distribution of platelets. As you might expect, the drying process will lead to some aggregation and completely excluding solvent molecules is difficult, but it is experimentally a simple method. Solution mixing has proven popular for GO and water soluble polymer systems for the ease of handling and preparation.

Similarly, melt mixing involves dissolving an exfoliated graphitic powder into a melt-polymer [74]. Avoiding a solvent makes this method attractive for scaling-up and industrial applications but the dispersion of platelets is significantly worse than those created by other methods. Again, this method has only been used effectively with GO or some other chemically modified graphene.

*In situ* polymerisation involves dissolving the pre-exfoliated graphene platelets in neat monomer, followed by polymerisation in the presence of the dispersed platelets. The dispersion of graphene is usually very good using this method as intercalated monomers can mediate complete exfoliation during polymerisation. *In situ* polymerisation is frequently extended to grafted systems: functional groups on GO can be used as starting points for grafting-from or grafting-to systems.

The methods described above are clearly a simplified overview of a vast area of research. Many other methods are being worked on in the hope of creating an easy method of dispersion. Here, systems involving chemically modified graphene have been neglected, though it is a diverse field because they are less relevant to computer simulation (in general we would like to work with the least number of species possible). Fortunately, melt mixing methods are steadily improving: the graphene-polymer interaction is now better understood and methods for creating extremely high shear forces to aid dispersion are being developed. Computer simulation has the most potential for predictive power and increasing understanding in melt mixing systems, firstly because it contains the fewest number of species and secondly, the system has many parallels with clay-polymer systems which have been studied computationally before [34].

To achieve a homogeneous material, 2D-nanoparticles must be randomly orientated otherwise the composite would have different mechanical properties along different axes [75]; although one study suggested the disparity would never be more than a factor of 2 [76]. This can be a problem when processing the composite, for example, injection moulding can create alignment of sheets along the direction of flow. When the filler is aligned in this way gas permeability is reduced in the direction normal to the sheets [31, 29]. Far from being a disadvantage, the improved

gas-barrier properties due to anisotropic dispersions could be one of the easiest and readily commercialised applications of graphene composites.

### 1.4.3 Graphene-Polymer Interaction

In the current literature, there are two approaches to explain the mechanism of graphene reinforcement of a polymer matrix. First, Maron *et al.* argue that nanocomposites should be considered as molecular composites or self-reinforced composites [77]. This approach takes into account the nanoparticle's role as a nucleation site for crystallisation and polymer confinement. However, it cannot explain the often disappointing reinforcement ability of graphene compared to the predictions made by the 'rule of mixtures' (see equation 1.1).

An alternative explanation is proposed by Young *et al.*[78], who acknowledge that the reported reinforcement contribution of graphene is often much lower than graphene's ideal 1.0 TPa. Effective transfer of this exceptional stiffness to the polymer network is impossible. Instead, the composite's stiffness, $E_c$, is a function of the matrix's stiffness, $E_m$, the volume fraction of the filler, $V_f$, and the aspect ratio of the sheet, $s$:

$$E_c \approx E_m \left[ 1 - V_f + \frac{s^2}{12} \frac{\eta_0}{(1+\nu)} V_f^2 \right]. \tag{1.2}$$

$\eta_0$ is the Krenchel orientation factor, if flakes are aligned along the strained axis $\eta_0 = 1$; for random flake orientations $\eta_0 = 8/15$ [79, 76]. This description is derived from continuum shear-lag theory and neatly explains that the reinforcement is due only to the nanoparticle's aspect ratio, in fact, there is no contribution from the nanoparticle's stiffness in equation 1.2.

It is noteworthy that, by this explanation, few-layer graphene (e.g. 5 sheets thick) is a better reinforcement particle than graphene because it will not crumple in the polymer and retain a higher aspect ratio. Young *et al.*'s theory explains the elastic modulus at low strains but does not account for graphene's effect on the composite's toughness [78]. It does suggest that simulations of graphene composites need to capture enormously disparate length scales, a challenge that will be

discussed in the next section.

Whether improving a composite's mechanical properties relies primarily on a strong matrix-filler interaction [80], or if the high aspect ratio of the nanoparticle is enough, is still under investigation. It is known that stress transfer occurs in graphene-polymer composites, and this can be measured by Raman spectroscopy [71]. Studies in this vein suggest a strong interaction between matrix and nanoparticle is required for the resulting composite to embody graphene's mechanical properties. One way of doing this is by using functionalised graphene; both experimental and theoretical studies have shown an increase in effective modulus with higher oxygen content in GO [81, 82]. If this interaction is deemed crucial, the type of functionalisation needed on a sheet will, of course, be system dependent.

The elastic modulus is not the only desirable property nanoparticles give to a material. In the case of thermosetting polymers, graphene acts as a barrier, interrupting the polymer network of cross-linked molecules, essentially weakening the matrix. However, the inclusion of graphene acts to deflect crack propagation and confine polymers; this is hoped to increase toughness and make the material much less brittle, a common drawback of using thermoset polymers like epoxy resins.

### 1.4.4 Bulk Nanocomposites

It is clear from the above analysis that using the simplistic rule of mixtures to estimate graphene's effectiveness as a polymer reinforcer is inadequate. A remarkable review of the literature was conducted by Papageorgiou *et al.*, who have collated the effective modulus of graphene in different polymer systems [30]. Graphs summarising their research are shown in figure 1.5.

The range of systems summarised in figure 1.5 is enormous. The systems vary in their matrices (from elastomers to thermosets), graphene preparation routes, percentage loading, processing techniques, flake size, and flake orientation. The spread is large on the log-log plot but the correlation is remarkable close to unity. This is in line with the group's claim that $E_{\text{eff}} \propto E_m$ — introduced as part of their work modelling graphene-nanocomposites with shear-lag theory, mentioned in the previous section. The data gathered by Papageorgiou *et al.* shows that polymers

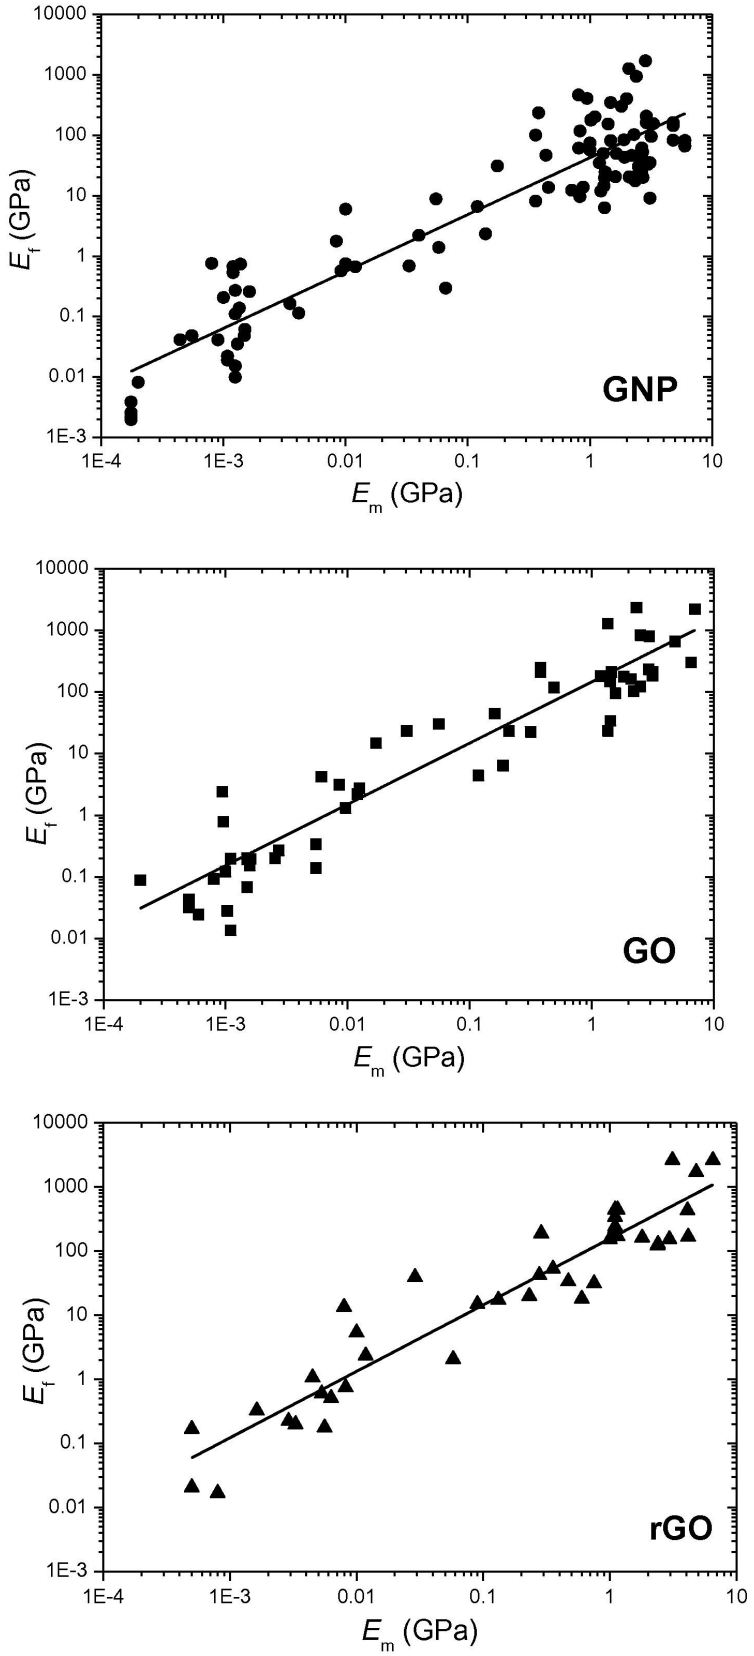


**Figure 1.5:** The effective modulus of graphene (GNP), graphene oxide (GO) and reduced graphene oxide (rGO) as calculated using equation 1.1 as a function of the modulus of the matrix. Taken from [30].

with lower moduli, such as elastomers, exhibit the largest proportional improvement in strength and stiffness on addition of graphitic nanoparticles.

## 1.5 Computational Studies

Nanocomposite materials owe many of their unique properties to the constituent's large aspect ratios, but the wide gap in length scales across the system presents a problem when it comes to simulation. Computational techniques have been well established for systems at certain length scales: *ab initio* methods, like density functional theory (DFT), calculate precise data about systems with a handful of particles; molecular dynamics (MD) uses classical mechanics to sample a wide range of systems and processes on the nanoscale; and continuum models, such as the finite element method (FEM), neglect atomistic detail but give information on engineering scales (see figure 1.6). To study graphene-polymer systems accurately a simulation must have atomistic detail to capture graphene correctly, but also sample large time and length scales to observe the slow diffusion of polymers and exfoliation/aggregation of graphene sheets.

The work in this thesis centres on particle based methods of simulation, mainly using classical all-atom MD, to describe a system. These simulations allow chemical specificity and atomic precision with an intuitive range of application. Particle based simulation methods, however, face a serious challenge: how does one simulate graphene whilst respecting its enormous aspect ratio? To simulate graphene flakes capable of giving the best reinforcement to polymer matrices, accurate descriptions of angstroms (the size of atoms) and above tens of microns (ideal flake sizes) are needed. That would mean simulating lengths spanning five orders of magnitude, this is simply out of reach for current methods; currently, even the largest MD simulations ever attempted only reach 100s of nanometres [83]. Therefore insights into the reinforcement mechanisms, exfoliation, dispersion and processing, which require long timescales to simulate, will require new approaches [84]. Coarse grained (CG) simulations are a typical way of accessing longer time and length scales. This coarsening of particles necessarily loses detail at the lower end,

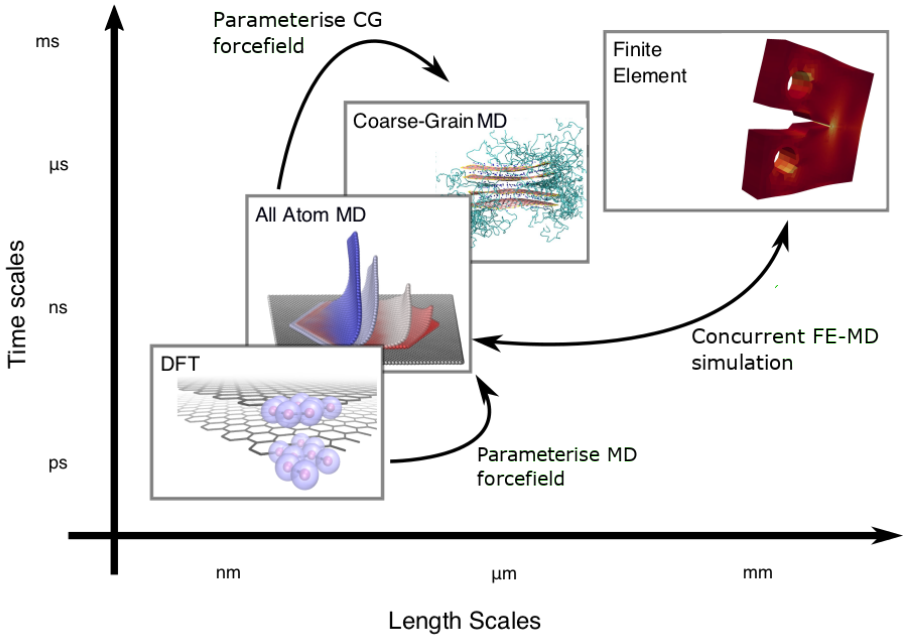


**Figure 1.6:** A guide to the different computational modelling techniques and associated time and length scales referred to in this thesis. The arrows indicate multiscale methods that have been developed as part of this work. Adapted from Ref.[87]

however, novel shaped particles or mixed precision within the simulation offer some ways to preserve the aspect ratio. These techniques often run into problems of load balancing as the computational demand varies between regions within the simulation domains, but progress has been made recently in tackling this problem [85].

Another complication arises when attempting to model the polymer matrix. Polymer chains are extremely large, often ranging from 100s to 100,000s of monomer units. The entanglement length of thermoplastics usually requires over 100 monomers [86], a length scale that is within the range of typical molecular dynamics simulations. However it is often desirable to simulate larger polymers (to reach stronger polymer materials), and avoiding self-interaction across periodic boundaries is an obvious limitation. Reptation (a diffusion process in polymer melts) of single polymer chains - each hundreds of nanometres long - in reality, require several microseconds to relax. It is important to be aware of these limitations when reading literature around polymer MD simulations.

As well as the disparity in length scales, polymer systems are anisotropic fluids which are well known to be difficult to model as the computational power needed to simulate them is so high. These systems bridge the gap between a regular crystalline solid with perfect long-range order, which can be modelled using a small, efficient periodic simulation, and an isotropic liquid with no long-range order, which can be modelled using a continuum model. Anisotropic liquids, therefore, require large simulation boxes to capture both their fluidity and short-range order.

Computer simulation has been used in many ways to probe different aspects of graphene. In this way, we can study graphene in a manner that is useful to researchers but impossible to measure experimentally.

A good review of methods for graphene simulation has been written by Otyepka *et al.* [88]. Their review details the common and recommended techniques for wave-functional, density-functional and molecular dynamics simulation of graphene. Graphene's interactions with other molecules are characterised by its strong, long-range dispersion forces: quantum mechanical (QM) models must then implement large basis sets and MD simulations should go beyond the typical pairwise intermolecular interaction (which has been suggested to be insufficient [89]). Including these measures obviously come with a large computational penalty; fast, accurate models at both levels of theory are needed to further this field.

### 1.5.1 Isolated Graphene

Computational studies were able to elucidate the theorised ripples in graphene [90]. The ripples are estimated to be 80 Å in length and 10 Å high [90, 91]. Two dimensional crystals were thought to be impossible because long wavelength fluctuations would destroy their long-range order [12], but the intrinsic ripples in graphene, which are an anharmonic coupling of bending and stretching modes, can suppress the long wavelength modes.

The delocalised $\pi$ electrons above and below the graphene sheet create a quadrupole normal to the basal plane. In an infinite graphene sheet this quadrupole cancels out for particles above or below the sheet, however, Kocman *et al.* showed that a graphene platelet must be micrometre size for the quadrupole effect to be

insignificant [91]. The quadrupole field is also non-negligible near the edges of a graphene sheet; a common practice in simulation is to terminate a sheet with hydrogen atoms, which would create a significant variation in the charge distribution which is often overlooked.

Another consideration is the electrostatic field created at the peaks and troughs of the ripples in graphene. This could significantly alter the way that graphene sheets interact as a typical MD simulation that uses point-charges, centred on atoms, would not capture any of the effects mentioned here.

### 1.5.2 Graphene-Graphene Interaction

We are particularly interested in the process of graphite exfoliation and aggregation. To understand how to create and maintain a graphene dispersion we must understand how to model the interactions between different graphene sheets.

It has become common practice to simulate graphene systems with MD forcefield potentials derived from organic molecules or other unrelated systems (see section 2.1) [88]. So far, to the best of our knowledge, no comprehensive testing or justification for transferring these parameters to graphene has been given in the literature. It is also common to simulate graphene with reactive forcefields like ReaxFF [92] or AIREBO [93], but these are significantly more costly to calculate.

Validating graphene simulations is difficult because so little thermodynamic data exists that could be used for fitting parameters. Perhaps the best experimental data for this purpose is the interlayer binding energy of graphite, although measurements vary hugely. Even many QM approaches fail to fall within the extraordinarily large experimental range [94] (see section 3.2, figure 3.1 and table 3.2 for more detail).

One can also compare with the out-of-plane layer ‘breathing’ modes of vibration in few layer graphene. This frequency can be measured experimentally using Raman spectroscopy [95], or estimated from the shape of a graphite layer binding energy curve [96].

It might be unsurprising then that no good parametrisation for MD forcefields has been carried. This summary certainly devalues any previous findings based on

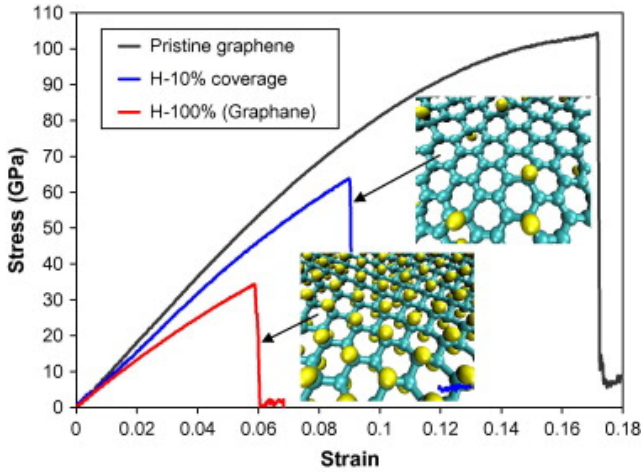


**Figure 1.7:** Schematic showing the effect of $sp^3$ defects on a graphene sheet's mechanical properties. Defects are hydrogen atoms introduced in random positions. Taken from Ref.[108].

parameters developed for biological molecules, because we simply do not know whether these forcefields are applicable to graphene.

Many other articles have shown that the pairwise potentials used in MD to model van der Walls (vdW) forces are insufficient in graphene and other carbon nanostructures [97, 89]. This has not stopped researchers using the traditional potentials to study graphene-graphene interactions [98, 99, 100, 101, 102], graphene crumpling [103] and graphene-$C_{60}$ interactions [104] *inter alia*.

QM attempts to model graphene have focussed on small aromatic molecules like coronene and extrapolated those results to graphene [105, 106, 88]. Periodic calculations get round this limitation and have been used to study adsorption onto graphene but cannot account for ripples in the structure. For a detailed overview of this area see Ref.[88].

### 1.5.3 Graphene Defect Studies

MD has been used to study the mechanical properties and fracturing processes of graphene. To do this a reactive forcefield must be used: unlike typical MD simulations, which use a harmonic potential to represent bonds, a Morse potential [103] or an adaptive bond order potential [107, 92] must be used to allow for bond breakage.

Using such a forcefield, Shenoy *et al.* studied the effect of $sp^3$ defects on

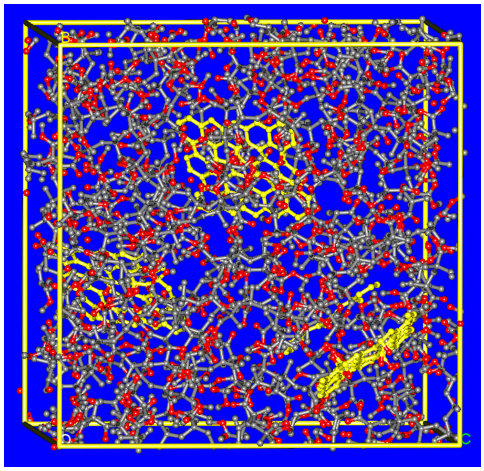

**Figure 1.8:** Molecular dynamics simulation snapshot taken from Ref.[82]. 6.5% weight loading of graphene oxide nanoplatelets (yellow) in PMMA (carbon: grey; oxygen: red).

graphene's mechanical properties [108] (see figure 1.7). Their studies involved putting graphene sheets under tension and randomly introducing hydrogen atoms to the sheet to represent $sp^3$ defects; this has direct relevance to graphene oxide defects. At just 5% hydrogen coverage a drastic decrease in failure strain and tensile strength were observed, reaching a 65% reduction by 30% hydrogen coverage, at which point the Young's modulus had decreased by 30%. Shenoy's results suggest that even a small number of $sp^3$ defects could dramatically decrease graphene's mechanical performance.

CVD often produces polycrystalline-graphene, a graphitic sheet made from grains of nanometre sized graphene crystallites with irregular structure at grain boundaries [109]. Molecular dynamics simulations of these systems predict a Young's modulus for this morphology of less than half that of pristine graphene, but also explain why polycrystalline-graphene has a remarkable ability to maintain its strength in the presence of large defects [110]. Grain boundaries reduce stress concentration by encouraging branching around the crack-tip so toughness of the sheet is preserved. This information's relevance to nanocomposites needs to be investigated but it suggests that it is unnecessary to strive for pristine graphene.

### 1.5.4 Graphene-Polymer Systems

Given the uncertainties in MD parametrisation for graphene outlined above it is surprising that many polymer-graphene systems have already been studied [111, 112, 113, 114, 115, 116]. MD (and CG) simulations use periodic simulation boxes, after the simulation has equilibrated the elastic modulus of the system can be measured by incrementally 'stretching' or 'compressing' the box and measuring the difference in energy [117, 82, 118]. However, the systems are equilibrated far above the glass transition temperature, to allow for adequate phase sampling, and MD simulations cannot realistically model the crystallisation of a polymer which occurs on the order of seconds. Comparisons with engineering materials can still be made in the case of large polymer chains, as the amount of crystallisation that takes place is small.

Adding different functional groups to a graphene sheet has improved a composite's mechanical properties, showing the advantages of using GO as a replacement for graphene. Simulations have shown that the Young's and shear moduli nearly doubled for composites with higher degrees of oxidation on graphene [82] (see figure 1.8). As mentioned above, a strong polymer-graphene interaction is thought to improve a composite's mechanical properties, it is obvious then that GO shows increased elastic modulus in PMMA, a polar polymer, over graphene. Perhaps surprisingly, the interaction energy increases by a similar amount when the functionalisation of graphene increases in a polyethylene system [45].

The two systems mentioned, Refs [82, 45], both purely MD simulations, are limited however by their size. They use very short polymer chains, 10 and 15 monomer units, we cannot assume these would not crystallise if allowed to cool under equilibrium conditions so these results may not relate to their respective solids. The graphene sheets used contain 66 and 1,600 carbon atoms (square sheets approximately 14 Å and 50 Å in length), these are far smaller than what experimentalists think is necessary for optimum reinforcement ($> 30\mu m$). Clearly, the edges of these simulated sheets take up comparatively more volume and therefore have more effect on the material than typical graphene sheets. The edges are terminated with C-H

bonds which is an approximation, in reality, they can be unsaturated carbon atoms [119], folded edges [120] or a mix of functionalisation depending on the sheet's preparation. The importance of these results to experimentally made composite materials needs to be evaluated with care.

CG models can be very specific to their purpose, a result of the number of approximations that are made in their design. A review on how such models can be built for polymer systems using *ab initio* through to empirical fitting is given in Ref. [121]. We are not aware of any CG system for graphene that do not group a number of graphene atoms into spherical particles, e.g. 4:1 [122] and 7:1 [123] regimes have been published using an inverse Boltzmann approach to test mechanical properties and polymer interaction respectively. To respect graphene's single-atom thickness and unusual dispersion forces a novel CG approach is needed that does not use spherical particles. Due to the lack of experimental data to compare to, these simulations would be very difficult.

A common technique to get around the limited flake size one can simulate with MD is to simulate an infinite graphene sheet extending through periodic boundaries. The mechanical properties derived from these simulations assume perfect stress transfer from matrix to graphene as any tension in the lateral directions will stretch the graphene lattice proportionally, without allowing for slippage or any rearrangement. Nonetheless, this has been used in many studies [124, 125, 126, 127, 128, 34].

### 1.5.5 Multi-Scale Simulations

Clearly, a lot of the limitations described above with simulating the chemical and material properties of graphene are associated with their scale. Researchers want to consider larger systems and probe longer timescales with greater accuracy. The main hurdle here is not simply the expense of using more processors to crunch the numbers in a massive simulation, but that codes rarely compute at double the speed when they are given double the computing power. This is referred to as 'strong scaling', and the difficulty in achieving codes that successfully strong scale has meant new approaches in simulation are being developed that are relevant to this work.

Coveney *et al.* have shown the power of a truly multi-scale method of material simulation [129]. Automating processes at different length scales that both run and parameterise each other in real time has opened the door to large simulations. In their study, DFT calculations fed all-atom MD which in turn was used to iteratively parameterise a CG simulation. Their system contained hexagonal clay platelets in a polymer matrix and clearly has many similarities with the graphene-polymer systems we have been considering.

Other approaches to extend the use of simulation to more tangible scales include using information from a large material simulation in a continuum model [130, 131]. Such models have been well developed by engineers and can predict properties for many materials and composites. Continuum models have already studied what the fracturing behaviour [28] and mechanical properties [132] of graphene composites might look like, but crucially these are based on methods that are not designed for graphene's subnanometre dimensions (this includes not accounting for sheet edges and sheet translucency [133]). Hadden *et al.* have got round this by simply ignoring the problematic edge effects, they simulate graphene sheets that extend through the periodic boundaries and used those simulations to parameterise an experimentally verifiable multi-scale simulation of graphene platelet reinforced carbon fibre epoxy resin [99].

## 1.6 Other 2D Materials

Graphene can claim the highest measured tensile strength, electrical conductance and aspect ratio of any 2D material, but the field of 2D materials has diversified since its inception in 2004 and there are several other materials of note that have their own advantages.

Clays have been studied as polymer reinforcements for many years [129, 33]. Clays are aluminosilicates which occur in hexagonal sheets. Naturally occurring clays have ionic substitutions in the lattice (e.g. Si for Al), charge-balancing counterions (e..g. $Na^+$) reside between the clay sheets. Ionic solvents therefore help to exfoliate the sheets and create a dispersion. Clays are usually limited to below 1μm

lateral dimension.

Hexagonal boron nitride (h-BN), has a hexagonal structure like graphene and consists of alternating boron and nitrogen atoms. Like graphene, it is difficult to isolate in its monolayer form [134]. It is an insulator, with a wide band gap, leading it to be referred to as 'white graphene'.

Molybdenum disulfide ($MoS_2$) contains a layer of molybdenum atoms and two layers of sulfur atoms, one above and one below. The molybdenum atom can bond to the sulfur atoms in both a trigonal prismatic and octahedral geometry. It occurs naturally in its layered form molybdenite, much like graphite. $MoS_2$ is part of a broader family of 2D materials known as transition metal dichalcogenide monolayers, $WTe_2$, $WS_2$ and $MoSe_2$ are typical examples.

MXenes are transition metal carbides with the form $M_{n+1}X_n$ [135]. They are created by etching out the 'A' layers from MAX phases typically with hydrofluoric acid [136]. For example, $Ti_3AlC_2$ is a MAX phase compound, made of hexagonally structured layers of $Ti_3C_2$ separated by layers of Al [137]; by carefully controlling the chemistry and temperature, the Ti and Al bonds are broken, isolating the $Ti_3C_2$ layers. The reaction results in a change in the surface chemistry, adding hydroxyl or fluorine groups. MXenes have high conductance and are promising candidates for energy storage and sensing devices, with a large array of morphologies and surface chemistries that are yet to be explored.

## 1.7 Summary

Strikingly, the feverish interest in graphene has resulted in little of practical use. The list of potential applications grows increasingly superlative but by most predictions, remain a long way off. Research by Young et al. [71, 81, 72, 76, 30, 79, 78, 65, 31, 18, 73] is leading the investigation into polymer-graphene composite dependencies and characterisation, which will surely help guide future experiments.

As for computational studies on graphene, the number of approaches used and systems being studied is growing quickly. However, the lack of validation and representative studies has held the field back from making any serious predictions

or discoveries. Exploiting the predictive power of computational methods to aid the design of graphene nanocomposites is currently not possible; simulation is still playing catch-up with experiment.

In this thesis, first I investigate the validity of using MD to study graphene. In chapter 3 I describe a new method for simulating graphene at the nanoscale. This led to discoveries relating to graphene's frictional and exfoliation properties, which I present in chapters 4 and 5. As graphene oxide is commonly used for material science I investigate how to best describe its nanoscale structure in chapter 6. I then attempt to address the problem of simulating graphene across multiple length scales by coupling two methods of simulation together in chapter 7.1.

Comprehensive methods of predicting the properties of complex materials, like graphene nanocomposites, do not yet exist. But a 'virtual lab' that could do this reliably is a very exciting prospect. In this thesis, I show the predictive power of computational modelling techniques and develop methodologies for generalising and expanding those in the future.

## Chapter 2

# Computational Methods

Several modelling techniques are discussed throughout this thesis and a technical introduction to the methods used is given below. The majority of the work revolves around molecular dynamics (MD), but other methods that offer a different balance of precision and efficiency, providing useful insight at different length and time scales, are also used. It is also important to put the methods used within the context of the other popular modelling methods being used in the field currently.

## 2.1 Molecular Dynamics

MD is a popular simulation technique for studying the configurations and motion of matter with atomic precision. Essentially, it is a tool that defines the way different point atoms interact (known as a forcefield), then steps the system through time according to Newton's equations of motion.

The majority of simulations discussed in this thesis were run using LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) [138]; a powerful and versatile chemical MD simulation program.

### 2.1.1 Potential Energy Functions

As part of our study, four all-atom force fields were compared against each other to test their viability for large and accurate graphene simulations. AMBER [139], OPLS [140], Dreiding [141] and COMPASS [142], were chosen as fast forcefields that have been used in existing articles on graphene simulation [88, 98, 143].

Given a chemical structure, each of these forcefields can produce appropriate

bonded potentials, non-bonded potentials and partial charges based on typical values for atoms and functional groups in similar environments. The functions can be based on empirical results or calculated using *ab initio* methods. These structures and functions are then fed into LAMMPS for simulation.

Each forcefield describes the interactions of atoms slightly differently and will be discussed in section 3.1. One of the simplest descriptions of atomic interactions—which many forcefields are based on—is described below as a way of introduction. The total potential energy of a system in a simulation, $V_{\text{tot}}$, is based on Lennard-Jones/Coulombic non-bonded interactions and harmonic penalties for bond-length and angle deviations:

$$V_{\text{tot}} = V_{\text{bonds}} + V_{\text{angles}} + V_{\text{pd}} + V_{\text{id}} + V_{\text{vdW}} + V_{\text{Coulomb}}. \tag{2.1}$$

The first four terms in equation 2.1 are bonded interactions. $V_{\text{bond}}$ is the bond energy and is represented by a simple harmonic potential:

$$V_{\text{bonds}} = \sum_{\text{bonds}} k(b - b_0)^2, \tag{2.2}$$

where $b$ is the bond length, $b_0$ is the reference bond length between atoms and $k$ is the force constant. Note that in LAMMPS, the factor of a half usually seen in equations describing a harmonic force is absorbed into the force constant.

Angle potentials are calculated using a similar harmonic potential, where $k^\theta$ is the force constant and $\theta^0$ is the reference angle, see figure 2.1:

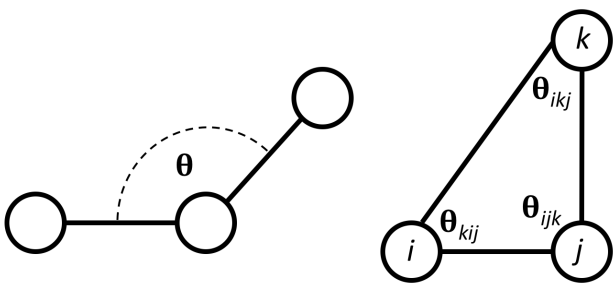


**Figure 2.1:** Explanation of the termination of $\theta$ in equation 2.3.

$$V_{\text{angle}} = \sum_{\text{angles}} k^{\theta}(\theta - \theta^{0})^{2}. \tag{2.3}$$

Torsional potentials for a proper dihedral (pd) are calculated by observing the angle created between two planes. Given four atoms bonded sequentially $ijkl$, the angle between planes created by atoms $ijk$ and $jkl$ is denoted $\phi$. $k^{\phi}$ is a force constant, $n$ is an integer and $d$ is either $+1$ or $-1$.

$$V_{\text{pd}} = \sum_{\text{proper dihedrals}} k^{\phi}[1 + d\cos(n\phi)]. \tag{2.4}$$

Torsional potentials for improper dihedrals are used to keep some atom groups in plane with each other or to stop molecules inverting to their mirror images:

$$V_{\text{id}} = \sum_{\text{improper dihedrals}} k^{\xi}(\xi - \xi^{0})^{2}. \tag{2.5}$$

For a central atom $i$, bonded to three other atoms, $jkl$, $\xi$ is the angle between the two planes containing atoms $ijk$ and $kjl$; and $\xi^{0}$ is its reference angle; $k^{\xi}$ is a force constant.

The last two terms in equation 2.1 define the non-bonded interactions in the system. Non-bonded potentials describe repulsive, dispersive and Coulombic interactions. The repulsive and dispersion potentials are combined into one Lennard-Jones potential, and by default do not need to be calculated for pairs where their interaction has already been accounted for by a bonded interaction:

$$V_{\text{LJ}}(r_{ij}) = \sum_{i=1}^{N_{\text{atoms}}} \sum_{j>i}^{N_{\text{atoms}}} \begin{cases} 0 & \text{if } i, j \text{ are connected by 3 or fewer bonds} \\ 4\varepsilon_{ij}\left(\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right) & \text{if } i, j \text{ otherwise} \end{cases}, \tag{2.6}$$

where $r_{ij}$ is the distance between atoms $i$ and $j$; $\sigma_{ij}$ is the distance at which the potential is at its minimum; and $\varepsilon_{ij}$ is the depth of the potential well. Usually, a forcefield will only define values for $\sigma$ and $\varepsilon$ for interactions between atoms of the same type, so when $i \neq j$, $\sigma_{ij}$ and $\varepsilon_{ij}$ can be calculated using either a geomet-

ric (equation 2.7) or arithmetic (equation 2.8) mean of coefficients defined in the forcefield:

$$c_{ij} = \frac{1}{2}(c_{ii} + c_{jj}); \tag{2.7}$$

$$c_{ij} = \sqrt{c_{ii} \cdot c_{jj}}. \tag{2.8}$$

Finally, the Coulombic potential between atoms with partial or full charges $q_i$, associated with them; $\varepsilon_0$ is the vacuum permittivity.

$$E_C = \sum_{i=1}^{N_{\text{atoms}}} \sum_{j>i}^{N_{\text{atoms}}} \frac{1}{4\pi\varepsilon_0} \cdot \frac{q_i q_j}{\varepsilon_r r_{ij}}. \tag{2.9}$$

The non-bonded terms become somewhat more complex when a cut-off distance is introduced, $r_c$, beyond which the force between two particles equals zero. This can help reduce the number of calculations at each step by ignoring small terms given by atoms that are far apart, but to smoothly bring the potential energy down to zero (i.e. give the function a continuous derivative) a smoothing-function, $S(r)$ must be used. Here we have used the '`lj/smooth`' function described by LAMMPS which computes a Lennard-Jones interaction up to an inner cut-off distance $r_{in}$ and new function between $r_{in}$ and $r_c$. We desire a new shifted potential, $V_s$, such that:

$$V_s(r_{ij}) = \begin{cases} V_{\text{LJ}}(r_{ij}), & 0 \leq r_{ij} < r_{in} \\ S(r_{ij}), & r_{in} \leq r_{ij} < r_c \\ 0, & r_c \leq r_{ij} \end{cases}. \tag{2.10}$$

The force must be smooth at all points, so we require:

$$\begin{aligned} S(r_{\text{in}}) &= -V_s(r_{\text{in}}), \\ S'(r_{\text{in}}) &= -V_s'(r_{\text{in}}), \\ S(r_c) &= 0, \\ S'(r_c) &= 0. \end{aligned} \tag{2.11}$$

Dashes denote differentiation in space. $S(r)$ is a is a third order polynomial calculated by LAMMPS derived from the conditions specified above.

$$S(r) = C_1 + C_2(r - r_{\text{in}}) + C_3(r - r_{\text{in}})^2 + C_4(r - r_{\text{in}})^3 \tag{2.12}$$

It should be noted that the smoothing function used here is specific to LAMMPS, and many other methods exist for bringing non-bonded interactions smoothly to zero. The details of the smoothing methods are often not specified in papers. It is commonly thought to not significantly change a simulation's behaviour, although the cutoff distance certainly does [144].

The optimum value for cutoff distances is system dependent and difficult to determine. We used an outer cutoff of 11 Å, and an inner cutoff of 9 Å, which are typical values. The inner cutoff must be chosen to be large enough so that the interaction is not overly affected by the smoothing approximation, but small enough so the function is not too steep, A difference of 2Å was chosen to balance these two criteria.

The Coulomb potential is calculated using equation 2.9 for $r_{ij} \leq r_c$. Beyond the cut-off distance interactions are calculated in reciprocal space using a particle-particle-particle-mesh (pppm) solver [145].

### 2.1.2 Time Integration

Now that all the potential energy functions are known, the force, $\mathbf{F_i}$, on an atom $i$, with position $\mathbf{r}_i$, can be calculated.

$$\mathbf{F}_i = -\frac{\partial V_{\text{tot}}}{\partial \mathbf{r}_i} \tag{2.13}$$

$\mathbf{F}_i$ is calculated from a sum of all non-bonded pair interactions (within the cut-off radius $r_c$), the pppm solver's long range charge interaction and bonded interactions (which includes terms that depend on 2,3 or 4 atoms). The process of MD simulation involves moving the particles according to Newton's equations of motion:

$$\frac{\partial^2 \mathbf{r}_i}{\partial t^2} = \frac{\mathbf{F}_i}{m_i}; \ \frac{\partial \mathbf{r}_i}{\partial t} = \mathbf{v}_i; \ \frac{\partial \mathbf{v}_i}{\partial t} = \frac{\mathbf{F}_i}{m_i}, \tag{2.14}$$

where $t$ is time, and $m_i$ and $\mathbf{v}_i$ are the mass and velocity of atom $i$ respectively. A standard velocity-Verlet algorithm will be used to integrate the equations of motion. This algorithm requires the positions, $\mathbf{r}$, velocities, $\mathbf{v}$, and accelerations, $\mathbf{a}$, of all atoms at time $t$. The algorithm will then increment the positions and velocities by $\Delta t$ simultaneously, according to the following relations:

$$\begin{aligned} \mathbf{r}(t+\Delta t) &= \mathbf{r}(t) + \mathbf{v}(t)\Delta t + \frac{1}{2}\mathbf{a}(t)\Delta t^2; \\ \mathbf{v}(t+\Delta t) &= \mathbf{v}(t) + \frac{\mathbf{a}(t)+\mathbf{a}(t+\Delta t)}{2}\Delta t. \end{aligned} \tag{2.15}$$

$\Delta t$ must be chosen carefully so that the simulation runs efficiently but still represents plausible motion, it must therefore be significantly slower than the fastest motion in the system. This is typically a carbon-hydrogen bond vibration. Aromatic CH bond vibrations are approximately 10 fs long. A time-step of 1 fs is chosen as no constraints were used.

The initial conditions need to be known: the exact way the initial atom coordinates are generated will be discussed later alongside the simulations; the initial velocities can be taken from previous simulations or random atomic velocities can be generated according to a Maxwell-Boltzmann velocity distribution. The accelerations are derived from the atomic positions alone, using equation 2.14.

### 2.1.3 Temperature and Pressure Coupling

The next step is to couple the simulation to temperature and pressure. The canonical ensemble (NVT), and isothermal-isobaric (NPT) ensemble are both used in this work and require different consideration. The abbreviations NVT or NPT refer to the parameters that are kept constant in the ensemble: number of particles, N; pressure, P; volume, V; and temperature, T. therefore allowing the total energy and either pressure or volume to vary. A simulation where there is no thermostat (the total energy of the system, E, remains constant) is called the microcanonical ensemble: NVE.

The absolute temperature of a system can be calculated from its kinetic energy $E_k$:

$$E_k = \frac{1}{2}\sum_{i=1}^{N} m_i v_i^2; \tag{2.16}$$

$$T = \frac{2E_k}{N_{\mathrm{df}} k_B}; \tag{2.17}$$

where $k_B$ is Boltzmann's constant $N_{\mathrm{df}}$ is the number of degrees of freedom in the system given by:

$$N_{\mathrm{df}} = 3N - N_c - N_{com}. \tag{2.18}$$

Here, $N$ is the number of atoms, $N_c$ is the number of constraints invoked; $N_{com} = 3$ because the velocity of the centre of mass of the system is removed.

The pressure is calculated using the pressure tensor $\mathbf{P}$, and the virial $\Xi$:

$$\Xi = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j>i}^{N} \mathbf{r}_{ij} \otimes \mathbf{F}_{ij}; \tag{2.19}$$

$$\mathbf{P} = \frac{2}{V}(E_k - \Xi); \tag{2.20}$$

$$P = \frac{\mathrm{tr}(\mathbf{P})}{3}. \tag{2.21}$$

where $N$ is the number of particles in the system; $E_k$ is the kinetic energy of the system; $\mathbf{r}_{ij}$ is the distance between particles; $\mathbf{F}_{ij}$ is the force felt between those particles; and $\mathrm{tr}(\mathbf{P})$, is the trace of matrix $\mathbf{P}$ i.e. the sum of all elements along the main diagonal. The virial $\Xi$ The simulations use the Nosé-Hoover temperature and pressure coupling algorithm [146, 147]. This algorithm introduces a friction factor into the equations of motion to correctly adjust the system temperature or pressure.

### 2.1.4 Energy Minimisation

Energy minimisation algorithms are useful to get a picture of a molecule's lowest energy and preferred conformer(s). In chemical simulation, minimisation can also be useful to bring a high energy system to a lower energy configuration e.g. given

an arbitrarily generated starting system with unrealistic intermolecular distances. In these studies the 'conjugate gradient' method was used. A simpler algorithm is 'steepest descent', which, in LAMMPS is implemented by moving the atom co-ordinates according an equation like this:

$$\mathbf{r}_{i,n+1} = \mathbf{r}_{i,n} + \frac{\mathbf{F}_{i,n}}{\max(|\mathbf{F}_{\mathrm{n}}|)} \cdot h_n; \tag{2.22}$$

where, $\mathbf{r}_{i,n}$ refers to the position of particle $i$ after $n$ minimisation steps, $\mathbf{F}_{i,n}$ is the total force on particle $i$, $\max(|\mathbf{F}_{\mathrm{n}}|)$ is the largest absolute force on any particle in step $n$, and $h_n$ is the maximum displacement, chosen initially to be very small. If the new positions, $\mathbf{r}_{n+1}$, give a lower energy structure, they are accepted and $h_n$ is increased for the next step, if not the new co-ordinates are rejected and $h_n$ is reduced. The calculation is repeated until $\max(|\mathbf{F}_{\mathrm{n}}|)$ is below a specified value. This process will often reach, and be 'trapped', in a local minimum; for complex molecules and systems it is not possible to search the entire potential hypersuface, so algorithms may not find the true global minimum.

The 'conjugate gradient' algorithm is similar to this but combines the direction chosen at each step with information from the previous step, according to the Polak-Ribiere formalism [148].

### 2.1.5 Simulation Specifications

Unless otherwise stated, all MD simulations presented in this thesis used the following set of parameters.

All simulations were carried out using the modelling software LAMMPS [138]. Simulations used periodic boundary conditions; the timestep was 1 fs; Coulombic interactions were calculated using a particle-particle particle-mesh with a precision of 0.0001 kcal mol$^{-1}$ Å$^{-1}$; the cut-off for Lennard-Jones interactions was 11 Å; bonds, angles, dihedrals and impropers were implemented as specified by the OPLS forcefield; CH intermolecular interactions were unchanged from the OPLS definition, only the graphene carbon intermolecular interactions were changed with GraFF.

Minimisations used a conjugate gradient method with a force tolerance of $1 \times 10^{-6}$ kcal mol$^{-1}$ Å$^{-1}$ and energy tolerance of $1 \times 10^{-6}$ kcal mol$^{-1}$.

NPT and NVT simulations used a Nosé-Hoover barostat and thermostat [149].

## 2.2 Density Functional Theory

No DFT calculations were performed as part of the research presented in this thesis, but the results of DFT calculations published in the literature formed the basis for much of the work in the following chapter [150, 151, 152, 94]. Therefore, the basics of DFT will be discussed here.

Amongst the many approximations made in typical MD models discussed above, perhaps the most egregious to chemists is that electrons are not directly considered. Two problematic consequences of this, that will become apparent in subsequent chapters, are that chemical reactions including bond-breaking cannot occur, and that charges are of a fixed value and centred on atoms.

DFT allows the consideration of electrons indirectly by operating on the spatial dependant electron density function. This results in a vastly cheeper calculation than comparative quantum mechanical calculations that include electron correlation like Hartree-Fock theory. The obvious difficulty is then modelling electron correlation and exchange and numerous approximations and empirical methods exist to calculate this energy. DFT falls into difficulty when calculating dispersion interactions for similar reasons and several ways to overcome this exist; however a discussion is beyond the scope of this thesis.

## 2.3 Coarse Grained Molecular Dynamics

This variant of MD is implemented in much the same way as normal MD, however individual atoms are grouped together into coarse grained beads. The size of these beads depends on the scale and precision of the system required but typically can be thought of as reducing a functional group into one bead. These beads are often spherical, though not always, and their interaction potentials with other beads can have directional dependance too.

Lowering the precision of the model in this way clearly reduces the number of

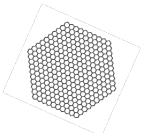

degrees of freedom and makes for a much faster equivalent calculation. However, parameterising the interaction potentials is very difficult and is often less transferable than typical MD force fields. This lack of transferability limits the forcefield to very specific chemical environments and even temperature ranges. Time resolution also becomes a problem, by coarsening the system one effectively reduces the friction between molecules, this leads to faster dynamics which are difficult to relate to real physical systems.

## 2.4 Finite Element Analysis

Finite element analysis (FEA) is a flexible class of methods for solving ordinary and partial differential equations (PDE) with complex boundary conditions. The method involves discretising the system into smaller parts, known as finite elements. The solution within each finite element is approximated by a set of shape functions. The element equations are recombined into a global system of equations that are solved by minimising an associated error function.

FEA will be used in chapter 7.1 where we use it to solve the conservation of momentum equation for objects undergoing engineering tests. To do this I use a library called deal.II [153], a comprehensive C++ library with many tools and routines for building and solving FEA problems.

## 2.5 Generating Starting Structures

Generating a starting point for a simulation is not a trivial task. A good starting configuration will resemble a physically realistic system and therefore reduce the time needed to equilibrate the system. Building a fluid particle simulation with a repetitive configuration, (e.g. identical molecules projected onto a lattice) is simple to generate but will need extra simulation time to remove periodicity and will initially not be dense enough.

### 2.5.1 Solvent Molecules

The approximate structure of small molecules can be generated by a molecular builder program, here Avogadro was used [154]. The user inputs the connectiv-

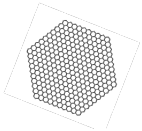

ity between atoms and the program uses a database of average bond angles and distances to generate a molecular structure.

Moltemplate, a module included in LAMMPS, was used to assign forcefield parameters to each atom in a unique environment. Topotools, a module within VMD (Visual Molecular Dynamics) [155] was used to convert between the Moltemplate output file to a LAMMPS input file. This structure and associated forcefield can then be minimised within LAMMPS for use in a simulation.

### 2.5.2 Graphene and Graphite

To generate different structures of graphene and graphite, a program was written to give a lot of customisability for the simulations [8].

The program can create graphene or graphite structures of any size, either crossing periodic boundaries, or with hydrogenated edges. Any number of structures can be combined into a single simulation input file.

Bond-lengths and interlayer-distances can be specified depending on the forcefield being used. All bonds, angles and torsions are generated by the program. Output is either a LAMMPS data file or in XYZ format for use with VMD.

## 2.6 Ensembles and Averages

Many simulations undertaken in this research are of a chaotic nature (their end state varies unpredictably with any perturbation in the initial state) so ensembles of simulations are required to understand the system. Unless otherwise stated, each replica is made unique drawing the atomic velocities from a Maxwell-Boltzmann distribution but are otherwise identical in their generation. A bootstrap with replacement [156] is usually performed on distribution of results gathered from the ensemble to quantify a confidence in any results that are derived. A bootstrap study provides an estimate of the confidence interval on an obtained average. This is similar to how a standard error, $\sigma_{\bar{x}}$, can be related to a confidence interval:

$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{N}}; \tag{2.23}$$

$$66\%\text{confidence interval} = \sigma_{\bar{x}}; \tag{2.24}$$

$$95\%\text{confidence interval} = 2 \times \sigma_{\bar{x}}. \tag{2.25}$$

However equation 2.23 only applies for normally distributed variables. The bootstrap method can give us information about a variable with an unknown distribution. For a dataset of size $N$, this is done by taking 'resamples' from the original dataset, at random with replacement, of size $N$. Typically, I took 10,000 resamples and the percentiles from the distribution of averages from all resamples gives the confidence interval.

## Chapter 3

# Developing A New Forcefield for Graphene and Graphite: GraFF

The initial remit of this research was to develop and deploy a multiscale approach to the modelling and simulation of chemically specific soft matter and graphene composite systems. Noting that the most significant roadblock to exploiting graphene's potential is producing it in high quantity and quality, I decided that attempting to simulate graphene exfoliation would be the first goal of my research. I set about studying the popular forcefields for performing MD simulations of graphene; as these all showed inadequacies, I decided to design my own which would allow fast and accurate simulations of graphene in exfoliation scenarios.

Liquid phase exfoliation of graphite perhaps holds the most potential for large scale commercial production of graphene [31, 32] (see section 1.1). It is a system which many hope computer simulation will help to understand and predict because of the difficulty in preparing the systems experimentally and the previous success of multi-scale modelling in elucidating anisotropic liquid systems [129]. Any new forcefield must be able to reproduce systems like this with a large number of atoms.

There have been a number of prominent papers which look at the mechanism of graphene exfoliation and the stability of graphene suspensions [98, 100, 101]. A review of the work carried out in this field has been written by Yang *et al.*[102]. The authors conclude that the field has not yet been able to help manufacturers because: (i) there is a lack of accurate force fields; (ii) too few solvent/surfactant systems

have been studied so general trends cannot be discerned; and (iii) there are poor coarse-grained models for amphiphilic molecules.

Shih *et al.*'s research [98] provides good insight into graphene dispersion stability in different solvents and exfoliation/aggregation mechanisms. Fu *et al.* have focussed on the exfoliation mechanisms of graphite [100]. These papers and other studies like them have all based their findings on traditional force fields developed primarily for biological systems and have used unphysical constraints to reduce computational cost. Both papers simulate graphene sheets that are rigid and cannot rotate, instead holding the flat sheets in place to investigate prescribed ranges of motion. I propose that investigating the forcefield is necessary for a unique material with electronic and mechanical properties vastly different from the small organic molecules these force fields were designed for.

Gobre and Tkatchenko [89] provide compelling evidence that the traditional practice of summing pairwise potentials to describe intermolecular interactions is insufficient in various carbon nanomaterials including graphene. The authors calculate the response of the fully interacting many-atom system by solving self-consistent screening equations, the potential being determined by assigning a quantum harmonic oscillator to each atom derived from the ground-state electron density obtained from DFT calculations. They discover unusual scaling laws for the intermolecular interaction behaviour of materials like graphene, graphite, carbon nanotubes and fullerenes.

Clay-polymer composites are a related system. The mechanism of exfoliating clay tactoids into a polymer has been seen with multiscale simulations [157]. These simulations show significant bending in the clay sheet and that flexibility is integral to the mechanism of exfoliation. Therefore, considering graphene exfoliation with fixed graphene sheets, like in other studies mentioned above, is not physically meaningful.

## 3.1 Existing Forcefields

It is common practice for graphene simulations to use a pristine graphene lattice terminated with hydrogens. This requires three types of atoms: graphitic carbon ($C_g$), benzylic-carbon ($C_b$) and benzylic-hydrogen ($H_b$). The two types of carbon atom only differ in their charge. As has become common practice, $C_g$ has been modelled as a neutrally charged $C_b$. This means that any polarity is only present on the very edge of the graphene flake.

I chose AMBER [139], OPLS [140], Dreiding [141] and COMPASS [142], as fast forcefields that have been used in existing articles on graphene. Note that it has become popular to simulate graphene using reactive forcefields such as ReaxFF [92] or AIREBO [93]. I chose not to include these in the following comparisons, as in general they add a 50× computational cost to an equivalent system, plus a smaller timestep often has to be used. This additional cost puts simulations of systems like liquid-phase exfoliation/dispersions out of reach, so were excluded them for our purposes.

Table 3.1 shows comparable parameters for the equations described in section 2.1.1. A non-bonded interaction cut-off, $r_c$, was not quoted by all papers so $r_c =$ 11.0 Å was used for all simulations.

### 3.1.1 AMBER

AMBER (Assisted Model Building with Energy Refinement) [139] has been developed primarily for biological purposes at the University of California. It has been used widely and successfully for modelling proteins and DNA. We use here the GAFF (General AMBER Force Field) subset, which provide parameters for small organic molecules.

The AMBER force field uses very similar forcefield to the one described in equation 2.1. It however includes half of the non-bonded terms (Lennard-Jones and Coulombic) between atoms separated by three bonds. This means that atoms $i$ and $l$ described by a proper dihedral interaction $V_{\mathrm{pd}}(\phi_{ijkl})$, also have a Lennard-Jones and Coulombic contribution weighted by a factor of 0.5.

| | AMBER | OPLS | Dreiding | COMPASS* |
|---|---|---|---|---|
| **Bonds** | | | | |
| $k_{CC}$ | 469.0 | 469.0 | 262.5 | 470.8361 |
| $r^0_{CC}$ | 1.404 | 1.400 | 1.390 | 1.4170 |
| $k_{CH}$ | 340.0 | 367.0 | 175.0 | 372.835 |
| $r^0_{CH}$ | 1.080 | 1.080 | 1.020 | 1.0982 |
| **Angles** | | | | |
| $\theta^0_{\text{all}}$ | 120 ° | 120 ° | 120 ° | 118 ° |
| $k_{CCC}$ | 85.0 | 63.0 | 25.0 | 61.0226 |
| $k_{CCH}$ | 35.0 | 35.0 | 25.0 | 44.3234 |
| **Dihedrals** | | | | |
| $k_{CCCC}$ | 5.3 | 7.25 | 3.125 | 8.3667 |
| $k_{CCCH}$ | 5.3 | 7.25 | 3.125 | 3.9961 |
| $k_{HCCH}$ | 5.3 | 7.25 | 3.125 | 2.3500 |
| **Non-bonded** | | | | |
| $\sigma_{CC}$ | 3.700 | 3.550 | 3.471 | 3.9150 |
| $\varepsilon_{CC}$ | 0.120 | 0.070 | 0.0951 | 0.0680 |
| $\sigma_{HH}$ | 3.080 | 2.420 | 2.420 | 2.8780 |
| $\varepsilon_{HH}$ | 0.100 | 0.030 | 0.0152 | 0.0230 |
| $\sigma_{ij}$ mean | arithmetic | geometric | arithmetic | equation 3.2 |
| $\varepsilon_{ij}$ mean | geometric | geometric | geometric | equation 3.3 |
| Partial Charges | 0.115 | 0.115 | 0.062 | 0.1268 |

**Table 3.1:** Parameters for graphene simulation with different forcefields. Only the first order dihedral coefficients are shown in this table. The definition of geometric and arithmetic means are given in equation 2.7 and 2.8 respectively; COMPASS uses 6th order combination laws, shown in equations 3.2 and 3.3.
*The COMPASS forcefield uses significantly different functional forms so one should not compare with these values. Bonded interactions go beyond a harmonic approximation, including cross-terms between bonds, angles and dihedrals.

### 3.1.2 OPLS

OPLS (Optimised Potentials for Liquid Simulations) [140] bears many similarities with the AMBER forcefield. It has gained popularity for being versatile and completely open source. We use the OPLS-AA (All-Atom) subset which considers each atom explicitly. As the name suggests, OPLS was developed to recreate experimental thermodynamic properties of liquids.

OPLS-AA implements the weighted non-bonded contributions to 1-4 atoms as in AMBER. The dihedral angles are described by a series of cosine functions as follows:

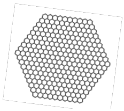

$$V_{\mathrm{pd}}(\phi_{ijkl}) = \sum_{n=1}^{4} K_n[1+(-1)^{n+1}\cos(n\phi)]. \tag{3.1}$$

### 3.1.3 Dreiding

Dreiding [141] was one of the first force fields to be developed. It was designed to be generic and transferable for small organic and main group inorganic molecules.

Dreiding forcefields are implemented as described in section 2.1.

### 3.1.4 COMPASS

COMPASS (Condensed-phase Optimised Molecular Potentials for Atomistic Simulation Studies) [142] was developed for a variety of molecules in the condensed phase.

The COMPASS forcefield includes several coupling terms between bond, angle and dihedral potentials. The simulations described in this report implemented the 'class2' force field which is described in full in reference [142]. The complicated coupling terms add many more variables with vague physical interpretations, and add more computation to each step but could describe important dynamics in the simulation. The COMPASS forcefield implements an unsmoothed 9-6 Lennard-Jones pair potential and includes full non-bonded interactions to 1-4 atoms.

The combination rules for the 9-6 Lennard-Jones parameters are given by:

$$\sigma_{ij} = \left(\frac{\sigma_{ii}^6 + \sigma_{jj}^6}{2}\right)^{1/6} \tag{3.2}$$

$$\varepsilon_{ij} = 2\sqrt{\varepsilon_{ii}\cdot\varepsilon_{jj}}\left(\frac{\sigma_{ii}^3\cdot\sigma_{jj}^3}{\sigma_{ii}^6\cdot\sigma_{jj}^6}\right) \tag{3.3}$$

## 3.2 Evaluating Existing Force Fields for Graphene

When parameterising a force-field it is vital to have experimental and higher level theory results with which to compare. Perhaps the reason no other attempt has been made at defining a parameterisation specific to graphene is because of the lack of such data. It is practically impossible to create a consistent distribution of graphene

**Table 3.2:** Comparison of experimental results for the energy of removing the top graphene sheet from a bulk graphite substrate

| Energy meV / atom | Experimental description | Reference |
|---|---|---|
| $61 \pm 5$ | Extrapolating thermal desorption spectroscopy data of small aromatic molecules from a graphite surface. | [158] |
| $31 \pm 2$ | Atomic force microscope study: sliding the top layer of a graphite stack over a step. | [159] |
| 35 +15/-10 | Transmission electron microscopy study of collapsed carbon nano-tubes. | [160] |

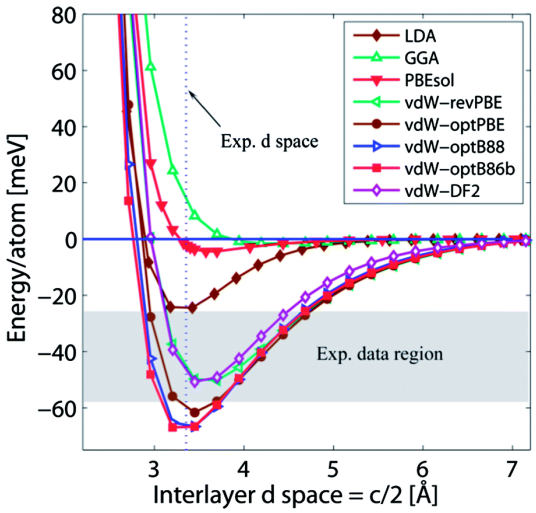


**Figure 3.1:** Comparison of different DFT functionals in representing the binding energy of a graphene sheet to bulk graphite.[94]

flake sizes experimentally; where a typical small organic molecule could be parameterised by comparing to its bulk properties e.g. in terms of density, enthalpy of vaporisation or solvation properties, such bulk properties are not available for graphene. Crystallographic data for graphite is all that is available: carbon-carbon distances of 1.421 Å and a layer-spacing of 3.354 Å. A force-field must also recreate graphite's minimum energy structure, where layers shift relative to each other in an alternating AB stacking configuration.

### 3.2.1 Interlayer Binding Energies

The best experimental measurement for the intermolecular forces that we can use to compare to simulations is the energy required to exfoliate a layer of graphene

**Table 3.3:** Results of an inter-layer binding simulation of graphite using different forcefields.

| Forcefield | Binding energy / meV | *d*-spacing / Å | C-C bond length / Å |
|---|---|---|---|
| Experiment | -66 to -20 | 3.354 | 1.421 |
| AMBER | -100 | 3.669 | 1.429 |
| OPLS | -55 | 3.483 | 1.415 |
| Dreiding | -81 | 3.396 | 1.384 |
| COMPASS | -61 | 3.355 | 1.391 |

from bulk graphite. Experimentally, results vary by huge margins as can be seen in table 3.2. A number of DFT techniques have been summarised by Wang *et al.* and are shown in figure 3.1. It can be seen that without specific consideration for long range dispersion forces, as arises in the GGA or PBE functionals, little to no energy barrier to exfoliation is observed [94]. From this we can see that so-called "higher-level" theory approaches also need to be treated with caution. Wang's analysis suggested that out of the functionals shown in 3.1, the vdW-optPBE [151] and vdW-DF2 [152] functionals were the most appropriate for studying graphite; this was based on the accurate lattice structure they reproduced, exfoliation energy and enthalpy of reactions.

To compare a classical force field to the results listed above a simple 8-layer graphite stack extending through the $x-y$ periodic boundaries was built containing 1920 atoms. The system size was chosen so that top layer 'saw' graphite within the vdW cut-off distance below it and there were no complications with periodic boundary considerations. The stack was separated by a vacuum in the $z$ direction. The graphite structure was minimised for each forcefield and then the energy was calculated as the top layer of the stack was raised. The results of this heavily constrained exfoliation and the obtained structural information are shown in figure 3.2 and table 3.3.

The classical forcefields tested here: AMBER, OPLS, Dreiding and COMPASS have all been used in published papers to study graphitic systems [102], yet clearly AMBER and Dreiding can be immediately discounted for their extremely large binding energies. OPLS and COMPASS have binding energies that fall into the experimental range and between the DFT results. The structural data match

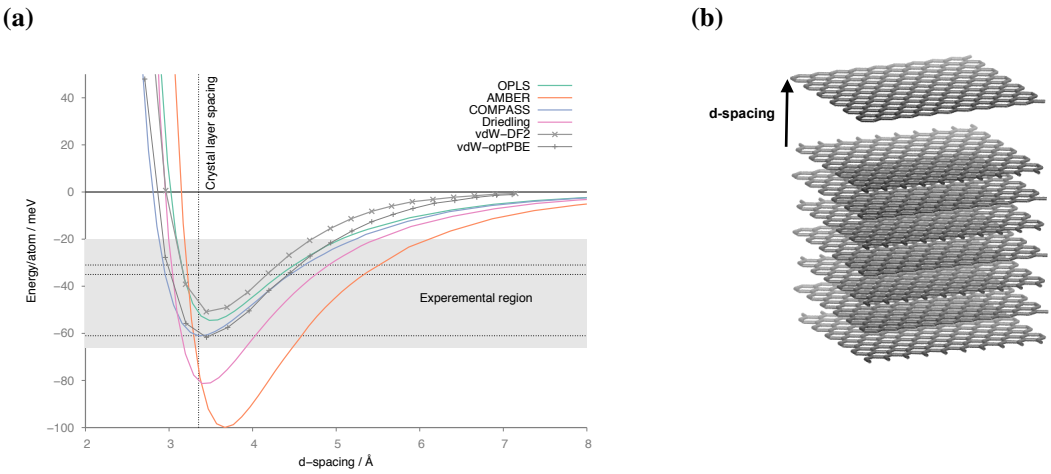


**Figure 3.2:** (a) shows comparison of the exfoliation energies calculated using existing forcefields (OPLS, AMBER, COMPASS and Dreiding), DFT calculations (vdW-DF2 and vdW-optPBE) and by experiment. The experimental region is very large. Classical forcefields predict a much larger long range interaction. (b) shows the simulation setup, the full stack contained 8 graphene layers.

fairly well to experimental readings, although AMBER predicts a larger than expected *d*-spacing.

### 3.2.2 Sliding Energy Barrier

Previous studies, for both graphene and clay cases, suggest that a sliding mechanism would be the most likely method of exfoliation [98, 157].

Experimentally, graphene is characterised by stick-slip behaviour which makes getting precise energies barriers for this motion difficult [161]. DFT methods have been able to provide some insight. Gao *et al.* have studied a case of two graphene sheets sliding past each other [150]. The group used a density functional approximation with a many-body dispersion interaction (DFA+MBD, following methods in Refs.[162, 163]). There is a small barrier to sliding, which is crucial for capturing the physical behaviour.

A simulation in LAMMPS was designed to replicate the above calculation, using the same forcefields discussed above. Two sliding regimes were tested: along both edges of the primitive unit-cell (i.e. in the armchair and zig-zag directions, see figure 3.3a). The forcefields fail to provide a significant barrier to sliding in the first section of the armchair direction, see figure 3.4a. The minimum at 0.33 (a displace-

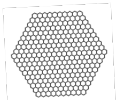

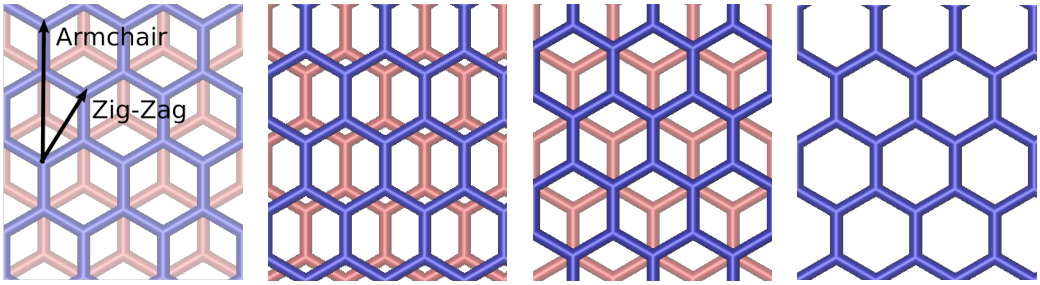


**(a)** AB **(b)** first maximum **(c)** BA **(d)** overlapping

**Figure 3.3:** Configurations of maxima and minima in figure 3.4a. (a) shows the direction the top sheet (blue) will move in the armchair direction (figure 3.4a), and the zig-zag direction (figure 3.4b). The other configurations correspond to points in figure 3.4a, at unit cell-fractions of: (a) 0.0; (b) 0.17; (c) 0.33 and (d) 0.67.

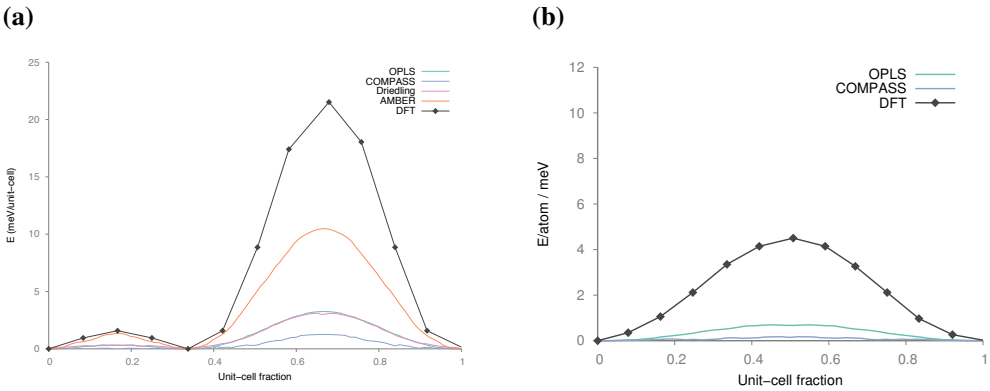


**Figure 3.4:** Energy barriers to sliding one graphene sheet over another in the two principle directions. (a) armchair direction; unit-cell fractions of 0 and 0.33 are symmetrically similar points. (b) zig-zag direction. For reference the experimental interlayer binding energy is -120 to -40 meV / Unit cell

ment of 1.4 Å) is a symmetrically identical configuration to the starting point (the graphite has gone from AB stacked to BA), therefore a graphene sheet could slide along an armchair-like path without encountering an energy barrier larger than the one seen between those two points.

**Table 3.4:** Results of a simulation where two sheets of graphene slide past each other using various forcefields.

| Forcefield | Energy barrier / meV |
|---|---|
| DFT [150] | 1.58 |
| AMBER | 1.35 |
| OPLS | 0.34 |
| Dreiding | 0.34 |
| COMPASS | 0.06 |

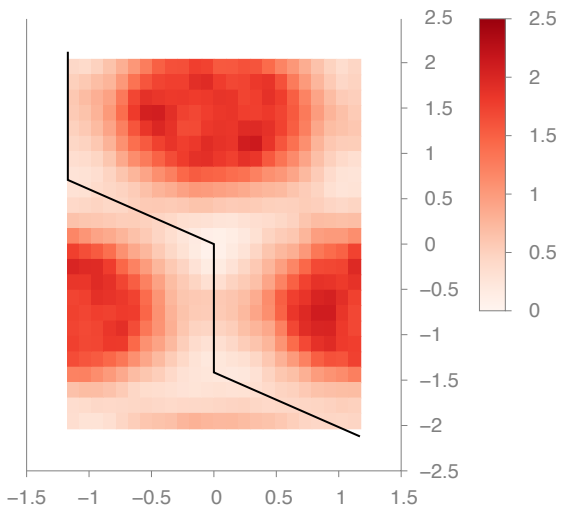


**Figure 3.5:** A statistical potential energy surface of two graphene sheets sliding on top of one another using the OPLS forcefield. The black line shows the minimum energy pathway from one global minimum to another symmetrically identical one. This low energy barrier allows sheets to zig-zag almost unimpeded in any direction.

This was highlighted by an MD simulation run at room temperature using the OPLS forcefield which exhibited sheets in a graphite stack sliding past each other unimpeded due to thermal energy alone. From this simulation, we can measure the relative positions of the sheets and therefore find the probability of finding the sheets at a certain displacment, $P(r)$. These values were accumulated into a histogram. The statistical potential, $V(r)$, was then derived using the following relation:

$$V(r) = -kT \cdot \ln(P(r)) - c, \tag{3.4}$$

where $k$ is the Boltzmann constant, $T$ is temperature and $c$ is a constant related to the partition function of the system. The statistical potential energy surface obtained from this simulation is shown in figure 3.5.

Similar behaviour was seen for all the forcefields we tested. At room temperature these forcefields cannot correctly describe graphene-graphene interactions as the sheets slide far too easily past each other.

## 3.3 A New Graphene Forcefield

It is clear that if one wants to simulate graphene exfoliation, the sliding motion needs to be correctly captured.

Simply changing the coefficients of the existing forcefields will not add a barrier to the sliding motion without increasing the adsorption energy. A 3-body dependent term is needed to add more detail to the potential energy surface above a graphene sheet.

Graphite is stacked in an AB configuration (some studies report that two layer graphene is AA stacked [150], but this is against general consensus [30]). In an AB structure, the smallest unit cell contains two atoms from each layer. In a given layer of each unit cell, one carbon lies directly inline with a carbon atom in the layer below, and one carbon sits above the centre of a hexagonal ring. A line between a pair of not inline carbon atoms creates a 23° angle with the basal plane, and are 3.64 Å away from each other. This information will be useful when parameterising the 3-body potential.

Angular dependent forcefields are often used to represent hydrogen bonds. A Dreiding description of a hydrogen bond is given by the equation:

$$V(r_{AD},\theta) = \begin{cases} 0, & 0^\circ < \theta < 90^\circ \\ V_{LJ}^{12-10}(r_{AD}) \cdot \cos^2(\theta), & 90^\circ < \theta < 180^\circ \end{cases}, \tag{3.5}$$

where:

$$V_{LJ}^{12-10}(r) = \varepsilon \left[ 5\left(\frac{\sigma}{r}\right)^{12} - 6\left(\frac{\sigma}{r}\right)^{10} \right], \tag{3.6}$$

$r_{AD}$ is the distance between the acceptor and donor, $V_{LJ}^{12-10}(r_{AD})$ is a 12-10 Lennard-Jones potential function and $\theta$ is the angle created by the acceptor-hydrogen-donor atoms. $\sigma$ and $\varepsilon$ are the zero-crossing distance and depth of the potential curve. This gives a maximum interaction when all three participating atoms are in line.

To add an angular dependence to the interaction between a pair of carbon atoms

in different graphene layers, $C_1$ and $C_2$, we use another reference carbon atom, $C_R$, bonded to $C_1$ (see figure 3.6). The most favourable interaction will be at 90° between the three atoms. We also would like to impose a cut off at 45° so that we avoid spurious interactions between atoms within the same layer.

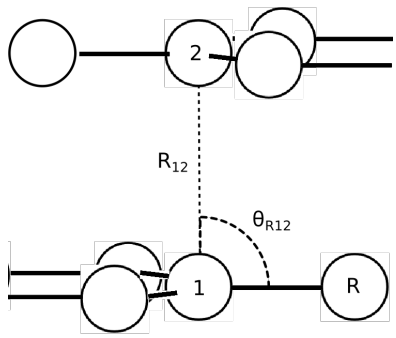


**Figure 3.6:** Schematic of the parameters used in the 3-body potential of GraFF. See equation 3.10 in the main text.

The conditions of the new Graphene ForceField (which we call GraFF), $V_{\text{GraFF}}$, are then:

$$V_{\text{GraFF}}(\theta_{R12}, r_{12}) = \begin{cases} 0, & 0° \leq \theta_{R12} < 45° \\ 0, & 135° \leq \theta_{R12} < 180° \end{cases}, \tag{3.7}$$

$$\frac{\partial}{\partial \theta} V_{\text{GraFF}}(\theta_{R12}, r_{12}) = \begin{cases} 0, & \theta = 45° \\ 0, & \theta = 135° \end{cases}, \tag{3.8}$$

$$\theta_{R12} = 90° \Rightarrow \min(\{\mathrm{V}_{\text{GraFF}}(\theta_{\mathrm{R12}}, \mathrm{r}_{12}) : \theta_{\mathrm{R12}} = 0°....180°\}), \tag{3.9}$$

where $\theta_{R12}$ is the angle made by $C_R\widehat{C_1}C_2$, and $r_{12}$ is the distance between $C_1$ and $C_2$.

To achieve this we weight a 12-6 Lennard-Jones potential by a $\cos^2(2\theta)$ term, so that the energy is at a minimum at 90°and is brought smoothly to 0 at 45°. This also makes the angular dependance of the potential resemble a $p$-orbital.

Note that the reference carbon atom is merely used to give the orientation. Without including it in the potential energy, I constructed the function below:

$$V_{\text{GraFF}}(r_{C_1C_2},\theta) = \begin{cases} 0, & 0^\circ < \theta < 45^\circ \\ V_{LJ}^{12-6}(r_{C_1C_2}) \cdot \cos^2(2\theta), & 45^\circ < \theta < 135^\circ \\ 0, & 135^\circ < \theta < 180^\circ \end{cases}, \tag{3.10}$$

where:

$$V_{LJ}^{12-6}(r) = 4\varepsilon \left[ \left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right]. \tag{3.11}$$

The forces on $C_2$ and $C_1$ are then:

$$\begin{aligned} \mathbf{F}_{C_2} &= \mathbf{F}_{\text{kernel}} + \mathbf{F}_{\text{angle}}, \\ \mathbf{F}_{C_1} &= -\mathbf{F}_{C_2}, \end{aligned} \tag{3.12}$$

where:

$$\begin{aligned} \mathbf{F}_{\text{kernel}} &= \frac{\partial}{\partial r_{C_1C_2}} V(r_{C_1C_2},\theta) \cdot \mathbf{r}_{C_1C_2}, \\ \mathbf{F}_{\text{angle}} &= \frac{\partial}{\partial \theta} V(r_{C_1C_2},\theta) \cdot \mathbf{r}_{tan}, \\ \mathbf{r}_{tan} &= (\mathbf{r}_{C_1C_2} \times \mathbf{r}_{C_1C_R}) \times \mathbf{r}_{C_1C_2}. \end{aligned} \tag{3.13}$$

where $\mathbf{r}_{C_1C_2}$ is a vector of unit length in the direction of $r_{C_1C_2}$.

This gives the force map shown in figure 3.7. While this definition fits the criteria we set for a potential it leads to some unphysical dynamics. One can imagine as two particles approach through space on unaligned trajectories, this potential creates an unimpeded rotational force. I had clearly overlooked some of the effects of such a potential.

I therefore implemented a potential which acts similarly to the Dreiding hydro-

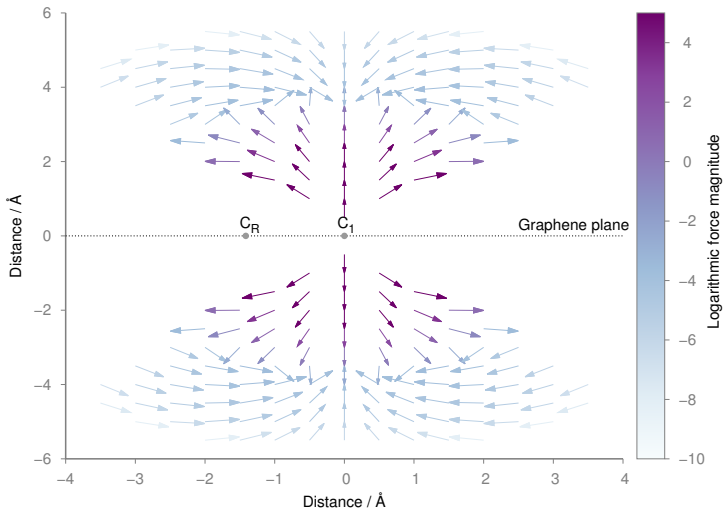


**Figure 3.7:** Force acting on atom $C_2$ calculated using equations 3.13, where $C_1$ lies at the origin and $C_R$ is at (-1.41,0).

gen bond potential [164] described above. This is a typical bonded angle potential which involves all three atoms (including $C_R$) in the resulting forces. $V_{\text{GraFF}}$ is now a function of $r_{R2}$, $C_R$ replaces the donor atom, $C_1$ the hydrogen atom, and $C_2$ replaces the acceptor.

$$V_{\text{GraFF}}(r_{C_RC_2},\theta) = \begin{cases} 0, & 0^\circ < \theta < 45^\circ \\ V_{LJ}^{12-6}(r_{C_RC_2}) \cdot \cos^2(2\theta), & 45^\circ < \theta < 135^\circ \\ 0, & 135^\circ < \theta < 180^\circ \end{cases} . \tag{3.14}$$

The forces are then:

$$\begin{aligned} \mathbf{F}_{C_R} &= \mathbf{F}_{\text{kernel}} + a_{11}\mathbf{r}_{C_RC_1} + a_{12}\mathbf{r}_{C_1C_2}, \\ \mathbf{F}_{C_2} &= -\mathbf{F}_{\text{kernel}} + a_{22}\mathbf{r}_{C_1C_2} + a_{12}\mathbf{r}_{C_RC_1}, \\ \mathbf{F}_{C_1} &= -(\mathbf{F}_{C_R} + \mathbf{F}_{C_2}), \end{aligned} \tag{3.15}$$

where:

$$
\begin{aligned}
\mathbf{F}_{\text{kernel}} &= \frac{\partial}{\partial r_{C_R C_2}} V_{\text{GraFF}}(r_{C_R C_2}, \theta)\mathbf{r}_{C_R C_2}, \\
\mathbf{F}_{\text{angle}} &= \frac{\partial}{\partial \theta} V_{\text{GraFF}/}(r_{C_R C_2}, \theta), \\
a_{11} &= a \cdot \cos(2\theta), \\
a_{12} &= -a, \\
a_{22} &= a \cdot \cos(2\theta), \\
a &= \mathbf{F}_{\text{angle}}/\sin(2\theta).
\end{aligned}
\tag{3.16}
$$

Now graphene atoms are simulated with a reduced 12-6 Lennard-Jones potential together with the new forcefield described above. Fitting to the sliding and exfoliating graphs discussed earlier (figures 3.2, 3.4a and 3.4b) the $\sigma$ and $\varepsilon$ values listed in table 3.5 were found to be the best fit.

This forcefield still favours the correct AB stacking in graphite. The $\cos^2(2\theta)$ term favours a 90° arrangement, so one might expect AA stacking. But the interaction between carbon atoms not directly above one another is only reduced to a factor 0.66; this reduced interaction still yields AB as the minimum energy structure because the 'shifted' carbon is above the centre of a ring, which means the interaction is between six atoms instead of directly above just one.

**Table 3.5:** Coefficients for GraFF simulations

| | $\varepsilon$ | $\sigma$ |
|---|---|---|
| $V_{LJ}^{12-6}(r)$ | 0.02 | 3.55 |
| $V_{\text{GraFF}}(r, \theta)$ | 0.025 | 3.354 |

The resultant force interactions calculated from equations 3.15 are shown in figure 3.8. Comparison to the other studies for exfoliating and sliding are shown in figures 3.9a and 3.9b.

Figures 3.9a and 3.9b show that the new forcefield is in good agreement with quantum-mechanical studies chosen by Wang *et al.* [165]. However it does not agree with quantum Monte Carlo calculations carried out by Fal'ko *et al.*, who predict the interlayer binding energy to be much lower [96]. The exfoliation energy

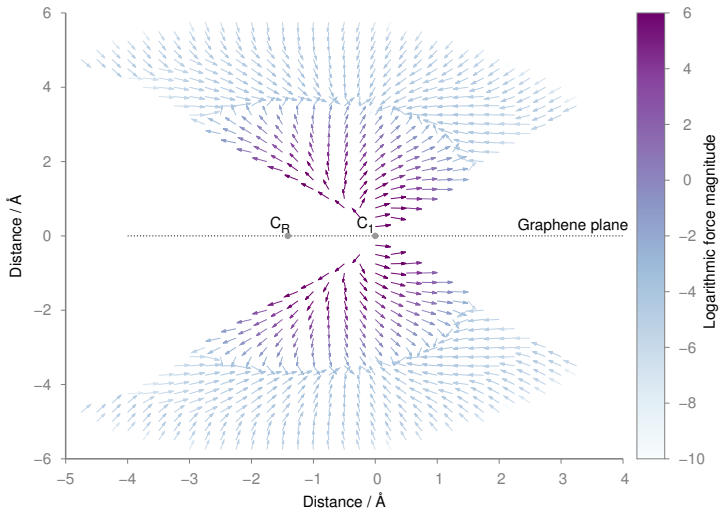


**Figure 3.8:** Forces acting on atom $C_2$ calculated using equations 3.15, where $C_1$ lies at the origin and $C_R$ is at (-1.41,0).

| Method | Adsorption energy [meV per atom] |
|---|---|
| **Experiment** | |
| Zacharia *et al.* [158] | $61 \pm 5$ |
| Liu *et al.* [159] | $31 \pm 2$ |
| Benedict *et al.* [160] | 35 +15/-10 |
| **Experimental range** | **25 — 66** |
| **DFT** | |
| vdw-DF2 [165] | 50.8 |
| vdw-optPBE [165] | 61.7 |
| **LAMMPS simulations** | |
| GAFF [139] | 99.9 |
| OPLS [140] | 54.5 |
| Dreiding [82] | 81.3 |
| COMPASS [142] | 61.0 |
| **GraFF** | **56.8** |

**Table 3.6:** Adsorption energy of the upper sheet of a graphite stack using various forcefields.

falls within experimental range too.

That the second peaks in figure 3.9b are not in very good agreement raises a moot point, there are significant discrepancies between the experimental values and their quantum calculations so there is no compelling reason to see a close agreement here. We can expect that this region would not be sampled much under normal conditions as it is higher in energy; accurately modelling the minimum energy pathway

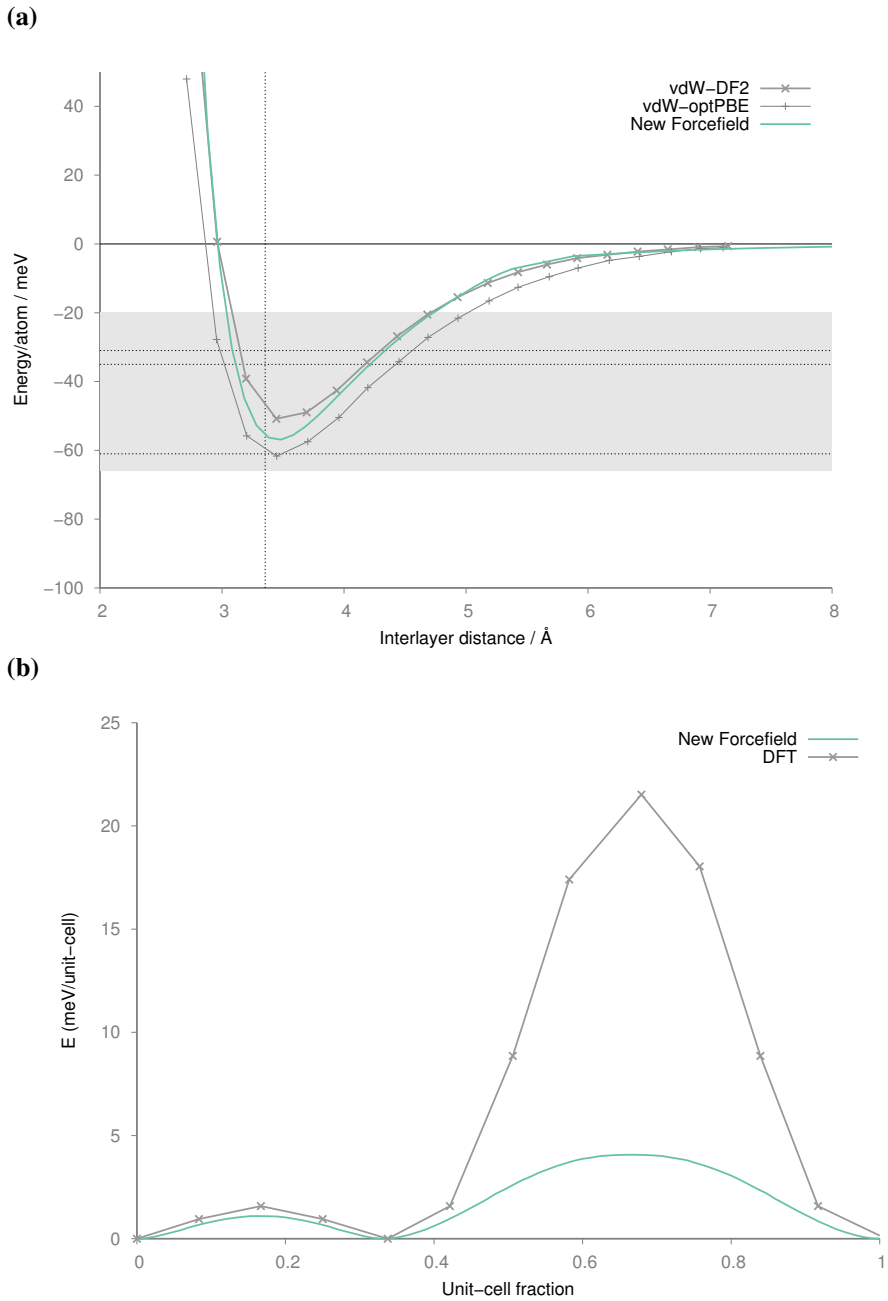


**Figure 3.9:** Comparison of the new forcefield described in equations 3.14 and 3.15 with (a) DFT exfoliation simulations (see figure 3.1) and (b) DFT sliding simulation (see figure 3.4a).

is deemed to be more important. To give some idea of the equilibrium population of the two flakes along the armchair direction of motion see Figure 3.10.

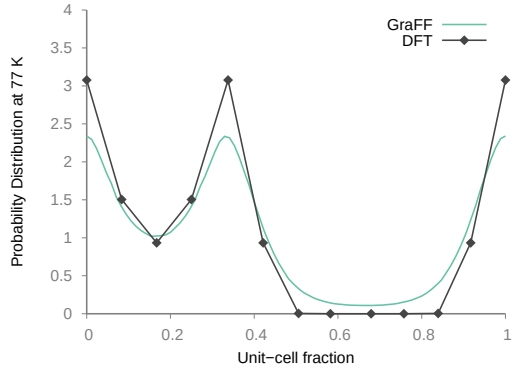


**Figure 3.10:** Equilibrium populations of two flakes along the armchair vector in Figure 3.3a we use the relation $P = \exp(-\Delta E/k_B T)$ where $\Delta E$ is given in Figure 3.9b. This shows that GraFF provides a sufficient energy penalty at unit-cell fraction between $0.45 - 0.85$ , leading to a very low population at these positions, matching the results from DFT.

In the original OPLS forcefield $\varepsilon_{OPLS} = 0.07$ kcal mol$^{-1}$ and $\sigma_{OPLS} = 3.55$ Å. Charges were $+/-0.115$ for terminating hydrogens and carbons respectively. For GraFF we combine a standard Lennard-Jones potential where $\varepsilon_{LJ} = 0.02$ kcal mol$^{-1}$ and $\sigma_{LJ} = 3.55$ Å with the weighted angular potential where $\varepsilon_{\text{Angular}} = 0.025$ kcal mol$^{-1}$ and $\sigma_{\text{Angular}} = 3.627$ Å; charges were unchanged. These parameters were chosen by recursive improvement of the adsorption energy and sliding potential, summarised in Tables 3.6 and 3.4. Increasing the weighting factor of GraFF increases the barrier to sliding while increasing the weighting of the ordinary Lennard-Jones parameter increases the adsorption energy.

Figure 3.11 illustrates the importance of capturing the smallest peak in Figure 3.9b. A flake can slide over another in the armchair pattern or zig-zag pattern shown in Figure 3.11a without encountering any higher energy barriers. By symmetry each segment in the paths shown in Figure 3.11a has the same energy barrier associated with it, shown in Figure 3.11b. The second peak in Figure 3.9b is due to completely overlapping sheets, i.e. AA stacking; this is much higher in energy and therefore unlikely to be explored.

We now have a description of graphene's intermolecular interaction that, ac-

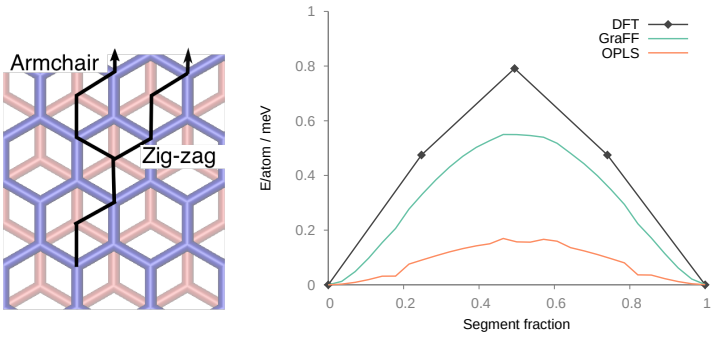


**(a)** **(b)**

**Figure 3.11:** Minimum energy pathways for a graphene sheet sliding over another. Each line segment in (a) has the same energy barrier, shown in (b).

cording to static calculations, recreates the energy barriers to adsorption and sliding. In the next chapter, I use this forcefield to recreate an experiment investigating graphene's friction. The motivation was to validate GraFF but some interesting conclusions were drawn about the mode of friction in atomically smooth materials. GraFF is used to describe graphene in chapters 4 and 5.

## Chapter 4

# Graphene Superlubricity

The previous chapter presented a new forcefield for describing graphene in atomistic simulations. Although DFT results were easy to compare with, validating the forcefield with experimental evidence was much trickier. Experimental absorption energies were available, but were far from in agreement with each other (table 3.3). Validating the sliding energy barrier introduced by GraFF was of course very important, I found a study by Feng *et al.* [161] which used a scanning tunnelling microscopy (STM) tip to propel graphene flakes along a graphitic bed and decided to see if GraFF would recreate their results. The results of this investigation, presented in this chapter, were reported first in Ref.[1]; the simulations and analysis presented are my own work under the supervision of James Suter and Peter Coveney. Along with finding good agreement with the experimentalists' results we were able to resolve the mechanism behind their observations and predict the consequences of the unusual properties of friction in nanoscale materials.

Graphite's lubricating properties have long been known, the simple explanation reads that the lubrication properties are due to the 'weak' interactions between individual layers [166]. However, these interactions are not weak enough to allow graphite to readily exfoliate into graphene on a large scale [42]. As has been discussed throughout this thesis, separating graphite layers down to a single sheet has become an intense area of research as scientists attempt to utilise graphene's superlative properties [32]. We have an intuitive understanding of friction on the macro-scale but have little understanding of the mechanisms involved in nano-

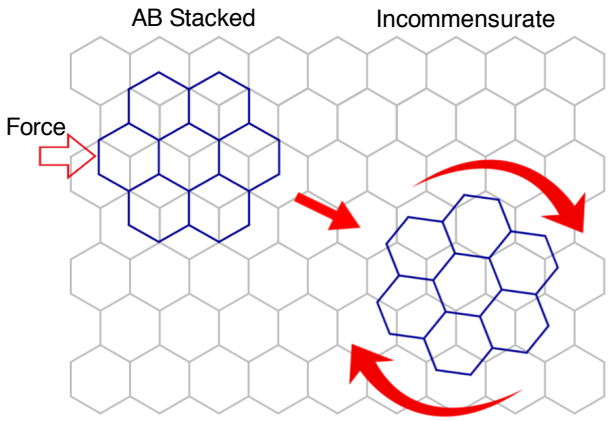


**Figure 4.1:** Schematic of the simulation setup. The flake, in blue, is pushed with a continuous force until it has been dislodged from its commensurate AB stacked configuration. Once free it is able to slide and rotate, essentially frictionlessly.

layered materials [167, 168, 169].

In a study by Feng *et al.* [161], an STM tip was used to push graphene flakes over a graphitic substrate to examine the frictional behaviour of graphene nano flakes. The flake was initially in an AB stacking orientation, the lowest energy configuration for graphite layers (see Figure 4.1). The authors found that once the graphene flake was pushed out of this preferred arrangement, the superlubric nature of graphene caused the flakes to 'slip' large distances of many times their diameter after they had been dislodged, before becoming 'stuck' in another commensurate position. They found that the flakes on average travel further at lower temperatures, which they attributed to the reduced thermal fluctuations in the substrate. If I could recreate this behaviour with GraFF we would have a lot more confidence in its ability to simulate graphene accurately.

## 4.1 Methods

My approach was to build a MD simulation that closely represented Feng's experimental setup: a graphene flake on a graphitic bed which is then pushed from its commensurate AB stacked position, then observed its dynamics until it again comes to rest. A simplistic schematic is shown in Figure 4.1.

Starting structures were built using the program described in Section 2.5.2 [8]. Graphene flakes were terminated with hydrogen atoms and were laterally 10 nm at

their widest point, containing 2520 atoms. Flakes of size 10nm are at the lower end of the range of flakes used experimentally by Feng *et al.*[161]

In the experiment, the substrate was epitaxial graphene on Ruthenium (001). Simulating metal surfaces is impossible with our chosen forcefield so I decided to approximate the substrates's properties. From STM images given in the paper it is clear that the substrate is not always a graphene monolayer but contains several layers. I then chose to simulate the substrate in two ways, one as a free standing graphene sheet, and one as a graphite stack. The graphite stack was held in place with a flat wall potential, which uses a 12-6 Lennard-Jones potential to provide a force perpendicular to the bottom plane of the simulation box. The graphite substrates were simulated with four graphene sheets stacked in an ABAB arrangement. The wall potential was matched to that of graphite in OPLS: $\varepsilon = 0.07$ kcal mol$^{-1}$; $\sigma = 3.4872$ Å.

Simulation dimensions with graphite substrates were $15 \times 15 \times 3.5$ nm; simulations with graphene substrates were $15 \times 15 \times 2.5$ nm, giving a vacuum spacing of at least 15 Å. For reference, in all simulations the substrate extended throughout the *xy* plane, with a vacuum space in the *z* direction.

Once these structures were generated, their structures were minimised within LAMMPS, with subsequent initial velocities generated from a Maxwell-Boltzmann distribution and equilibrated in the NPT ensemble for 500 ps at the required temperature.

Then I had to find a way to mimic an STM tip pushing a graphene flake from its commensurate position. A rough estimate provided by the authors for the force applied to the flake by the tip was 500 pN for a typical flake of 4000 atoms. They also estimated that the tip displaces the flake initially by 2 Å corresponding to an energy of the order of 1 eV. Due to the nature of the experimental setup, this interaction is very difficult to characterise accurately and there is no guarantee of consistency between repeat experiments. The nature, magnitude, and duration of the force applied to the flake are not known precisely. In addition, the graphitic substrate is also far from perfect as it is in our simulations, which may impact the adsorption energy.

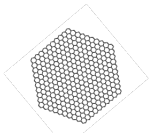

In my simulations I simply chose the smallest force necessary to consistently displace the flake from a commensurate position. The force was applied until the flake had moved 2 Å from its starting position, copying that done in their experiments. The minimum force required was found to be f pN per atom or 10 nN in total for the flakes we used, made of 2520 atoms. This is comparable to the value estimated in experiment (500 pN). Below the minimum force of 4.2 pN per atom, the flake stays in its commensurate position; the flake heated up because of the applied force and passes energy to the substrate, which subsequently also increases in temperature but ultimately stays in the same position.

The only variable we can compare between our simulations and the experiment is the straight line distance travelled by the flake. This distance will be used to evaluate the simulation. It was calculated by measuring the distance travelled by the flake's centre of mass, relative to an atom in the underlying substrate to remove the effects of any drift from results.

After the initial equilibration, the rest of the simulation was conducted in the NVE ensemble. This ensemble was chosen as adding force to the atoms in the flake injects energy into the system which a thermostat would immediately try to remove by altering the velocities. The NVE ensemble preserves this uneven energy distribution so we can observe how the flake dissipates its energy to the substrate. The overall temperature of the simulation did not rise by much because the substrate acted as a heat sink for the simulation.

## 4.2 Results and Discussion

In this section, I will first discuss the unusual dynamics of a graphene flake in motion over the substrate. I show that the results compare well with the experimental observations, and give hypotheses for the key mechanisms controlling the flake's behaviour. Once we are confident in our simulation setup we can probe other scenarios like collisions and exfoliation.

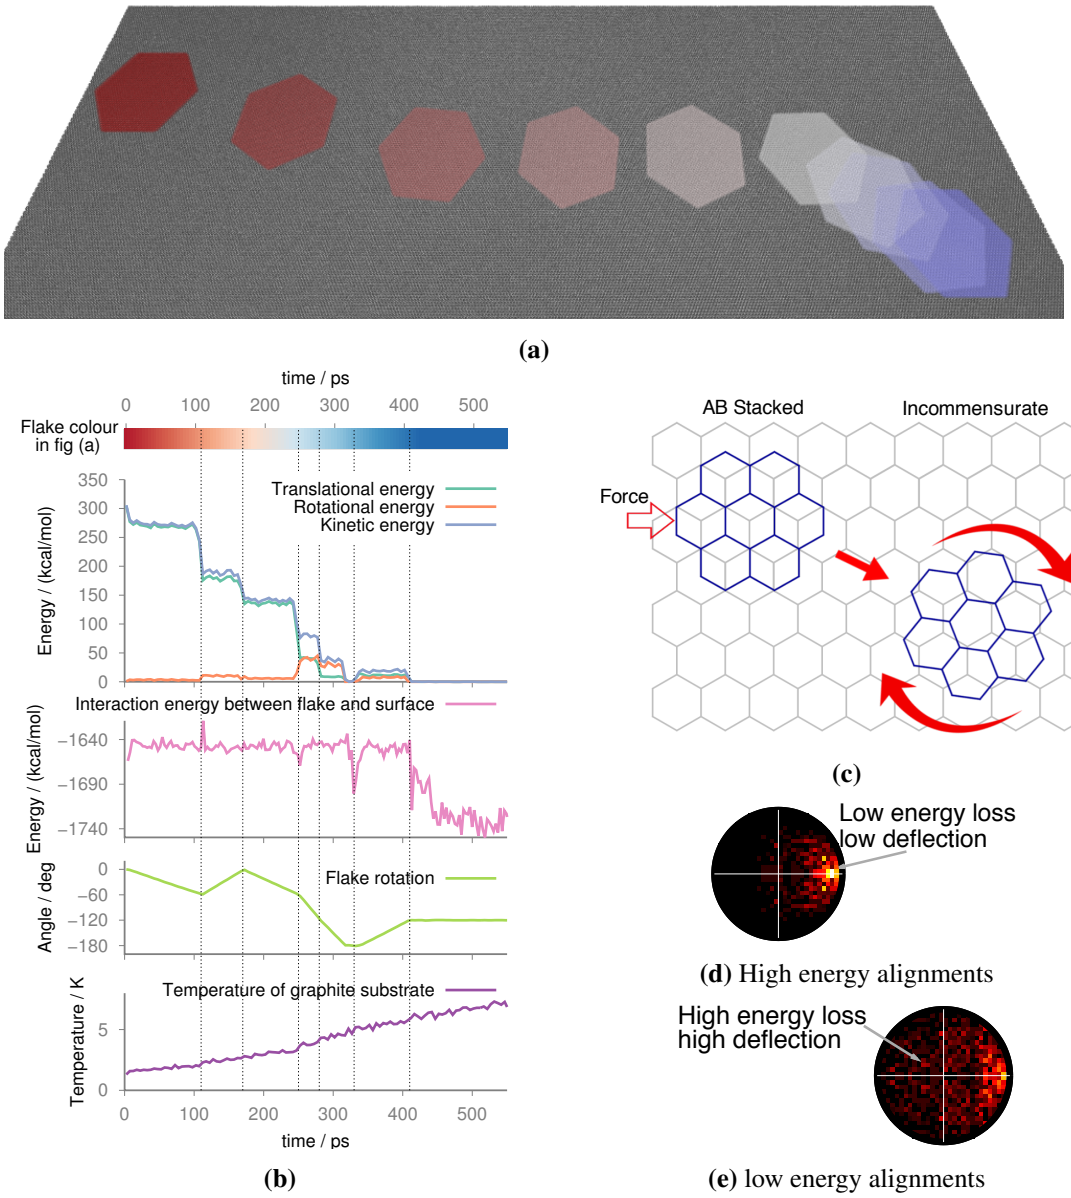


**Figure 4.2:** (a) A representative trajectory from a large ensemble of simulations of a 10 nm graphene flake on a graphite surface after being pushed out of a commensurate position. In this instance, the trajectory lasted 500ps until the flake was stationary. The flake travels from left to right, with colour corresponding to time (red at the start, blue at the end). A snapshot of the flake is shown every 50 ps. The flake slides and rotates freely when unaligned with the surface lattice and is only deflected when it is aligned. The internal kinetic energy of the flake is shown in (b). The flake loses energy which is dissipated to the substrate in alignment events which are represented by dashed vertical lines in (b), until it comes to rest in a commensurate position. (c) Schematic of the simulation setup. The flake, in blue, is pushed with a continuous force until it has been dislodged from its commensurate AB stacked configuration. Once free it is able to slide and rotate, essentially frictionlessly. The angle of deflection and energy lost in each alignment event are compiled into histograms in panels (d) and (e). Alignment events are shown when the kinetic energy of the flake is (d) above 100 kcal / mol, and (e) below 100 kcal / mol.

### 4.2.1 Friction Mechanism

The simulations showed that the motion of a graphene flake propelled over a graphitic substrate is characterised by fleeting alignments with the underlying lattice during which the flake is scattered and energy is converted between translational and rotational energy. These alignment events are how friction is manifested on the nanoscale as energy is dissipated to the bulk during each event (see Figure 4.2b; the alignment events are indicated by the dotted vertical lines), causing the flake to slow down and energy to be transferred to the substrate. When in an incommensurate position, the flake 'sees' an almost uniform potential energy surface (PES) in the lateral plane [170], as shown by the flat interaction energy between the flake and the substrate in Figure 4.2b. The flake therefore slides and rotates unimpeded. Gold nanoparticles have been observed to show similar behaviour, where the particles travel ballistically with extremely low friction, a process which was dubbed superlubricity [169]. Whenever the flake's orientation is an integer multiple of 60° — due to the hexagonal symmetry of graphene — there is a fluctuation in the interaction energy between flake and surface; this is because, in the narrow window when the flake is aligned, the PES has a defined well. This results in a deflection of the centre of mass trajectory, and exchange of translational and rotational energy.

When the flake has a low enough kinetic energy the wells in the PES can trap the flake. The difference in the non-bonded energy between the flake and surface for commensurate and incommensurate arrangements is approximately 100 kcal mol$^{-1}$. When the internal kinetic energy of the flake (the blue line in Figure 4.2b) is above 100 kcal mol$^{-1}$, deflection after alignment events tends to be small; below, larger deflections are more probable as the flake moves more slowly until it encounters a commensurate position in which to settle. This is shown in figure 4.2d and 4.2e, where we plot the distribution of kinetic energies and deflections of a flake after alignment events, the radial distance is being given by $E_{\text{incident}}/E_{\text{deflected}}$, and the angle from the $x$-axis is the deviation in the centre of mass trajectory, i.e. a point on the circumference of figure 4.2d would have lost no kinetic energy after aligning with the substrate lattice, and a point on the left of the $y$-axis would have

been deflected by greater than 90°.

The histograms shown in figure 4.2d and 4.2e are a compilation of the deflection and energy loss caused by every alignment event in an ensemble of 180 simulations described in the next section. The flake is said to be 'aligned' if its orientation is within 10° of a multiple of 60°. The energy of a flake, $E$, is the sum of $T_{translational}$ and $T_{\text{rotational}}$. $E_{\text{incident}}$ is the energy of the flake when it first becomes aligned, $E_{\text{deflected}}$ is the energy immediately after or when the flake comes to a standstill. A point on the histogram is defined in polar coordinates: the angle is the deflection of the centre of mass trajectory during an alignment event and the the distance from the centre is $E_{\text{incident}}/E_{\text{deflected}}$. Histogram bins are $0.067 \times 0.067$. The colour ranges from 0 counts (black) to 50 counts (white)

An illustrative trajectory of a superlubric graphene flake is provided in an accompanying video `superlubricity.avi`. This clearly illustrates the motion of the graphene flake: initially moving in a straight line, early alignment with the underlying lattice creates small deflections in the path of the graphene flake, rotating the flake by 60° each time. As the energy of the flake is dissipated in these events, the flake slows down and the alignments create much larger deflections, with the graphene flake centre-of-mass motion almost representing a random path. Eventually the flake comes to a stop in a commensurate position.

It should be noted that standard forcefields that only use simple Lennard-Jones non-bonded interactions [88, 171] did not reproduce these effects. Our simulations show that, using these forcefields, the flakes slip after propulsion and slide for as long as we simulated them (after 10 ns the flakes had travelled over 1000 nm, already an order of magnitude greater than that seen in experiments), therefore fatally underestimating the friction.

In figure 4.2b the translational and rotational energies were calculated between 5 ps intervals to reduce noise. Using a higher resolution (0.1 ps, see figure 4.3) makes the spikes in the interaction energy between the flake and substrate more visible as these are very fleeting events. The higher resolution also gives a better indication of the energy dissipation to the substrate: the bond energy in the substrate

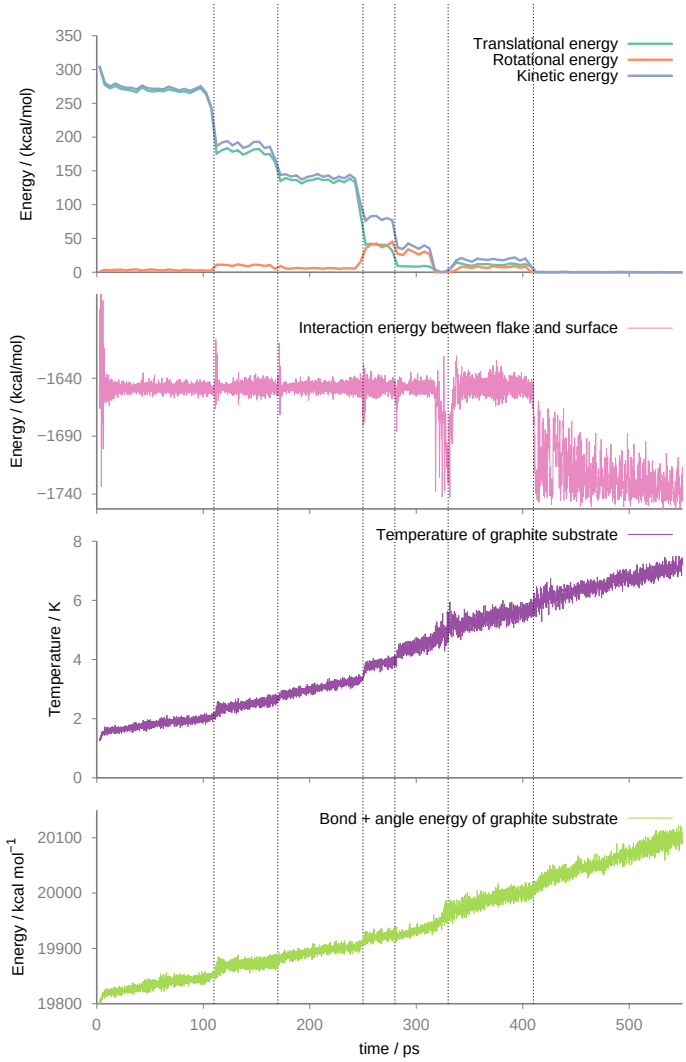


**Figure 4.3:** Energies of a 10 nm flake on a graphite substrate at 1 K, the same trajectory that is represented in Figure 4.2. The bottom 3 traces have a resolution of 0.1 ps.

jumps at each alignment event and steadily increases as energy is transferred from the flake.

### 4.2.2 Distances Travelled

The only experimentally measurable variable is the straight line distance travelled by the flake after propulsion. This is the characteristic quantity I will use to study this system and compare with Feng *et al.*[161].

Once I had first observed the unusual trajectory of a graphene flake and seen that the distance travelled was comparable to experiment it was important to run

ensembles of simulations to understand and get statistics on the system. At first I thought there may be some correlation between the angle at which the force was applied and the flake's trajectory. A large ensemble of 180 replicas was undertaken: three simulations for propulsion applied along each degree between 0° and 60°. This produced no discernible correlation; it also became apparent that this was a very chaotic system and the number of variables should be minimised in order to get reasonable results, so I stopped investigating the angle at which the force was applied.

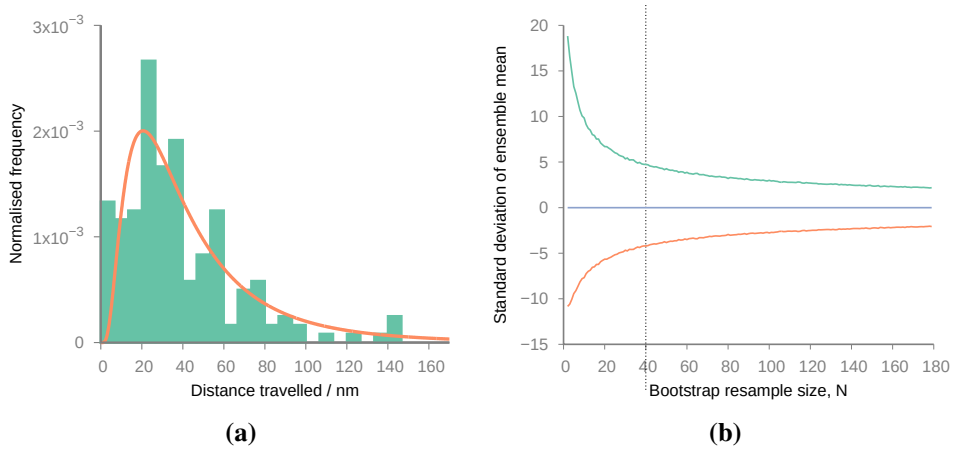


**Figure 4.4:** (a) Distribution of distances travelled by graphene flakes following propulsion. This histogram has been taken from an ensemble of 180 replicas and fitted to a log-normal distribution. (b) The 68% confidence interval of averages obtained from different bootstrap resample sizes. We deem the point of diminishing returns to be at a sample size of 40 replicas; after this point adding more replicas makes little difference to the uncertainties.

Using the ensemble of 180 replicas, I could determine an approximate distribution of results gained from this simulation and estimate the number of replicas it is preferable to simulate to minimise uncertainties of reported averages. A histogram of straight-line distances travelled by flakes within this ensemble is shown in Figure 4.4a. The quantity follows a log-normal distribution [172]. How the confidence interval depends on resample size when using the bootstrap method (see section 2.6) is shown in Figure 4.4b. $N = 40$ is indicated by a vertical dashed line; this is the size of ensemble (i.e. number of replicas we used in subsequent studies) that we concluded offers the best balance between computational efficiency and accuracy.

| **Experiment** | |
|---|---|
| Temperature [K] | Distance [nm] |
| 5 | 95 |
| 77 | 35 |

| **Simulation** | | |
|---|---|---|
| Temperature [K] | Graphite substrate Distance [nm] | Suspended Graphene Substrate Distance [nm] |
| 1 | 68.8 (+12.8/-10.8) | 47.9 (+5.0/-4.5) |
| 100 | 35.4 (+4.1/-3.7) | 29.3 (+5.2/-4.4) |
| 200 | 44.7 (+6.8/-5.9) | 25.9 (+2.8/-2.5) |

**Table 4.1:** Average straight-line distances travelled by a graphene flake on a graphitic substrate after being dislodged from a commensurate position. The flakes used in the experiment by Feng *et al.*[161] ranged from 8-18 nm while in our simulations we used 10 nm flakes. Surface roughness, caused by increasing the temperature or reducing the substrate's rigidity, increases friction. Errors are calculated from a bootstrap of 40 replicas explained in the methods section; errors were not reported for the experiment.

The standard error with 40 replicas is +4.7/-4.2 nm. We deem the point of diminishing return arrives at about $N = 40$. Quantities used in Table 4.1 are derived from ensembles of 40 replicas; errors are given as the confidence interval of one standard deviation, derived from a separate bootstrap on that ensemble with a resample size of 40.

Values given in table 4.1 are the arithmetic mean of the straight-line distance travelled. This is so that one can compare exactly with the experimental results and Feng *et al*'s averaging methods. However, for simulations we can also measure the total distance travelled, not just a straight line between the start and end points, taking into account the changes in direction after alignment events. This was calculated by taking the sum of the straight-line distance travelled by the centre of mass of the flake every 5 ps. The total distance travelled gives a better indication of the friction between the substrate and projectile. Also, the geometric mean is a more representative average of a log-normally distributed population. Table 4.2 characterises the same populations described in table 4.1, but using the geometric mean of the total distances travelled, which we believe to be the best representation of our data.

| Temperature / K | Graphite substrate Distance / nm | Suspended Graphene Substrate Distance / nm |
|---|---|---|
| 1 | 72.3 (+6.6/-6.0) | 49.6 (+4.2/-3.9) |
| 100 | 41.7 (+4.6/-4.2) | 25.2 (+3.8/-3.3) |
| 200 | 45.7 (+4.9/-5.5) | 31.6 (+3.4/-3.0) |

**Table 4.2:** Geometric means of the total distance travelled by a 10 nm graphene flake on graphitic substrates.

For a comparison of how using the total distance travelled changes the results compared to the straight line distance see Table 4.3, which shows the geometric averages of the straight line distance travelled.

| Temperature / K | Graphite substrate Distance / nm | Suspended Graphene Substrate Distance / nm |
|---|---|---|
| 1 | 54.2 (+5.7/-5.1) | 40.1 (+3.9/-3.6) |
| 100 | 25.2 (+4.2/-3.6) | 19.1 (+3.2/-2.7) |
| 200 | 30.5 (+5.0/-4.3) | 20.0 (+2.7/-2.4) |

**Table 4.3:** Geometric means of the straight-line distance travelled by a 10 nm flake.

The distributions of the total distance travelled that formed the averages in Table 4.2 are shown in Figure 4.5. The distributions of the straight-line distances that formed the averages given in Table 4.1 and Table 4.3 are shown in Figure 4.6.

#### 4.2.2.1 Finite Size Effects

| Temperature / K | 15 × 15 nm Substrate Distance / nm | 30 × 30 nm Substrate Distance / nm |
|---|---|---|
| 1 | 68.8 (+12.8/-10.8) | 58.8 (+4.2/-3.9) |
| 100 | 35.4 (+4.1/-3.7) | 27.5 (+4.9/-4.1) |

**Table 4.4:** Size effect simulations to be compared with table 4.1. Flake propelled on a graphite substrate, values are the arithmetic means.

So far the results reported have been for the distance travelled for a graphene flake on a substrate of 15 × 15 $nm^2$ in size. To verify that the trends we observe are not influenced by finite size effects we repeated the simulations of 40 ensembles at 1 and 100 K with a graphite substrate using a much larger substrate of 30 × 30 $nm^2$, while keeping the size of the flake identical. As can be seen in Table 4.2 the

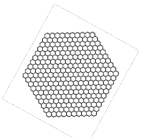

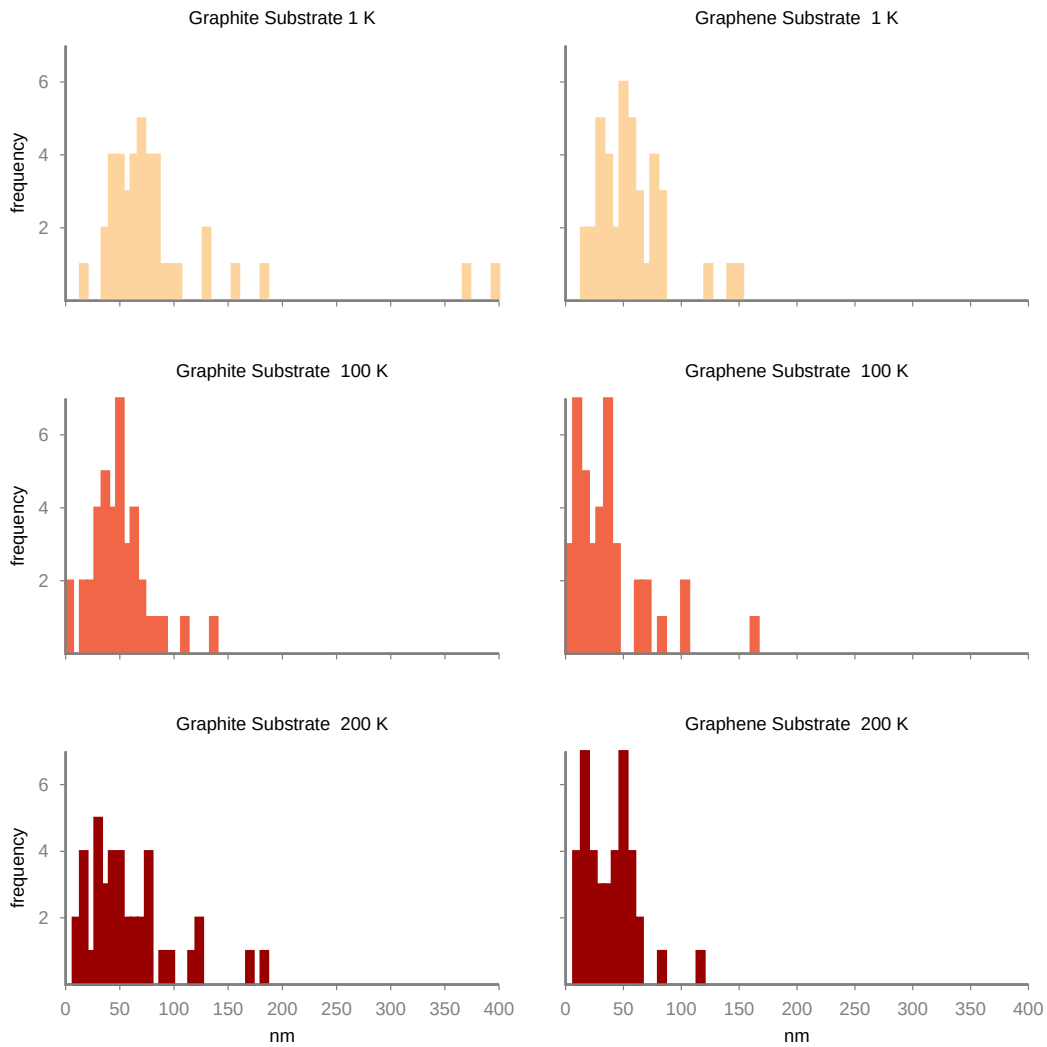


**Figure 4.5:** Distributions of total distances travelled by a graphene flake after propulsion on different substrates. Two outliers in the 1 K graphite simulation do not disproportionately affect the overall temperature trend we observe. Including the outliers, the geometric mean is 72.4 nm; ignoring them, the mean is 66.8 nm.

temperature trend is preserved at the larger size and the distances travelled are of comparable value. There is some sensitivity to system size, which may be due to the presence of longer wavelength undulations, but this effect is within statistical error.

The distributions of distance travelled for the $30 \times 30$ nm$^2$ substrate are given in Figure 4.7.

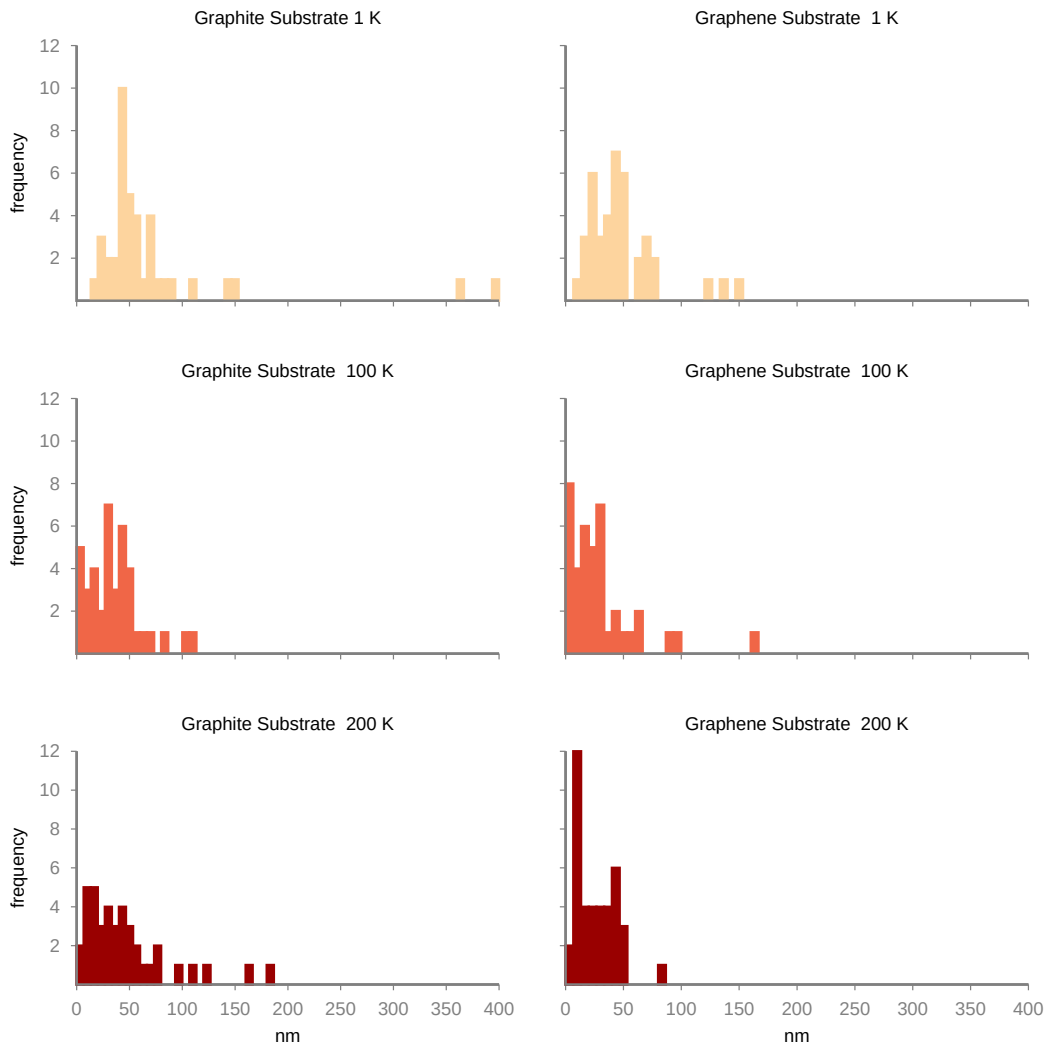


**Figure 4.6:** Distributions of straight-line distances travelled by a graphene flake after propulsion on different substrates.

### 4.2.3 Temperature and Surface Roughness Dependance

The simulations, recreated the trend observed by Feng *et.al*: the straight-line distance travelled by the graphene flake is much further at lower temperatures. This is the case whether the substrate is graphite (modelled as 4-layer graphene) or a suspended single graphene sheet (see Table 4.1).

The geometric averages in Table 4.2 indicate that the temperature trend — that flakes travel further at lower temperatures — observed in experiment is only a low temperature effect (<100 K). From 100 to 200 K we see a slight increase, although within error, in the distances travelled. The increased internal energy at

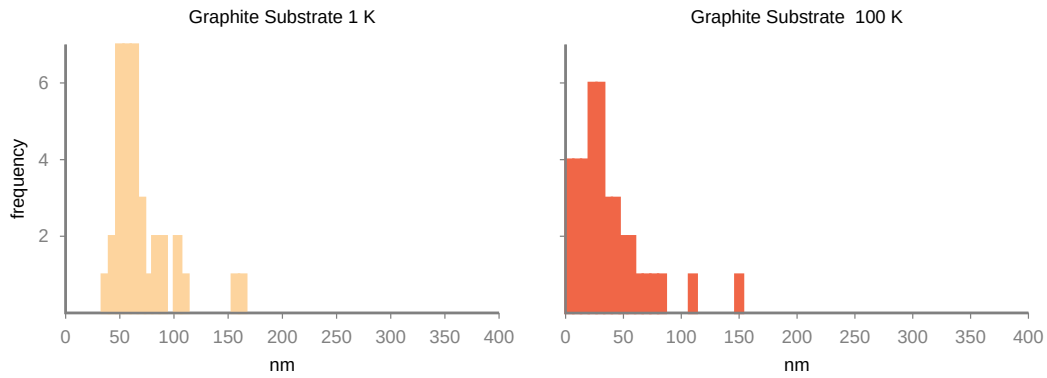


**Figure 4.7:** Distributions of total distances travelled by a graphene flake after propulsion on a graphite substrate which is 30 × 30 nm.

higher temperatures of the flake and substrate mean the flake does not settle as quickly, increasing the distance travelled (i.e. towards the end of the flake's motion, at higher temperature it has more energy to move out of the defined commensurate potential energy well).

As is shown in Table 4.1, temperature affects the distance travelled by the flake. Higher temperatures induce a range of frequencies of undulation in the surface (see Figure 4.8b) which impede a graphene flake from sliding across them. This is illustrated by the kinetic energy of the graphene flake at higher temperatures, which decreases between alignment events due to dissipation of energy to the graphite substrate. The total distance travelled by a flake at higher temperatures is therefore lower. This effect is more pronounced when flakes travel over a suspended graphene sheet, rather than over graphite, since thermally induced undulations are larger in the more flexible substrate, as can be seen in the greater amplitudes in the Fourier transforms of the surface height function given in Figure 4.8b. Our results indicate that the temperature trend observed in experiment is only a low temperature phenomenon. Between 100 to 200 K we see no change within statistical error in the distances travelled. This is due to a competing effect where the increased internal energy at higher temperatures of the flake and substrate means the flake does not stop as quickly, increasing the distance the flake travels. We tested our system for finite size effects in the previous section; by doubling the substrate size we show that trends are preserved and numerical values are within error for a larger system.

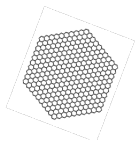

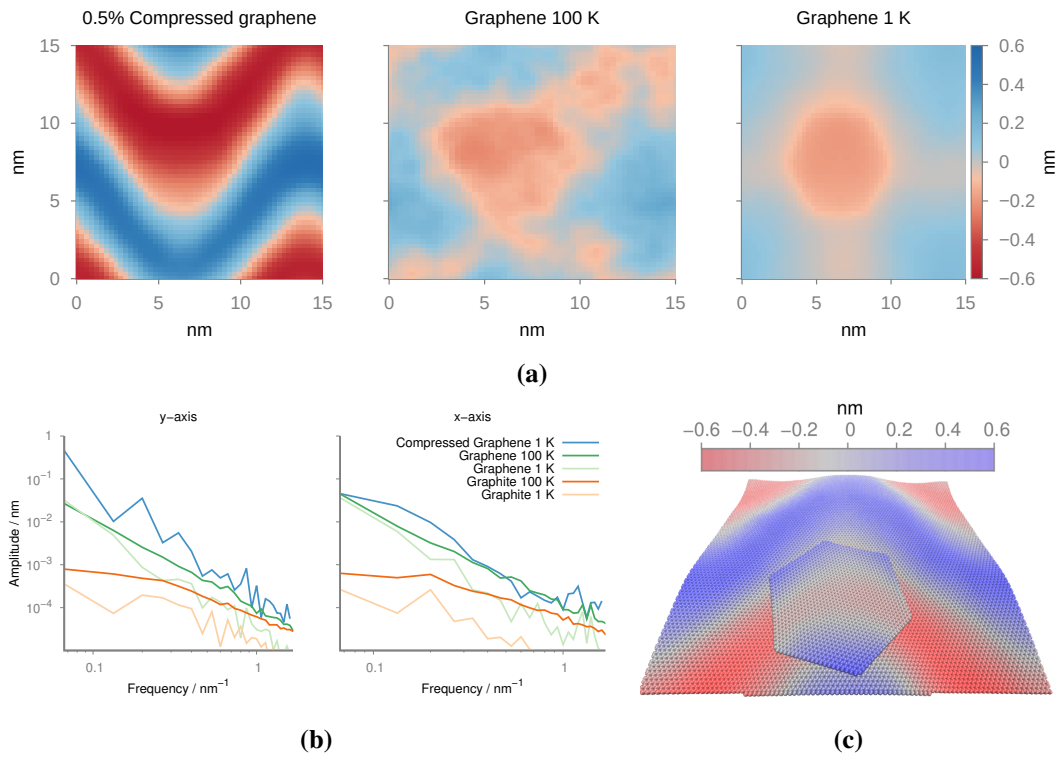


**Figure 4.8:** (a) Surface roughness for different suspended graphene simulations. Compressing the substrate produces large smooth undulations. Higher temperatures produce random fluctuations. At 1K there are so few fluctuations that the puckering of the surface due to the graphene flake can be seen. (b) Spectral intensity of a 2D Fourier transform of the different surfaces in the $x$ and $y$ directions. The frequencies are averages of all replicas in each ensemble. Graphite is clearly orders of magnitude smoother than the suspended graphene sheets. Vertical axes are identical. (c) A graphene flake 'surfing' in the trough of a travelling wave made in a compressed suspended graphene sheet. Colour represents the height displacement as in (a). Trapped in this trough, it does not align with the surface and so is free to slide without losing energy, see also the accompanying video: `surfing.avi`.

The Fourier transforms given in Figure 4.8b were calculated using the following method. The equilibrated starting coordinates for a replica were used, i.e. before the propulsion. The atoms in the top sheet of the graphite or all atoms in a graphene sheet were binned into a 50 by 50 array, with a bin size of 3 Å. The average $z$ displacement from the sheet's centre of mass in each bin created the height function array to transform. A two-dimensional discrete fast Fourier transform spectrum of the array was calculated. The amplitudes were obtained by dividing the spectral components along the $x$ and $y$ axis by the unit area of 2500. The $x$ and $y$ axes were

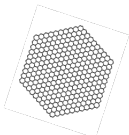

used as these are the only truly periodic components of the spectrum. The amplitudes given in the Figure are the result of averaging the Fourier spectrum for all 40 replicas in each ensemble.

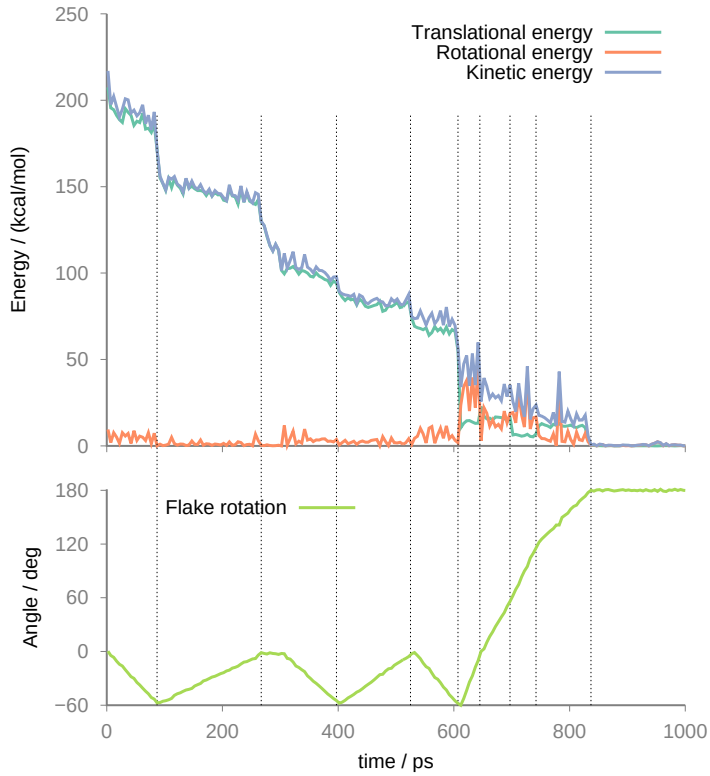


**Figure 4.9:** Energies and orientation of a 10 nm graphene flake on a graphite substrate at 200 K. Alignment events have a similar effect as in the lower temperature simulation shown in Figure 4.2b but between alignments the energy decreases as the flake encounters thermal undulations in the substrate.

This can be seen in figure 4.9 where in between alignment events the kinetic energy fluctuates and decreases. The increased noise can be expected from the thermal fluctuations. The decrease in energy between events shows that energy is dissipated to the surface at higher temperatures in a different way than at lower temperatures.

To confirm whether it is the undulations in the surface I also introduced undulations through compressing the substrate. In Figure 4.8a one can observe that compression introduces pronounced long wavelength traveling waves in graphene.

Experimentally, suspended graphene is known to produce dynamic ripples [173] which exhibit similar soliton-like behaviour to the waves in our simulation. The simulation was performed at 1K using a compression of 0.5 %, which is within previous experimental compression ranges [174].

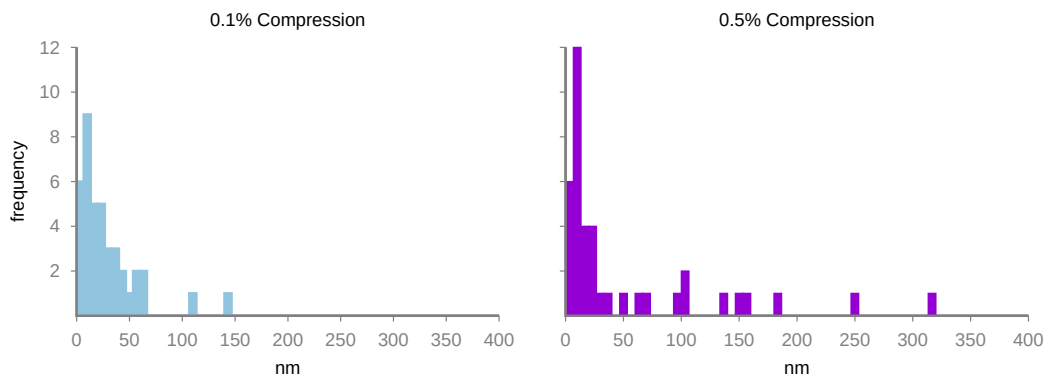


**Figure 4.10:** Distribution of distances travelled by a graphene flake after propulsion on a compressed graphene substrate.

The distribution of distances travelled by a flake after propulsion on a compressed graphene substrate is different to that found in the simulations with a graphite substrate. The majority of flakes stop sliding shortly after propulsion, by the mechanism of interacting with undulations and ending in commensurate positions, as described above.

However, a minority of flakes sit in the trough of a travelling wave in the substrate and are carried much further. See the accompanying video of such an event: `surfing.avi`. The distributions of distances travelled for different substrate compressions are shown in Figure 4.10 and Figure 4.11 gives a schematic description of the proposed mechanism.

While these waves impede the motion of the flake in the majority of simulations in our ensemble, another transport mechanism dominates in a small number of simulations, where the flake 'surfs' in the trough of travelling waves. On compressing a graphene sheet the distance travelled by a flake following propulsion is observed to increase. For this to occur the flake should be of comparable size to the wavelength of the undulations, and may provide a means of separating graphene flakes by size. A similar mechanism has been observed before in simulation of

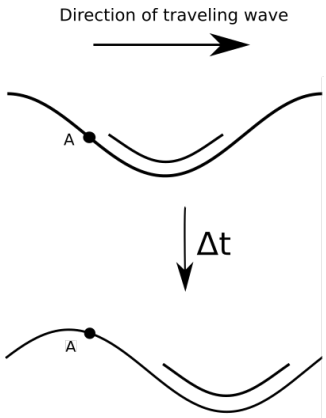


**Figure 4.11:** A graphene flake trapped in the trough of a travelling wave formed in a suspended graphene sheet. Atom A only moves vertically and is initially on the leading edge of the travelling wave. As atom A moves up, between the two snapshots, it interacts repulsively with the flake. Therefore the flake is forced along the wave's direction. This effect is maximised when the flake's width is half the substrate's wavelength, as in this instance, as an atom on the trailing edge of the substrate wave will move down, and into a more attractive region of the interatomic potential. See also the accompanying video `surfing.avi`.

water droplets on graphene [174], but not between 2D materials.

This mechanism could be a means to separate graphene flakes by size. More simulations would be required to test this theory out fully, but if one could control the wavelength of travelling waves in graphene substrate then flakes would be transported based on their size. If the mechanism described above is correct, flakes would preferentially travel if their size was close to $\lambda(n+1/2)$, $n \in \{0,1,2...\}$.

From the above simulations we have characterised three friction and transport mechanisms for superlubric graphene flakes sliding on a graphitic substrate. Alignment events transfer energy between (translational and rotational) kinetic modes within the flake, and between the flake and substrate. This mechanism is temperature independent and is only reliant on lattice matching of substrate and projectile. The second is due to short range undulations caused by thermal fluctuations in the substrate that lead to non-elastic 'collisions' with out-of-plane substrate atoms. Finally, we predict that, by compressing the substrate, one can introduce long wavelength undulations that can 'carry' a flake, vastly reducing the friction it experiences. Therefore, at low temperatures the dominating source of friction a

flake experiences would be due to alignment events; raising the temperature would increase the friction experienced. However, long range undulations may actually reduce friction if controlled properly.

### 4.2.4 Colliding Flakes

To further test the interactions of graphitic systems we simulated various collisions of flakes and different exfoliation scenarios.

The starting configuration was a single periodic graphene sheet acting as the substrate. The sheet was $15 \times 30$ nm$^2$ in size held in place with a flat 12-6 Lennard-Jones potential acting as a wall at the bottom of the simulation box, as used before. Two 10 nm diameter flakes were arranged along the long axis, separated by 15 nm. The configuration was equilibrated at 1 K.

Pushing the flake with the same procedure described in the above methods, led to the propelled flake bouncing off the stationary one, while losing a small amount of energy.

Pushing a flake with the same force as before but for 3.5 ps instead of 2.5 ps produced the simulation shown in the video `collision.avi`. The propelled flake is flexible enough and has sufficient energy to climb on top of the stationary flake. While on top no alignment events take place: the moving flake cannot align with the different substrates because they are AB stacked. The flake is propelled with significant force but does not dislodge the other flake.

Colliding flakes into each other highlights the flexibility of graphene. Given a force of 4.2 pN to each atom in a 10 nm flake for 2.5 ps — which is typical for the sliding simulations above — a flake collides and 'bounces off' another stationary flake. Given more force the flakes can slide on top or under one another. The AB stacking of graphite sheets means that the flake straddling another flake or a step defect cannot be fully commensurate with both layers; this reduces the contribution of alignment events to the friction experienced.

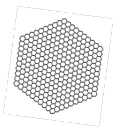

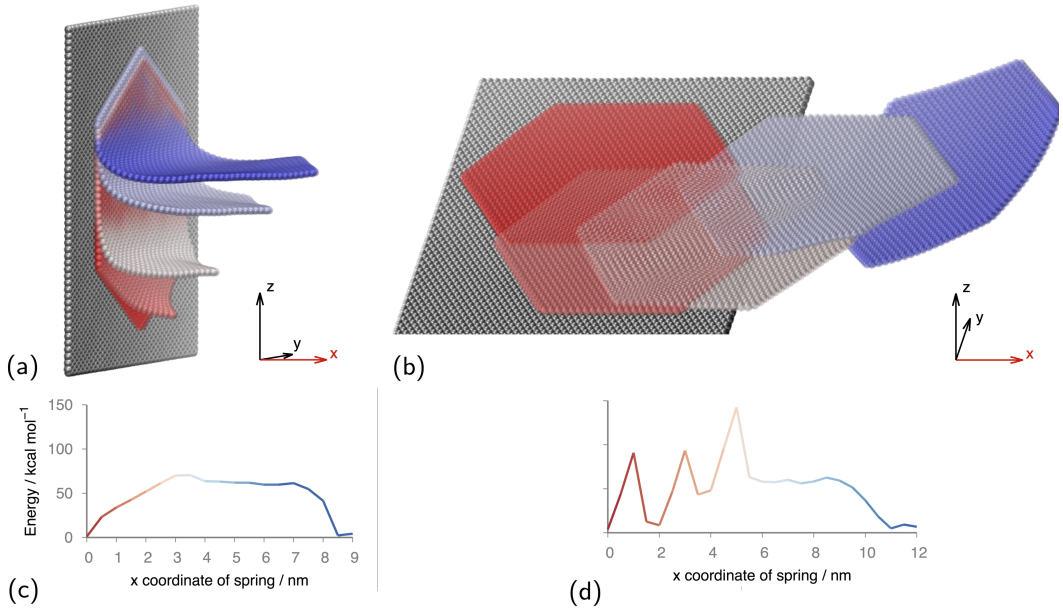


**Figure 4.12:** Steered simulations of two exfoliation mechanisms. 10 nm graphene flakes are pulled by a corner from a graphitic substrate with a harmonic spring along the $x$ axis. The spring was allowed to move in the $y$ and $z$ planes; the energy in the spring per atom is shown in plots (c) and (d). (a) and (c) exhibit a peeling mechanism. b) and d) manifest a shearing mechanism. Spikes in the spring's energy are due to the flake becoming trapped in commensurate positions with the substrate.

### 4.2.5 Steered Exfoliation Simulations

Next, I wanted to investigate if what we have learned about graphene's superlubric nature will affect the exfoliation mechanism of graphene. Figure 4.12 shows different regimes of exfoliating graphene flakes from a graphitic substrate. A harmonic spring, attached to a ring of six carbon atoms at the corner of the graphene flake, was used to induce exfoliation. Perpendicular configurations were used to compare peeling and shearing modes of exfoliation. A peeling mechanism mimics mechanical exfoliation with 'release tapes' [175]. Using a weighted histogram analysis method [176] we found that the work done to exfoliate the flake is 40% less via peeling than shearing, contrary to the common belief that a shearing mechanism would be more favourable [100].

The steered exfoliation simulations used similarly equilibrated systems in the NVT ensemble. The substrate was fixed so the substrate did not move with the

spring. The substrate was a 15 × 15 $nm^2$ graphene sheet. The peeling exfoliation had a simulation box of 15 × 15 × 12 $nm^3$; the shearing exfoliation had a simulation box of 22 × 15 × 2 $nm^3$ (see Figure 4.12), corresponding to a 12 nm vacuum space. The spring was attached to the 6 carbon atoms in a ring at the corner of the flake and had a spring constant of 5 kcal $mol^{-1}$. The spring was moved 1 Å every 100 ps along the exfoliation pathway and allowed to move freely in the other directions. This results in an effective pulling speed of 1 m $s^{-1}$. The averages for each position were taken over the last 50 ps at each point. Free energies were calculated using a weighted histogram analysis method (WHAM) [176].

We attribute the difference in the work required to exfoliate the flake to the friction that must be overcome when sliding the flake over the substrate. As the flake is sheared it can fall into different commensurate positions; each time it does so the spring must overcome this barrier to move the flake. Graphene is a flexible 2D material; bending the flake is associated with a lower energy barrier. I had not attempted to validate the bending energy of a graphene flake as described by OPLS, so these results are not definitive. By peeling the flake, the end still attached to the substrate does not need to leave a commensurate position. When designing exfoliation processes, including surfactants and intercalation agents for graphene liquid phase exfoliation, this peeling mechanism should be targeted.

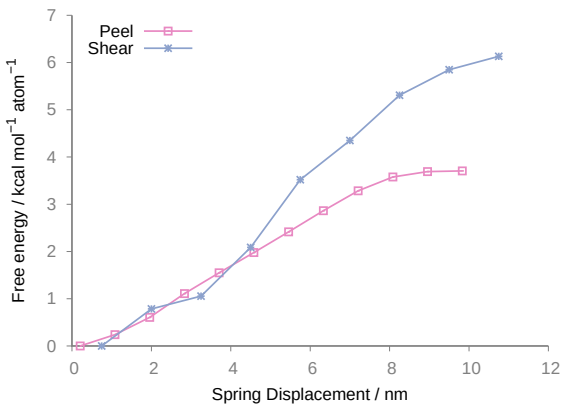


**Figure 4.13:** Free energies per atom of exfoliating a graphene flake from a graphitic substrate via different mechanisms.

Figure 4.13 shows the results of a weighted histogram analysis, comparing the free energies associated with peeling and shearing mechanisms of exfoliation. The

spring energy was divided by 2520 to make the quantity independent of size. The tolerance was 0.00001 and number of Monte Carlo trials was 20 (these are program specific parameters to reference [176]).

A comparative steered simulation was carried out where the spring was moved 5 Å every 100 ps, resulting in an effected pulling speed of 5 m $s^{-1}$, instead of 1 m $s^{-1}$. The peeling mechanism required only 12 % less work. The reduced difference in work done is due to the sheared flake having less time to fall into commensurate positions along its exfoliation pathway; less friction is therefore observed between the flake and substrate. Nonetheless, this does show that the peeling mechanism remains favoured at different pulling speeds.

## 4.3 Summary

In summary, GraFF has been shown to recreate several properties of graphene friction seen experimentally. The simulations have revealed the unusual sliding mechanisms and temperature dependence of graphene in a superlubric state. It has been shown that a full description of friction in graphene systems predicts new behaviours and that exfoliating graphite via peeling is the energetically favoured mechanism. In the next chapter, I investigate the exfoliation of graphene further, by simulating its exfoliation with an adhesive polymer explicitly.

## Chapter 5

# Micromechanical Exfoliation of Graphene

Having shown that GraFF offers useful insight into graphene-graphene interactions in the previous chapter, and having verified that it can model the adsorption and friction between sheets, I wanted to investigate a more realistic and complex exfoliation environment. The last chapter dealt with graphene in vacuum, exfoliation was induced by attaching a 'virtual' spring to a corner of the graphene flake. Here, the exfoliation will be facilitated by an adhesive polymer, mimicking the famous sellotape exfoliation that first produced graphene [11]. The work in this following chapter was first presented in Ref.[3]; the simulations and analysis was carried out by myself under the supervision of James Suter and Peter Coveney.

As has been discussed throughout this thesis, methods of reliably synthesising pristine graphene are scant within academia and industry and urgently need improving [18, 177, 178]. Micromechanical cleavage of graphite, using sticky tape to prise apart graphene layers, is still widely used in academic research and in industrial environments. The adhesive substrate provides a means to directly apply force to individual graphene layers to prise them apart, which is suited to produce very high quality large graphene sheets [179]. It is still the primary way of manipulating graphene sheets for probing its mechanical [14] and electrical [180] properties. However, it is an imprecise method and much effort is expended on trivial tasks like producing graphene as opposed to graphite, locating suitable graphene flakes on the

exfoliation tape, removing the substrate material, and cleaning off debris. There is clearly scope for improvement, I decided to see if we can learn anything about this process now we can resolve graphene's frictional properties using GraFF.

Other mechanical exfoliation techniques have shown promise as scalable methods for producing graphene [179, 181]. It is commonly understood [177] that mechanical exfoliation is propagated either normal to the graphene plane, known in engineering terms as a mode I fracture and achieved via sonication [177], or lateral to the plane in a shearing motion, known as mode II fracture and realised in processes such as ball-milling [177] or pressure driven fluid dynamics [182] (see figure 5.3a). To bring them to fruition it is important to improve the primitive understanding we have of the mechanical exfoliation mechanisms. Indeed it is remarkable how little attention has been paid to comprehending this key process.

Liquid phase exfoliation of graphene has been studied with MD before [183, 184] but without paying attention to graphene's unusual intermolecular forces [1, 185, 168, 167]. However, the studies [183, 184] did not allow graphene to move, simulating essentially static sheets and moving them manually to study the exfoliation.

## 5.1 Simulation Methods

The procedure used for simulating the exfoliation of graphene by polymer tapes is shown in figure 5.1. A graphite stack, comprising 7 graphene flakes, was compressed under NPT conditions between two 4 nm thick polymer layers. This configuration was then pulled apart in the NVT ensemble by increasing the box height at a constant speed to exfoliate the graphite. This setup is simple to perform experimentally, but has to our knowledge not been investigated at an atomic scale before.

The polymer 'sticky tape' was modelled using polymethylmethacrylate (PMMA) and polydimethylsiloxane (PDMS). The precise composition of commercial sticky tape brands is not publicly known but they certainly contain acrylate polymers, which have been used as pressure sensitive adhesives for many decades. PDMS is a common polymer used for graphene transfer [186]. To make the poly-

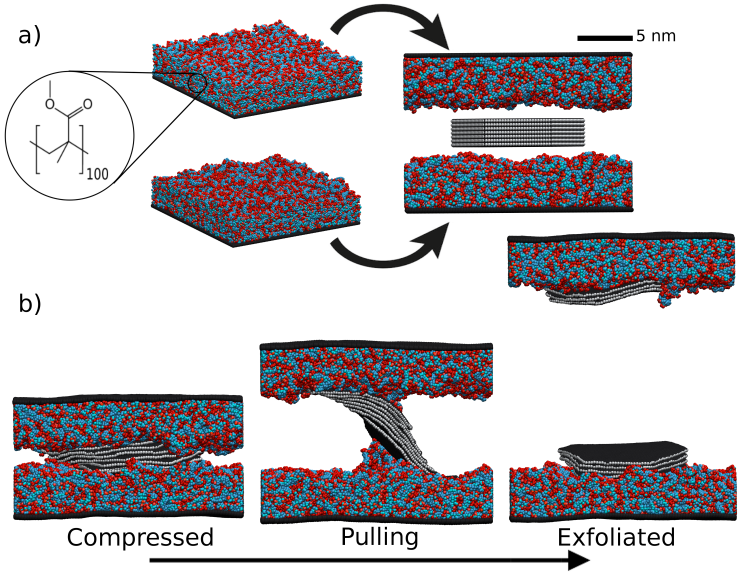


**Figure 5.1:** a) Two slabs of 88 poly(100)methylmethacrylate molecules were equilibrated on a rigid surface. The sticky tape exfoliation simulations are built by placing a stack of 7, 12 nm wide hydrogen terminated graphene flakes between two polymer tapes. Blue - polymer backbone; red - polymer branch; grey - graphite carbon; white - hydrogen edges. b) The simulation is compressed and equilibrated under NPT conditions, then pulled apart by increasing the vertical box dimension constantly at 10 $ms^{-1}$ until exfoliation occurs.

mers behave like a tape we build it on top of a continuous lattice extending through the *xy* plane; this represents the hard elastic polymer to which the adhesive polymer is attached.

The polymers were built using an in house algorithm written by James Suter [129]. This algorithm grows polymer chains systematically, one monomer unit at a time, using a distribution of torsion angles between monomer units to build realistic randomly orientated conformers. A benefit of this algorithm is that it can build polymers around other objects in the simulation (in this case the continuous lattice at one edge of the simulation), which many commercial polymer builders cannot do. Given a desired number of polymer chains, the end groups are scattered randomly into the simulation box, avoiding any overlap. Each polymer is extended by one monomer, the position of which is drawn from a known distribution of torsional angles from a previously simulated polymer chain. Each addition may be tried several times to find a configuration that does not have an energy over a certain threshold. A polymer chain may become 'trapped', where the chain cannot be extended further,

the algorithm can then delete some recently added monomer units to find a new configuration. It is hoped this algorithm created entangled and realistic polymer configurations.

The polymers were built on top of the substrate at a 0.5 relative density because PMMA has bulky side groups and the algorithm needs the extra volume to propagate the polymer chains and form entanglements. Polymers were made of 100 monomer units, which is probably below the molecular weight of commercial sticky tape but is above the entanglement length [86] so that adhesive properties are hopefully retained.

First, the polymer tape's energy was minimised in LAMMPS, then compressed along the $z$ axis over 200 ps to a density of 1. The simulation box was then increased by 25 Å perpendicular to the tape and allowed to equilibrate at 300 K for 2 ns. The compression was necessary to remove voids within the polymer within a practical timeframe. Expanding the simulation box allowed the polymer to relax after compression and form a naturally undulating surface, giving structures like to those shown in figure 5.1a.

Two different equilibrated polymer tapes were then rearranged to face each other and a graphite slab placed in between them with 10 Å initial spacing between polymer and graphite. Graphite starting coordinates were built using an in house script [8]. The polymer velocities are retained from the equilibration. The tape walls and PMMA/PDMS were then simulated together for 200 ps using NVT. Next, the sticky tape sections were simulated using NPT and allowed to move together in the $z$ direction; concurrently, the graphite was simulated using NVT — this was done because we found it more stable when introducing the graphite into a system which had already been equilibrated at 300 K. Finally, the whole system was simulated with the same integrator under atmospheric pressure for 500 ps, The $xy$ dimensions were kept constant and the barostat equalised pressure by changing coordinates along the $z$ axis.

Introducing graphene between the two polymer layers sometimes gave rise to undesired systems. When the box was compressed under NPT conditions the

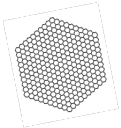

polymer was often found to 'topple' over the graphite stack. I attribute this to the flakes being small and the rather unphysical way these simulations are setup: those stacks that toppled over were excluded from the ensembles as they invariably looped across the periodic boundaries. This was done systematically: if two sheets that started off adjacent to each other drifted apart so that there was less than 10 kcal $mol^{-1}$ dispersion interaction between them the stack was deemed to have toppled over and was discarded.

For the exfoliation simulation step the boxes were then strained in the $z$ direction — perpendicular to the graphite layers — using a constant strain velocity of 10 m $s^{-1}$, within the speed at which one can pull sticky tape. When the simulation had zero stress in the $z$ direction, over an average of 50 ps, i.e. a break had occurred, it was allowed to run for a further 200 ps before stopping.

A summary of all the simulations performed is given in table 5.1. The first simulations undertaken used graphene 6 nm wide graphene flakes and polymer tapes that were 10 x 10 x 4 $nm^3$. In table 5.1 these are labelled 1 and 2, comparing different strain rates. Simulation 2 was used to estimate the variance in the expected exfoliation, discussed below. From simulations 1 and 2 it was concluded that using strain rates of 1 or 10 $ms^{-1}$ did not have a large effect on the outcome, therefore the following simulations used a strain rate of 10 $ms^{-1}$ for computational efficiency. To investigate finite size effects in simulation 3 we doubled the lateral dimensions of the polymer tape; and in simulation 4 we doubled the size of the polymer tape and doubled the graphene flake width to 12 nm.

Simulations 4 and 5 have the largest flakes and polymer box sizes, and compare poly-methyl-methacrylate with poly-dimethyl-siloxane. These two simulations form the basis for most of the discussion in the next section. The other simulation details are listed in table 5.1 for comparison.

Graphite and graphene non-bonded interactions were handled using the GraFF forcefield described in the previous chapters [1, 10]. All other interactions, in the polymer and polymer-graphene interactions were modelled using the OPLS forcefield [140, 187, 188, 189]. It should be noted that, using a typical forcefield such

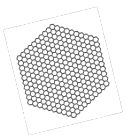

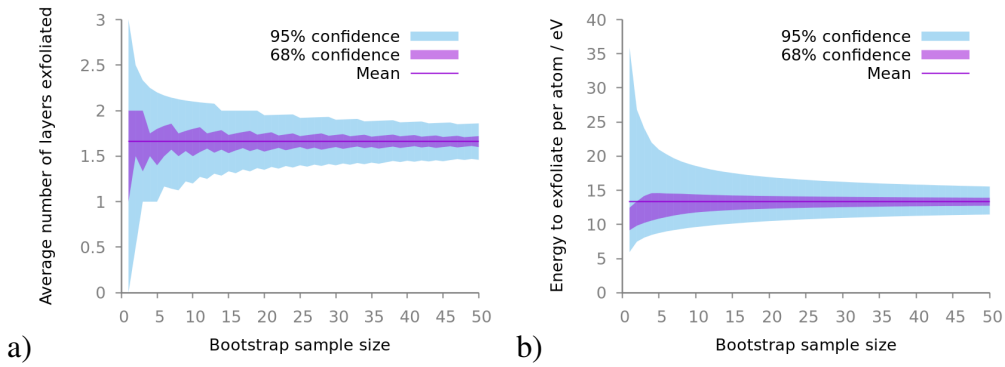


**Figure 5.2:** A bootstrap with replacement study on simulation 2 (see table 5.1) with varying the resample size. Our simulations with larger boxes used around 15 replicas per ensemble. Here, one can see 15 is a reasonable point at which there are diminishing returns concerning the confidence in the average a) number of layers exfoliated and b) energy required to exfoliate the stack.

as OPLS to parameterise these simulations, exfoliation has never previously been reported: the graphite stays attached to one polymer slab, and a gap always opens up between the other polymer layer and the graphite.

### 5.1.1 Ensemble Size

As seen previously these steered simulations have chaotic dynamics and are thus extremely sensitive to their initial conditions [190]. For all systems we therefore simulated ensembles to acquire reproducible results. Each replica in the ensemble was built separately: the polymer slabs vary in their spacial arrangement not just their initial atomic velocities, to ensure a range of configurations.

A large ensemble of 50 replicas was studied to characterise the system's global behaviour (simulation 2 in table 5.1). Each replica started with unique, uncorrelated atomic positions and velocities; the velocities were drawn from a Maxwell–Boltzmann distribution. The average number of flakes exfoliated and the energy required per atom to cause the exfoliation were used as the characteristics by means of which to quantify a particular replica. A bootstrap with replacement using 10,000 samples was performed on this sample to quantify the confidence in the results that were derived, shown in figure 5.2.

### 5.1.2 Thermostat

The 'pulling' phase of this simulation (figure 5.1) is carried out in the NVT ensemble. By changing the box size we are injecting a lot of energy into the system. Originally I wanted to conduct the simulation in the NVE as in the superlubricity experiments described in the previous chapter, but if we didn't couple the atomic velocities to a thermostat the temperature would increase drastically. I tried coupling the polymer to a thermostat, so that the graphene flakes would be simulated using NVE and not artificially lose energy, however the middle flakes are often not in contact with the polymer and the flakes gained too much energy and broke apart due to their high temperature.

A possible simulation setup that could be studied in future would be to include a carrier gas that would act as a physical thermostat (or heat sink) for the graphite [191, 192]. Page *et al.* suggest this is a way to preserve dynamics while keeping good temperature control in non-equilibrium MD simulations [191]. It would require a Monte Carlo step to add/remove gas particles as the simulation box changed size.

## 5.2 Results and Discussion

As mentioned above, mechanical exfoliation usually proceeds by either a normal or shear mechanism (see figure 5.3a). However, we find that micromechanical cleavage with PMMA exhibits a mixture of the two modes which we refer to as a peeling mechanism (see figure 5.3b). The adhesive polymer is a flexible platform, this provides a unique setup which allows the exfoliation force to be applied perpendicular to the graphite and allows the graphite to reorient to increase the motion's shearing component. Graphene exhibits very little friction between adjacent layers [161] and an extremely low bending energy [193] which make it unique amongst laminar materials. Friction and bending are so easily overcome that when the force is applied in the normal direction the sheets are able to bend and slide in a peeling motion. This mechanism reduces the steep potential energy gradient that purely normal mode exfoliation would otherwise engender.

| Simulation Label | 1 | 2 | 3 | 4 | 5 | 6 |
|---|---|---|---|---|---|---|
| Forcefield | OPLS+GraFF | | | | | OPLS |
| Pull speed / $ms^{-1}$ | 1 | 10 | 10 | 10 | 10 | 10 |
| Box size / nm | 100 | 100 | 200 | 200 | 200 | 200 |
| Flake diameter / nm | 60 | 60 | 60 | 120 | 120 | 120 |
| Polymer | PMMA | PMMA | PMMA | PMMA | PDMS | PMMA |
| Energy to remove per atom / meV | 12.0 | 13.3 | 7.46 | 13.7 | 13.2 | n/a |
| # Replicas in ensemble | 5 | 50 | 10 | 15 | 14 | 10 |
| Average layers exfoliated | 0.8 | 1.66 | 2.0 | 0.87 | 2.4 | 0 |
| # No exfoliation occurred | 1 | 4 | 1 | 5 | 1 | 10 |
| # 1 - Graphene exfoliated | 3 | 17 | 3 | 8 | 1 | 0 |
| # 2 - bilayer graphene exfoliated | 1 | 21 | 1 | 1 | 3 | 0 |
| # 3 | 0 | 8 | 5 | 1 | 9 | 0 |

**Table 5.1:** A summary of all the simulations we undertook. The simulations are chaotic (i.e. they exhibit extreme sensitivity to initial conditions), so ensembles are required to understand and compare systems. Simulations 1-4 compare finite size and pulling speed effects. These results are reasonably similar, the same mechanism being observed across all systems. Simulations 4 and 5 use the largest flakes and largest polymer tapes and therefore are the most useful for experimental intuition. "OPLS+GraFF" implies that polymer interactions are described by the OPLS forcefield, while graphene and graphite were modelled using GraFF. Simulation 6 compares simulation 4 to one with a more commonly used forcefield; OPLS employed alone is inadequate as no exfoliation is observed.

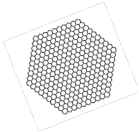

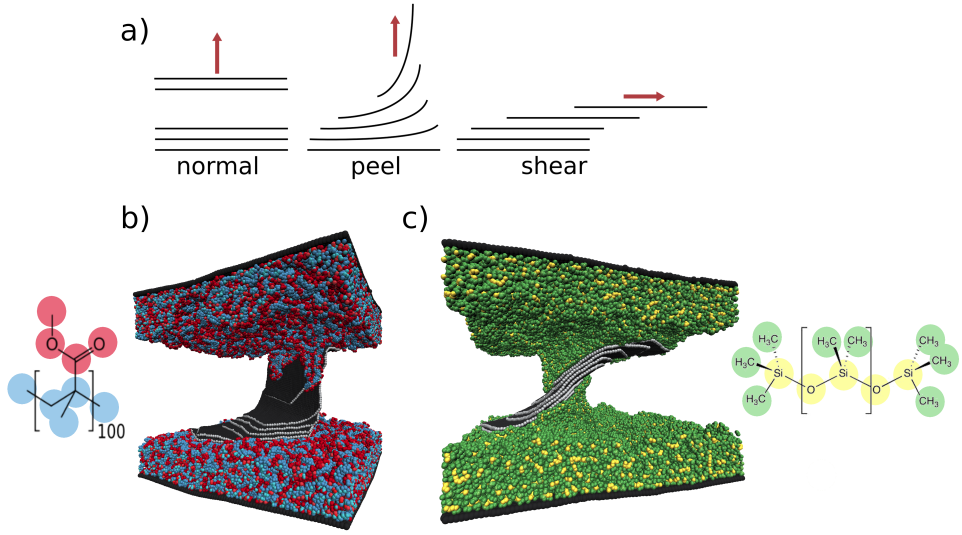


**Figure 5.3:** a) Schematic showing different modes of exfoliating laminar materials. Graphite is shown being exfoliated by PMMA in b) and PDMS in c). PDMS is more fluid and rearranges to stay in contact with the graphite, encouraging more shear component of the exfoliation. PMMA peels graphene layers from the top of the stack, whereas PDMS shears the layers apart, this goes some way to explaining the difference in exfoliation seen in figure 5.4

To study the chemical specificity of this process we also simulated exfoliation using PDMS as the adhesive. PDMS is a popular alternative to commercial sticky-tape adhesives for exfoliating graphene [186]. In our setup PDMS is less viscous than PMMA, allowing for more rearrangement around the graphite during the compression time. This is because we are using polymer chains much shorter (for both PMMA and PDMS) than would be used practically so a lot of their properties will not be comparable to experiment. It is also common to add cross-links in PDMS to increase its rigidity and reduce the debris left behind on the graphene, I did not attempt this in my simulations.

Being less viscous, the PDMS can also mold around the exfoliating graphite as it starts to peel and shear (see figure 5.3c). As the stack begins to shear we observe polymer chains attach to several sheets, making them less likely to exfoliate to a single layer.

The difference in the exfoliation mechanism can be quantified by investigating the dispersion forces between the graphite stack and the two different polymer adhesive tapes. In the simulations, PDMS adhesive tapes exhibit stronger disper-

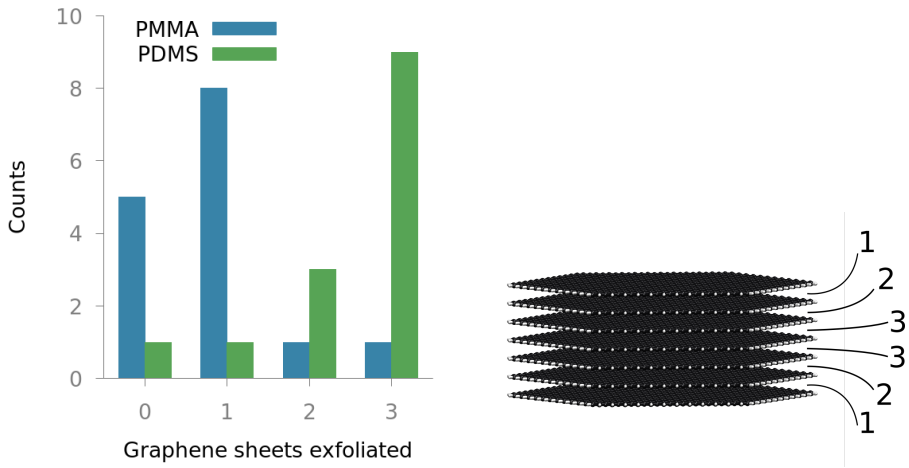


**Figure 5.4:** Frequency of different types of exfoliation by different polymers from molecular dynamics simulations. The histogram is made from ensembles of 15 replicas of PMMA and 14 replicas of PDMS. Graphene sheets exfoliated indicates the number of layers in the smallest stack after the original 7 layer graphite is exfoliated. 0 implies no exfoliation occurred.

sion interactions with graphite than PMMA. The total interaction energy between an outer 12 nm wide flake of graphene and polymer just before pulling is -77 $\pm$ 2 eV for PDMS and -59 $\pm$ 6 eV for PMMA. However, the interaction between a layer of polymer and a static, quasi-infinite graphene layer shows the opposite trend; the interaction energy between polymer and graphene is -0.838 $\pm$ 0.004 eV nm$^{-2}$ for PDMS and -0.880 $\pm$ 0.008 eV nm$^{-2}$ for PMMA. This discrepancy is due to the ease with which PDMS rearranges around the graphite stack during the compression stage and increases the amount of polymer in contact with the graphene.

Values for the interaction energy between a graphene flake and an adjacent polymer layer before the simulation starts pulling are the average and standard deviation of simulations 4 and 5, referred to in table 5.1. The interaction between a polymer layer and a quasi-infinite graphene layer that extended through periodic boundaries is taken from 5 different polymer layers equilibrated on a static graphene sheet. The polymer layer was 4 nm thick in a 20 x 20 x 6 nm$^3$ simulation box. The errors are the standard deviation of the distributions obtained.

The increased contact with graphite discourages the formation of graphene; instead PDMS favours breaking the graphite in the middle of the stack. From the distribution described in figure 5.4, the average number of layers exfoliated in the

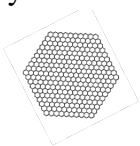

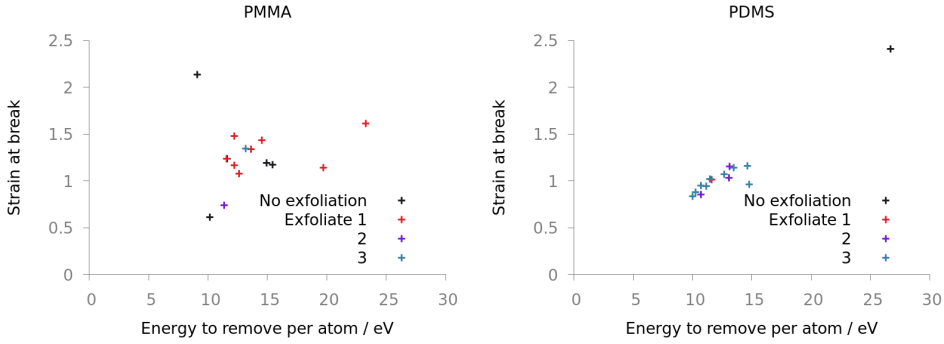


**(a)** **(b)**

**Figure 5.5:** Comparing the strain required to fracture simulations 4,5 (see table 5.1). PDMS produces a tight distribution compared to PMMA, the one exception being a simulation where no graphene was exfoliated.

simulations is $0.7 \pm 0.3$ for PMMA and $2.4 \pm 0.8$ for PDMS; clearly we are observing two distinct mechanisms (errors are the 95% confidence interval of the average on each distribution in figure 5.4).

The biggest difference between the polymers we investigate is their viscosity. This is a dynamic property, therefore, the timescales used could change the exfoliation outcome. Using a shorter dwell time could reduce the amount of polymer that can mold to the graphite, promoting graphene exfoliation. In a similar way, using faster shearing speeds could reduce polymer rearrangement around the graphite, again promoting graphene production. In general, higher viscosity polymers should be invoked to promote graphene exfoliation. The use of cross-linked PDMS should further improve performance [186].

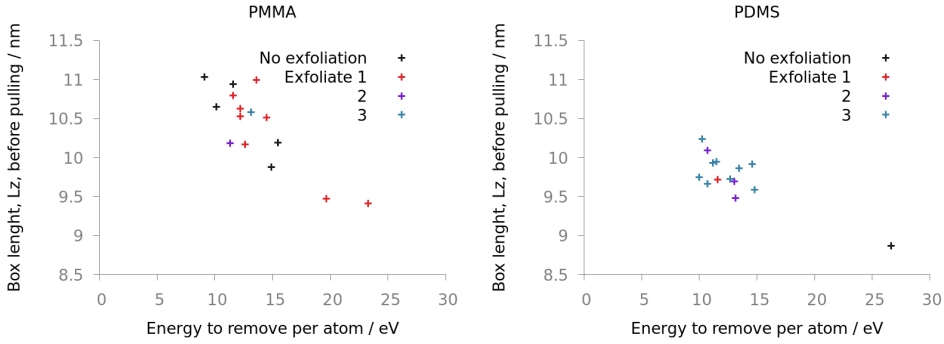


**(a)** **(b)**

**Figure 5.7:** Comparing the height of the simulation box before pulling of simulations 4,5 (see table 5.1). As discussed in the main text, PDMS is more fluid and therefore the polymer can deform around the graphite, the deformation makes the simulation boxes shorter after the compression stage.

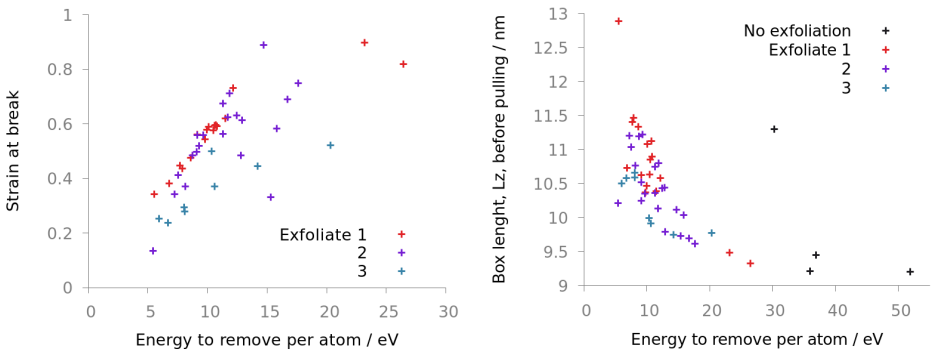


**(a)** **(b)**

**Figure 5.6:** A break down of the exfoliation replicas taken from simulation 2 (see table 5.1). a) shows that exfoliation of graphene, red crosses, (apart from two exceptions) follows a smooth trend, due to the peeling mechanism discussed in the main text. Exfoliation of 3 sheets follows a trend with a lower gradient, because more energy is required to bend 3 sheets than 1. b) shows the height of the box before the pulling simulation started; the distribution is quite broad and shows minimal correlation between compression and exfoliation outcome. Replicas that saw no exfoliation have a high stress because they often pulled polymers from the opposite tape with them, which requires a lot of rearrangement energy.

Some of the anomalies shown in figures 5.6 and 5.5 are due to polymer chains bridging between tapes and requiring a much larger total energy to break apart as the polymer chains are dragged with the increasing box size. A snapshot of such a simulation is shown in figure 5.8

We can now begin to understand the different ways in which graphene can ex-

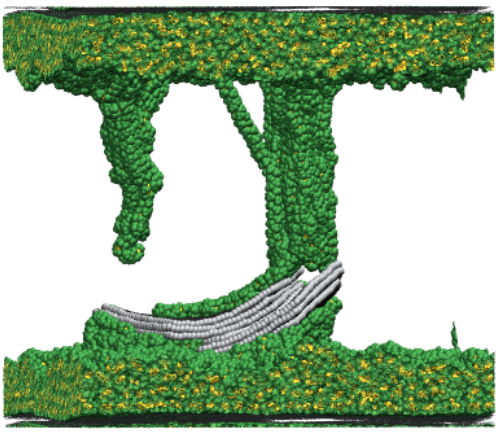

**Figure 5.8:** Example of a polymer chains bridging the simulation box, increasing the energy required to break apart the system.

foliate. Normal mode fracture is seen on the macroscale, where the friction between large sheets does not allow them to slide past each other; on the nanoscale, shearing is observed with fluid polymers like PDMS; and as the substrate become less fluid a peeling mechanism is seen, which favours synthesis of graphene and relies on low friction between layers. The peeling mechanism we observe with PMMA is desirable because it more reliably produces graphene. The component of peeling will be proportional to the flake's surface area as this mechanism requires low friction between layers. If the sheets can easily slide past each other, only a small number of sheets will exfoliate with the polymer as fewer sheets have lower bending energy.

Indeed, we know that the bending stiffness of graphite increases exponentially with the number of layers [194]. The peeling mechanism we have identified requires the exfoliating layers to bend; it therefore becomes exponentially harder to 'pick up' more graphene sheets. This reduces the number of cleavages required to reach a single graphene layer.

It is widely perceived that the reason graphene readily exfoliates with sticky tape is because the interlayer forces between graphene sheets are far weaker than those of graphene to polymer. This can be satisfactorily explained by the high dispersion forces found in materials like PMMA which do indeed contribute more non-bonded energy per unit-area than graphene. However, a more nuanced under-

standing is needed to explain why graphite is so stubborn in resisting solution phase exfoliation in similar polymers and solvents since, by similar arguments, exfoliation should be enthalpically favoured. In a liquid phase it may not be possible to generate the necessary shear forces to slide large adjacent sheets past each other, and the entropic penalties associated with exposing the solvent to more surface may be too high. Of course, exfoliation of other laminar materials in polymer is often observed [157], so a complete understanding still does not exist.

The systems simulated here have been designed to investigate the original experiment by which graphene was first discovered by Geim et al. [11]. Using MD gives us atomistic insight into the system but is limited in the time and length scales that can be measured. We simulate in vacuum graphene flakes that are 12 nm in diameter, at the lower end of what is found experimentally. These nano-flakes are still of interest, for example when demonstrating superlubricity of graphene [161]. Simulating sheets that are orders of magnitude larger would require novel multiscale modelling methods [129]. Moreover, real graphene particles are not so regularly shaped, and these aspects will also affect the behaviour.

## 5.3 Monte Carlo Model

Using a mathematical model we can extrapolate our results to show how the mechanism of exfoliation affects the experimental outcome of repeatedly exfoliating a graphite stack. Starting with a graphite stack of $N$ layers, (assuming graphite flakes are on the order of 100 $\mu$m in size, $N$ is then approximately 30,000); by repeatedly exfoliating from the top, we want to know what is left behind. If exfoliation of graphite caused the stack to fracture in the middle, it would take $c = \lceil \log_2(N) \rceil$ cuts to guarantee production of graphene (upper square brackets denote the ceiling function, rounding the number of breaks up to the nearest integer larger than the bracketed value).

However, it is reasonable to assume that the graphite breaks in a more stochastic process. Next, we assume that when the graphite stack is cut the break is equally likely to occur between any two graphene layers. We treat the problem as a Markov

process: there are $N$ possible states, one corresponding to every possible number of graphene layers in the stack. State 1 (N=1, graphene) is the absorbing state (also called the terminating state), all other states are transient. We therefore have $n$ transient states, $n = N - 1$. The probability at each iteration to transition from state $i$ to $j$ is:

$$b_{ij} = \begin{cases} 1/(i-1) & \text{, if } i > j; \\ 1 & , i = j = 1; \\ 0 & \text{, otherwise.} \end{cases} \tag{5.1}$$

We construct the transition probability matrix $\mathbf{P}$, an $N \times N$ matrix, from the elements $b_{ij}$, as follows:

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1/2 & 1/2 & 0 & \dots & 0 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 1/(N-2) & 1/(N-2) & 1/(N-2) & \dots & 1/(N-2) & 0 & 0 \\ 1/(N-1) & 1/(N-1) & 1/(N-1) & \dots & 1/(N-1) & 1/(N-1) & 0 \end{bmatrix} \tag{5.2}$$

The column and row number of $P$ correspond to the number of layers of graphene before and after exfoliation respectively. One can see that all states except for pure graphene will produce a graphite stack with fewer layers.

A general absorbing markov problem, with one absorbing state as we have here can be written as:

$$\mathbf{P} = \begin{bmatrix} 1 & \mathbf{0} \\ \mathbf{b} & \mathbf{B} \end{bmatrix} \tag{5.3}$$

where **B** is the matrix grouping the transition probabilities among the transient states, in this case it is the $n \times n$ matrix that makes up the bottom right portion of matrix **P** written above. **b** is the $n$-dimensional column vector grouping the probabilities from any state to the absorbing one, in this case it is the first column of **P** except the first entry which is the absorbing state. Finally, **0** is the zero vector.

$$\mathbf{b} = [1, 1/2, 1/3, ..., 1/(N-2), 1/(N-1)]^{\mathrm{T}}. \tag{5.4}$$

**B** represents the transition probabilities for all transient states (i.e. states 2 to N). Note that as **P** is a transition probability matrix each row must sum to 1.

The initial state of the system is a stack of $N$ sheets, i.e. the system has 100% chance of being in state N, the initial distribution can be described by a vector of length $n$:

$$\mathbf{a} = [0, ..., 0, 1]. \tag{5.5}$$

Let $\tau$ be the number of cuts, $c$, that have occurred when graphene is created (absorption into state 1). $\Pr(\tau = c)$ is then the probability that it took $c$ cuts to produce graphene. The PDF is then:

$$g(c) = \Pr(\tau = c) = \mathbf{a}\mathbf{B}^{c-1}\mathbf{b}. \tag{5.6}$$

This provides a distribution describing the expected number of exfoliations from a graphite stack to produce graphene; see figure 5.9. Our assumption was that the stack is equally likely to break between any two sheets; using our simulation results we theorise graphite stack is more likely to break near the outer sheets using certain polymers. This model can account for such a scenario by modifying Eqn.1, for example:

$$b_{ij} = \begin{cases} \pi(j-1, i-1, a) & \text{, if } i > j; \\ 1 & , i = j = 1; \\ 0 & \text{, otherwise.} \end{cases} \tag{5.7}$$

where $\pi(k,t,a)$ is the symmetric beta-binomial distribution, a distribution with discrete finite support that favours the lowest and highest values; $t$ is the number of trails (number of possible states after exfoliation), $k$ the new state, and $a$ is the shape parameter $0 < a < 1$.

A more intuitive understanding of the processes involved can perhaps be better achieved by treating the problem as a continuous one. Considering that our graphite stack has height $h$, after $c$ cuts it will have height $h_c$. If by exfoliating from the stack it is equally likely to break anywhere along its height, we have:

$$h_1 = h_0 X_1; \tag{5.8}$$

$$h_c = h_{c-1} X_c; \tag{5.9}$$

$$= h_0 \cdot \prod_{i=1}^{c} X_i = h_0 X; \tag{5.10}$$

where $X_i$ is a uniformly distributed random variable. The PDF of $X = \prod_{i=1}^{c} X_i$ is given by:

$$f_X(x) = \begin{cases} \frac{(-1)^{c-1} \log^{c-1} x}{(c-1)!} & , 0 < x < 1; \\ 0 & , \text{otherwise.} \end{cases} \tag{5.11}$$

N.B. $\log^a x \equiv (\log x)^a$. For a proof of the above see Dettmann *et al.* [195].

If $h_g$ is the height of a single graphene sheet, this continuous approximation is valid whilst $h \gg h_g$. We can find the chance that graphene has been produced after $c$ cuts by finding the probability $h_c \leq h_g$. The probabilities $\Pr(0 < h_c \leq h_g)$ and

$\Pr(0 < f_X(x) \leq h_g/h_0)$ are equivalent; the following integral [196] will therefore give us the probability of graphene being produced from a graphite stack after $c$ exfoliations:

$$G(c) = \int_0^{\frac{h_g}{h_0}} f_X(x)dx = \frac{(-1)^{c-1}}{(c-1)!} \cdot \int_0^{\frac{h_g}{h_0}} \log^{c-1} x\, dx = \frac{(-1)^{c-1}}{(c-1)!} \cdot S_{c-1}\left(\frac{h_g}{h_0}\right); \tag{5.12}$$

$$S_n(x) = x \sum_{k=0}^{n} (-1)^{n-k} \frac{n!}{k!} \log^k x. \tag{5.13}$$

Solutions to the above for $h_0 = 100\ \mu$m, $h_g = 3.35$ Å are shown in figure 5.9. Integrating a function like $\log^n x$ from 0 is made possible with the following steps. First we make a substitution of $x = e^t$:

$$I_n = \int_0^1 \log^n x\, dx = \int_{-\infty}^0 t^n e^t\, dt. \tag{5.14}$$

This allows us to use the formula for integration by parts: $\int u\, dv = uv - \int v\, du$. By making appropriate substitutions we get:

$$u = t^n,\ du = nt^{n-1}, \tag{5.15}$$

$$dv = e^t dt,\ v = e^t, \tag{5.16}$$

$$I_n = t^n e^t - \int nt^{n-1} e^t\, dt, \tag{5.17}$$

$$I_n = t^n e^t - nI_{n-1}. \tag{5.18}$$

Continue in this way until no more iterations are neccesary, this arrives at the solution in equation 5.13.

The above describes repeated exfoliation from a single graphite stack. So far we have considered only cleaving sequentially from one stack; it is straightforward experimentally to exfoliate in a parallel fashion by shifting the sticky tape at each iteration. By this method one can repeatedly cleave every stack that is created.

We find the probability of obtaining graphene by this method by considering the probability a single stack is not graphene: $H(c) = 1 - G(c)$. Cleaving in a parallel fashion there will be a maximum of $2^c$ stacks, each of which could be graphene; so the probability that graphene exists after $c$ cleavages is:

$$G^{\text{parallel}}(c) = 1 - H(c)^{2^c} \tag{5.19}$$

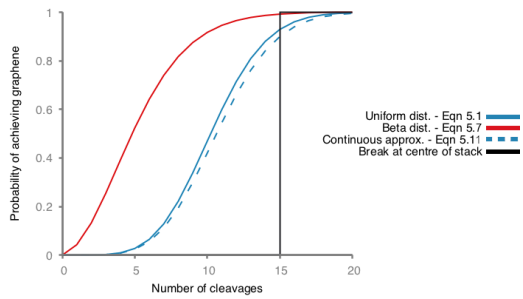


**Figure 5.9:** Probability of synthesising graphene from a 30,000 layer graphite stack after some number of exfoliation steps using different numerical and analytical methods. Beta binomial dist. $\alpha = 0.12$.

Figure 5.9 demonstrates how the way in which one exfoliates graphite can have a large impact on the effort required to synthesis graphene. Using the peeling mechanism which promotes exfoliation of smaller stacks can drastically reduce the number of cleavages required. However the shape parameter for the beta-bionmial distribution is not known accurately and is only representative, without detailed experimental research we cannot with confidence predict this value, or assume it is independent from $N$.

Fitting to the PMMA distribution shown in figure 5.4 gives $\alpha = 0.12$, a fit is shown in figure 5.10a. Extrapolating that data to 30,000 sheets is shown in figure 5.10b. Using this distribution produces the red line in figure 5.9. The expected number of exfoliations to produce graphene from a 100$\mu$m graphite stack is 11 for methods that break the stack at a random point but would be much lower if the peeling mechanism is employed that encourages exfoliation at the edges (when $\alpha = 0.12$, the expected number of cuts is 5 to produce graphene).

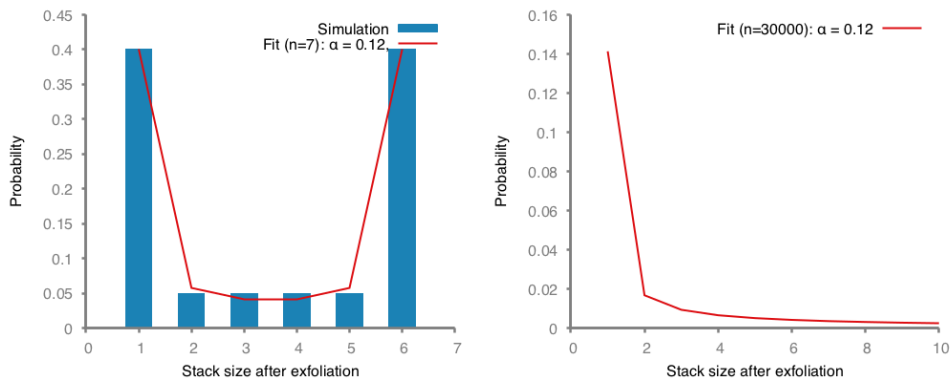


**(a)** **(b)**

**Figure 5.10:** a) Given that exfoliation occurs, fitting the betabinomal distribution to the PMMA mechanism of exfoliation gives a shape parameter of $\alpha = 0.12$. b) Extrapolating to a realistic stack of graphite, discussed in the main text, using the same shape parameter.

## 5.4 Conclusions

In this study we have, for the first time, provided atomistic insight into the mechanisms of mechanical exfoliation of graphene. To observe this phenomenon careful consideration must be given to the graphene intermolecular interactions as standard forcefields are not able to simulate the behaviour with sufficient accuracy; here we have used the recently developed forcefield GraFF [1, 10]. I have also presented a transferable mathematical model for describing different modes of graphene exfoliation. Graphene's low bending energy and low friction between sheets allows graphite to shear even when the force is applied normal to the graphene plane, facilitating the production of graphene via a peeling mechanism. To promote graphene production, experimentalists should use rigid or viscous substrates. Some variants of mechanical exfoliation are strong candidates for large scale production of graphene; however, understanding the fundamental mechanisms behind it is essential if reliable large scale manufacturing methods are to be found.

Extrapolating these simulations to much larger graphite flakes is only meant to be representative. MD is limited in its ability to simulate much larger systems, and I think the most that could be extracted from this study using only MD was achieved. In future, comparing the findings and theoretical model to experimental measurements, of the mechanism or probabilities of different number of layers af-

ter exfoliation, would be very interesting. I would like to see how significant the finite size effects would be, and test the hypothesis that the likelihood of exfoliating graphene could be controlled by the substrate's viscosity.

## Chapter 6

# The Structure of Graphene Oxide

One of the biggest challenges to exploiting graphene's desirable properties is synthesising it in large quantities. To circumvent this roadblock, many scientists replace graphene with graphene oxide (GO) in their materials. Because of its prevalence, especially in nanocomposite research, it was important for our investigations to be able to build accurate models of GO. The research in this chapter was originally motivated by wanting to build such a model, it also led to interesting analysis of the two-phase graphene oxidation reaction. The results were first published in Ref [2], the research and analysis was carried out by myself, Peter Coveney assisted in writing the manuscript.

Graphene oxide is an amorphous 2D material which has found widespread use in the fields of chemistry, physics and materials science due to its similarity to graphene with the benefit of being far easier to synthesise and process. However, the standard of GO characterisation is very poor because its structure is irregular, its sensitivity to the preparation method, and its propensity to transform due to its reactive nature. Atomistic simulations of GO are common but the nano-structure in these simulations is often based on little evidence or thought. There is no precise consensus about the nanostructure of GO [197]. The Lerf-Klinowski model [55] (figure 6.1a) is widely recognised and has formed the basis of much scientific research [198, 199]. This model assumes an uncorrelated random distribution of epoxy and alcohol groups on the surfaces, with alcohol and carboxyl groups around the edges. However, correlation between oxidised sites seems chemically intuitive:

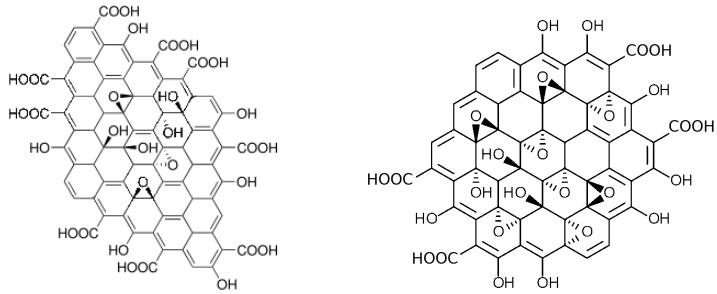


(a) (b)

**Figure 6.1:** a) Lerf-Klinowski model of the structure of GO [55]. Their work established that the major functional groups are carboxylic acids and alcohol groups on the edges, and epoxy and alcohol groups on the surface. This basic pattern has been confirmed by many experiments [55]. b) A circumcoronene molecule oxidised using our algorithm. This model relies on the same assumptions as to which functional groups exist, but the groups are added sequentially based on the relative reactivity of unoxidised sites.

isolated carbon double bonds are more reactive than conjugated/aromatic systems [200]; indeed, several experiments have shown the presence of oxidised and unoxidised regions [201, 202, 203, 204, 59, 205]. A comprehensive understanding of how this pattern could evolve does not exist.

Yang *et al.* [58] enlightened this discussion by studying the various reactive intermediate structures that could occur in graphene oxidation, using quantum mechanical calculations. They predict that oxidation is so overwhelmingly favoured adjacent to already oxidised carbons that separate large oxidised and aromatic regions are inevitable. The behaviour of the material will clearly depend on the distribution and morphology of these regions. Until now, however, simulations aimed at understanding the nanoscale electronic and mechanical behaviour of GO have used approximate models based on the Lerf-Klinowski model [198, 199]. Notwithstanding this, this work posits that randomly distributed oxygen containing groups represent an unnecessarily poor approximation for the description of GO.

Graphene oxide is most often made by the Hummers' method [206, 207], where potassium permanganate oxidises graphite in an acidic solution. This method will typically make GO with a carbon to oxygen ratio (C/O) of 2 [208]. C/O is a popular metric to characterise GO because it is experimentally easy to obtain and

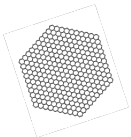

gives a simple measure of the extent of oxidisation. The rate of oxidisation at a graphitic site via permanganate, $MnO_4^-$, is predominantly influenced by the stability of the intermediate structure: graphene-$MnO_4^-$ [58, 209]. Yang *et al.*[58] found that the intermediate state is made more stable by the breaking of adjacent $\pi$-bonds, steric availability, and hydrogen bond formation with the $MnO_4^-$ ion. An important conclusion from Yang's work and the analysis to follow here is that, once a graphene sheet has been oxidised, the rate of oxidation adjacent to an oxidised site is very likely to be more than $10^{20}$ times faster than at a pristine graphene site. An initial oxidation reaction on pristine graphene then acts as a nucleation site, from which more oxidation can proliferate. The disruption of the $sp^2$ network, and the structure of that disruption, is well known to have an effect on the mechanical and electronic properties of the resulting material [108, 179]. Therefore, we must study the structure of the GO produced by this method to get a better understanding of its properties.

In this chapter, I describe a method developed to build large atomistic models of GO based on the local reactivity of graphene systems. I then use the results from this model to study the continuum percolation threshold of GO systems.

## 6.1 Atomistic Model

I used a machine learning approach to extend the subset of reactive sites Yang studied to any possible reactive site that could be encountered on a GO sheet. Through this method we can generate GO structures based on empirical and theoretical observations rather than a random generation which is currently the norm. This method is encapsulated within a program that systematically oxidises graphitic structures for atomistic simulation [8]. The program is open source and can generate structures for a variety of simulation requirements; here we will describe and assess the structures generated.

Yang *et al.* used DFT to calculate the relative stability of different GO-$MnO_4^-$ structures and thereby deduced the rates of reaction of various oxidation reactions, see figure 6.2. I refer the interested reader to their work [58]. I characterised the

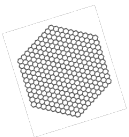

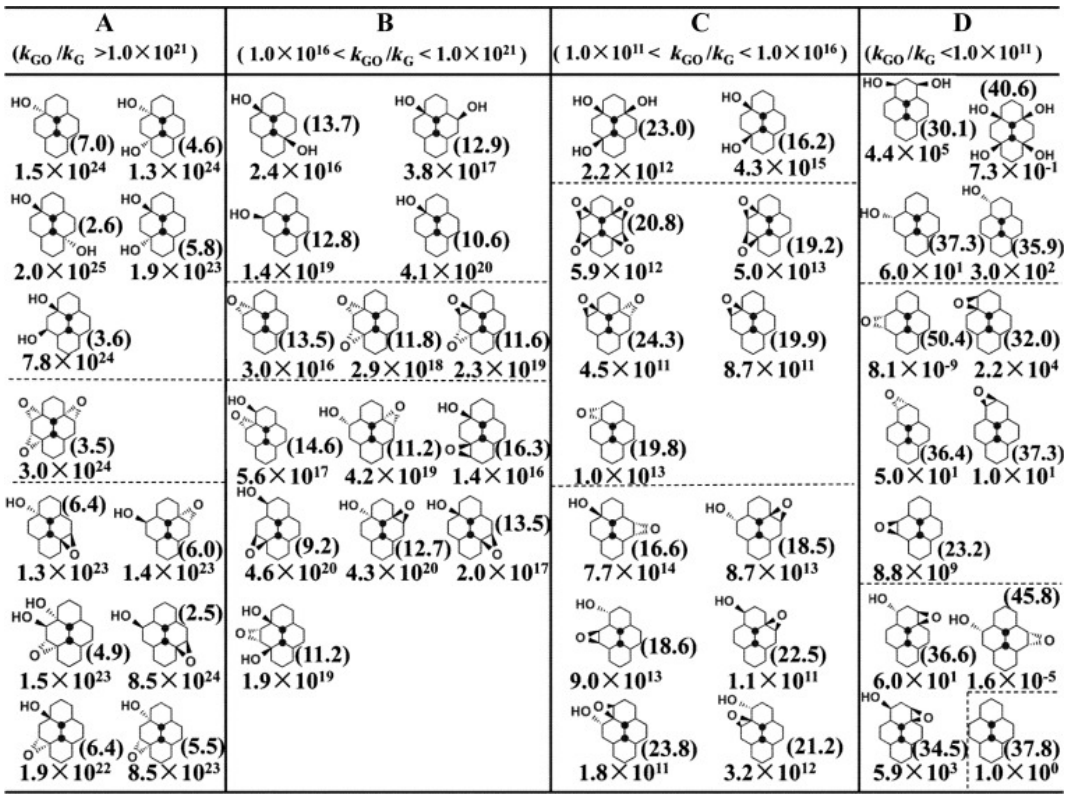


**Figure 6.2:** Relative rates of oxidising different GO sites by a $MnO_4^-$ ion attaching at the centre of each. Pristine graphene has a relative rate $k = 1$ (bottom right). The values in brackets denote the reaction energy barrier (kcal mol$^{-1}$). Figure taken from reference [58]

intermediate's structure as follows.

The $MnO_4^-$ ion attaches to a pair of bonded carbon atoms, the ion sitting on one face of a graphene(oxide) sheet. There are then 4 first neighbour carbon sites and 8 second neighbour carbon sites, the oxidation state of which significantly effect the stability of the intermediate structure. We record the number of first and second neighbour carbons with an alcohol group and epoxy group, and on which side of the flake they reside relative to the $MnO_4^-$ ion. Given that approximately 16000 combinations of these parameters exist [58] and we only have a small set of 52 to train with, we must reduce the system's characteristics. We simplify the problem by assuming the direction of neighbouring functional groups is not important, and we only consider whether they are first or second neighbours. This assumption may become an issue when there are multiple neighbouring oxygen groups, but it has a good foundation based on Yang *et al.*'s work. They deduced that the reactivity is a

function of broken adjacent $\pi$-bonds, steric availability, and hydrogen bond formation with the $MnO_4^-$ ion — it should be noted that they did not include directionality as an important factor. It is difficult to say if these factors are directly dependent on the relative positions of multiple functional groups; the most problematic case might be that steric availability may differ if two groups are adjacent or on opposite sides of the permanganate ion. However, reducing the feature set was necessary and I assumed it did not impact the local structure significantly.

Now, each potential reactive site has 8 features: the number of alcohol or epoxy groups in first or second neighbour position and on the same or opposite side of the flake, relative to the $MnO_4^-$ ion.

Two methods for predicting the rate of a GO site were developed and are outlined below. One is an empirical method based on the findings of Yang *et al.*, the second uses a random forest ensemble learning method to predict the rates.

### 6.1.1 Empirical Model

The three most important factors to a given site's reactivity are: broken adjacent $\pi$-bonds; steric availability; and hydrogen bond formation with the $MnO_4^-$ ion, according to Yang *et al.*'s research. In line with their findings, the following assumptions are introduced:

- Steric hinderance is only significant from first neighbour oxygen containing groups

- Hydrogen bonding only stabilises the intermediate from a second neighbour alcohol group

- Disruption to the conjugate $\pi$-network is independent of which face the oxygen containing group is on

The features of each state are then:

$$
\begin{aligned}
p_{OH} &= \text{\# first neighbour C-OH (alcohol groups)} \\
p_O &= \text{\# first neighbour C-O (epoxy groups)} \\
s_{OH} &= \text{\# first neighbour C-OH on the same face as the } MnO_4^- \text{ ion} \\
s_O &= \text{\# first neighbour C-O on the same face as the } MnO_4^- \text{ ion} \\
h_{OH} &= \text{\# second neighbour C-OH on the same face as the } MnO_4^- \text{ ion} \\
p &= p_{OH} + a_1 p_O \\
s &= s_{OH} + a_2 s_O \\
h &= h_{OH}
\end{aligned}
\tag{6.1}
$$

Then define three functions that effect the rate of oxidation at a particular site:

$$
\begin{aligned}
P(p) &: \text{Disruption to } \pi\text{-network} \\
S(s) &: \text{Steric hinderance} \\
H(h) &: \text{Hydrogen bonding to intermediate structure} \\
k &= 10^{P+S+H}
\end{aligned}
\tag{6.2}
$$

where $k$ is the rate constant. Raising to the power 10 is necessary to regularise the parameters of our model.

There are 5 features per site, using a second order polynomial to approximate the three functions above:

$$
\begin{aligned}
P(p) &= c_1 p + c_2 p^2 \\
S(s) &= c_3 s + c_4 s^2 \\
H(h) &= c_5 h + c_6 h^2
\end{aligned}
\tag{6.3}
$$

Here we have 8 parameters to fit: $a_i\{i = 1,2\}$ and $c_i\{i = 1,\ldots,6\}$. This is still very large and one must be aware that agreement will be approximate at best. We

will compare the results of this model with others later.

Using `scipy.optimize.curve_fit` [210] we obtain the following parameters:

$$
\begin{aligned}
a_{1,2} &= \{1.52,\ 0.68\} \\
c_{1-6} &= \{-3.35,\ 0.25,\ 23.24,\ -5.36, 11.26,\ -4.22\}
\end{aligned}
\tag{6.4}
$$

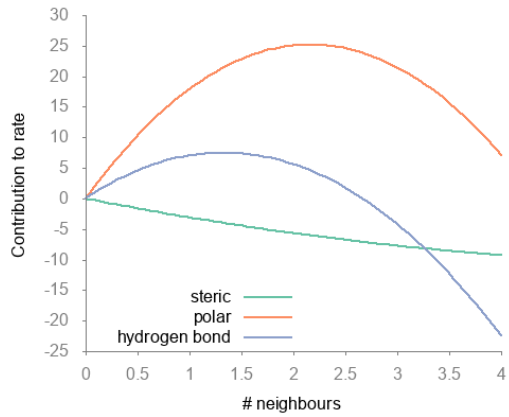


**Figure 6.3:** Functions used to estimate reactivity of graphene (oxide) sites. Functions refer to those given in equation 6.3 and 6.4.

Plotting the functions given in equation 6.3, gives the curves in figure 6.3. These broadly fit my expectation and chemical intuition.

### 6.1.2 Random Forest Model

I then tried to use a random forest ensemble learning method to predict the rates of oxidation at all possible sites. Given such a small training set, I found that some machine learning techniques did not perform well to predict the reactivity of different sites. The available data is far too sparse to train a neural network for example. However, a decision tree or random forest (RF) approach worked well (probably because the features were discrete. i.e. the number of first neighbour alcohol groups above the plane is an integer ranging from 0 to 4). There are 8 features: 2 different oxidation types, first/second neighbour, above or below the flake. We used the Scikit-learn software to generate our RF model [211]: the RF had a maximum depth of 4 and the output is an average of 500 estimators.

| Feature | | | Importance |
|---|---|---|---|
| Group | Neighbour | Position | |
| -OH | $1^{st}$ | below | 0.2880 |
| -OH | $1^{st}$ | above | 0.2120 |
| -O- | $1^{st}$ | below | 0.1872 |
| -O- | $2^{nd}$ | below | 0.0937 |
| -OH | $2^{nd}$ | above | 0.0825 |
| -OH | $2^{nd}$ | below | 0.0588 |
| -O- | $2^{nd}$ | above | 0.0523 |
| -O- | $1^{st}$ | above | 0.0255 |

**Table 6.1:** Feature importance for the random forest regressor described in the main text.

It was also necessary, when training the RF, to use the logarithm of the reactivities. This is a method to regularise the data set and is necessary for values such as these, which span many (31) orders of magnitude; otherwise the largest values would completely dominate the RF fitting process.

One way to characterise a RF regressor is to output the feature importances, shown in table 6.1. This calculated from how high up the 'tree' a feature is used to categorise the data, giving an indication of which feature has the most effect on the outcome. We can see that first neighbour alcohols have the highest importance. Surprisingly, a first neighbour epoxy group has little importance, we can explain this by proposing its steric hinderance to oxidation is similar to its polarity contribution.

### 6.1.3 Validate models

To validate our model we used various metrics to measure the model's success. The first was the coefficient of determination, $R^2$:

$$R^2(X) = 1 - \frac{\sum_i (y(X_i) - \overline{y(X)})^2}{\sum_i (y(X_i) - f(X_i))^2}. \tag{6.5}$$

For a data set $X$, the true rates are denoted $y(X)$, and the model predictions $f(X)$.

For this validation our data was split into a training, $X_{\text{train}}$, and test set, $X_{\text{test}}$, with a split of 39:13 (75% : 25%). The model was fitted to the training set and then assessed with the test set. This was repeated on several randomly generated

| | Empirical model | Random Forest | Extra Trees | MLP |
|---|---|---|---|---|
| $R^2(X_{\text{train}})$ | 0.67 | 0.72 | 0.66 | 0.87 |
| $R^2(X_{\text{test}})$ | 0.43 | 0.29 | 0.18 | -0.13 |
| Predict top $S_5(X_{\text{test}})$ | 0.59 | 0.53 | 0.56 | 0.45 |
| Average position $S_5(X_{\text{test}})$ | 1.64 | 1.82 | 1.77 | 2.03 |
| Predict top $S_{20}(X_{\text{all}})$ | 0.20 | 0.33 | 0.37 | 0.31 |
| Average position $S_{20}(X_{\text{all}})$ | 3.21 | 2.66 | 2.45 | 3.20 |

**Table 6.2:** A summary of the performance of different models. Random Forest and Extra Trees model use a depth of 4, MLP (Multi-Layer Perceptron) uses one hidden layer containing 8 nodes. $R^2$ values are the average of 50 fits on randomly generated $X_{\text{test}}$ sets; predict top and average positions were generated from 100 sub sets of each of the 50 fits.

training/test sets to cross-validate our scores.

Our RF model achieves $R^2(X_{\text{test}}) = 0.29$, which shows some level of correlation but is far from the ideal value of 1. Note that $R^2(X_{\text{train}}) = 0.72$, because we have reduced the features of the system, the model will never predict exact values for all of the training set because some of the information has been lost.

It is more important that this model correctly predicts the most reactive sites than the least, so the sum of all residuals is likely not to be the best measure of the model's suitability. The second method we use to evaluate the model is how well it predicts the most reactive structure from a subset. Taking a subset of size $n$, $S_n(X)$, of the set $X$, we can rank our model by how often it correctly predicts the most reactive sites and the average ranking of the site it predicts to be the most reactive.

For a subset $S_5(X_{\text{test}})$, the RF model predicts the most reactive site correctly 53% of the time. For a larger subset (reusing some of the training data points to increase the size) $S_{20}(X)$ it predicts correctly 30% of the time. By these measures the RF model performs far better than other methods we tried such as the empirical model or multi-layer perceptron neural networks.

Multi-layer preceptors (MLPs) of several different architectures were attempted and none achieved a good fit to a test set. Neural networks typically need large data sets and the available data was clearly not enough. Smaller networks did not improve the test set scores, so even though a good fit with the training data was achieved we do not believe the model was overfitting.

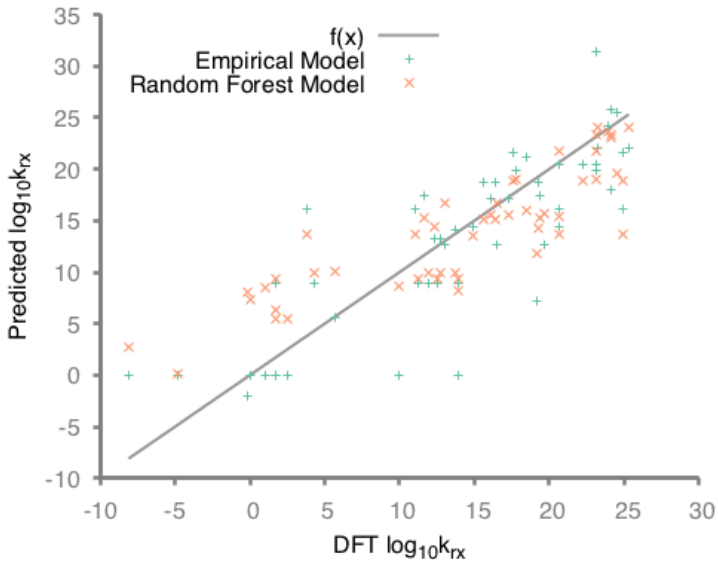


**Figure 6.4:** Residual plot for the empirical and random forest models described above. These predictors were trained on all 52 samples.

RF (and extra tree) models show promise. Although their $R^2$ of test sets were consistently lower than the empirical model their ability to predict relatively the most reactive site was comparable amongst test sets and significantly better for larger sets. A summary of all the test scores for each model are given in table 6.2.

Residual plots of the both the empirical and RF models are shown in figure 6.4. These plots show the loose correlation between the rate calculated with DFT and the predicted rate from the models described here.

Both the RF and empirical model were used to simulate GO growth. A nice feature of the empirical model is that is designed in such a way that prevents over fitting and one can intuit how the rates are calculated. However, it only takes a limited number of features, whereas the RF can include more.

An outline of the way we oxidise a graphitic structure, using the program given in Ref [8], is shown in figure 6.5. As is described by Yang *et a* [58], we assume that whether an alcohol or epoxy group forms after a permanganate ion binds to a C-C bond is random and equally likely. The function that decides whether to nucleate a new island takes two arguments: a user specified 'nucleation frequency'

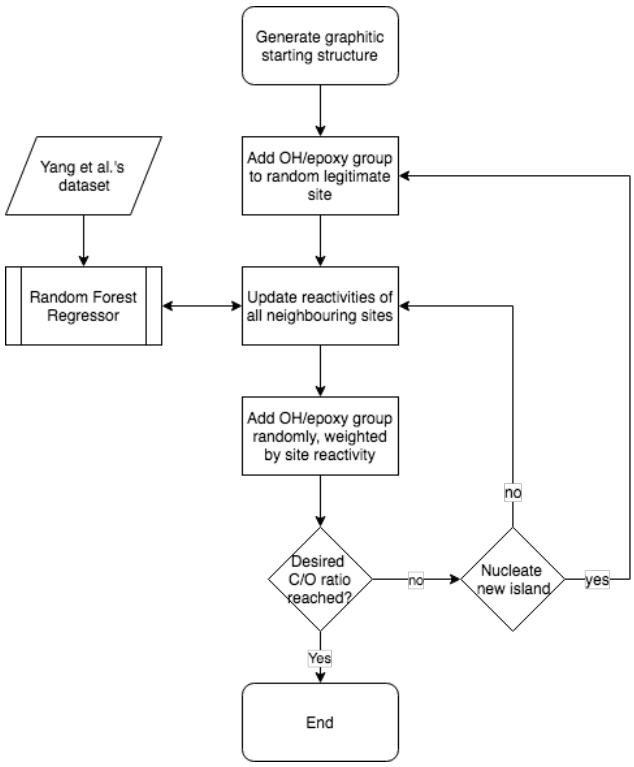


**Figure 6.5:** Simple flowchart representing the general method taken to oxidise a graphitic structure using a random forest regressor to calculate site reactivites.

(this parameter is described in the next section as $k_n$); and the total reactivity of the system. We use the total reactivity of the system to predict the time elapsed between adding OH or epoxy groups. New islands are added according to a Poisson distribution, with mean equal to the 'nucleation frequency' multiplied by the time between adding functional groups.

An example of a very small graphene flake oxidised using our program is shown in figure 6.1b. At this scale it looks similar to the Lerf-Klinowski model but the location of the oxygen containing groups is highly correlated. The most obvious difference comes when larger areas are oxidised, as seen in figure 6.6: the large oxidised region propagates from its nucleation site, structures emerge such as two phases of oxidised and unoxidised domains, and aromatic pockets within the oxidised island.

The structures generated are qualitatively similar to high-resolution microscopy images of GO [201, 202, 203, 204]; specifically, amorphous alcohol and epoxy groups make up the oxidised regions with unoxidised islands on the

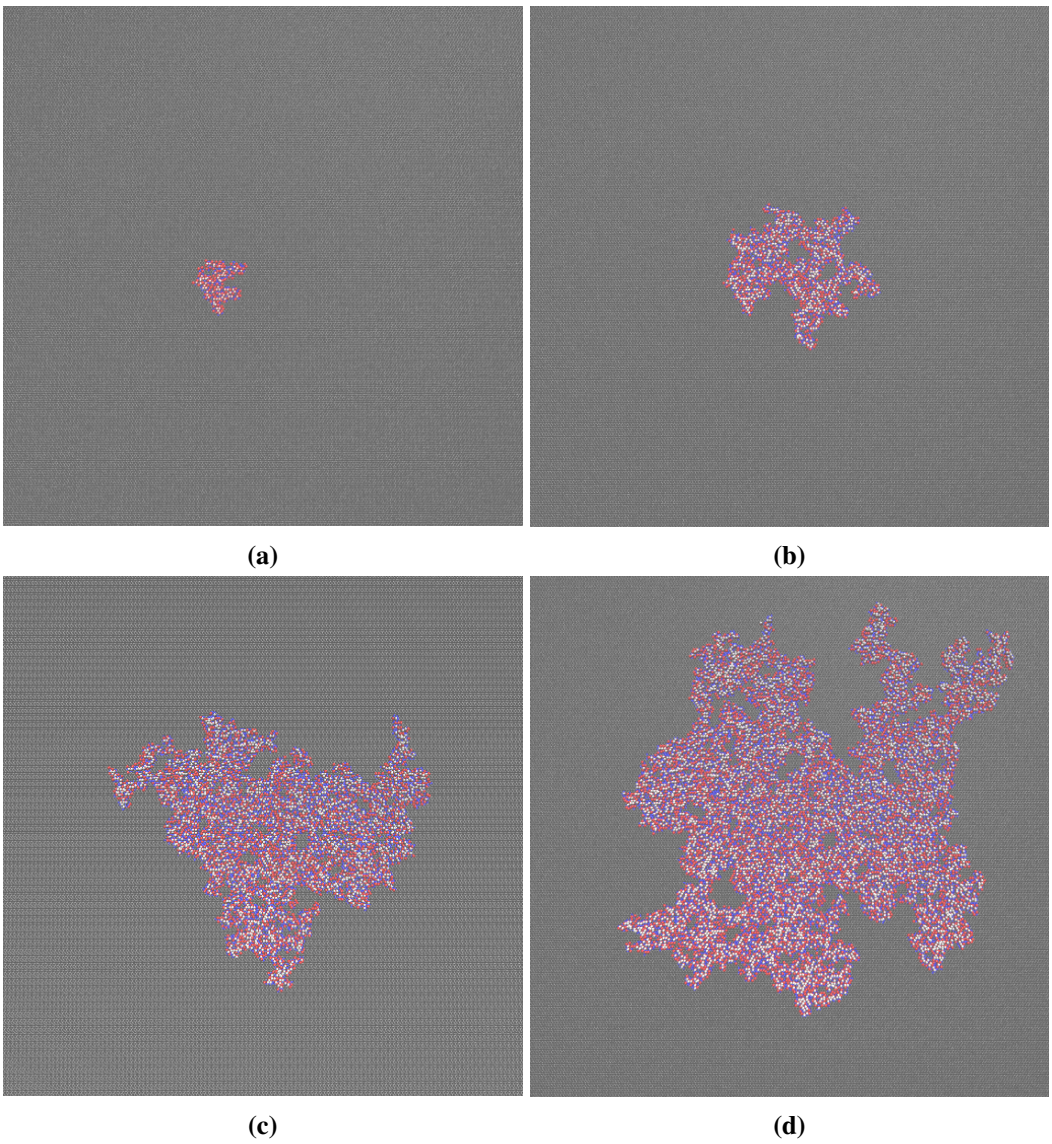


**Figure 6.6:** A representative example of a $50 \times 50nm^2$ graphene sheet nucleated and oxidised systematically using the approach described in the main text, evolution of the oxidised region is depicted through (a)-(d). Aromatic carbons: grey; oxidised carbons: blue; oxygen: red; hydrogen: white

nanometre scale. A random placement of oxygen containing functional groups, as described by the Lerf-Klinowski model, would not recreate these inhomogeneous phases.

The training data available is probably biased toward highly reactive sites (because the most reactive sites were of interest in the original study), and so the termination of our builder is less reliable. The average predicted reactivity of oxidised sites starts to decrease at a carbon oxygen ratio of $C/O \approx 2$, comparable to experi-

ment, for which simple oxidation of graphene normally gives the same ratio [208].

Clearly the models are far from perfect predictors of the reactivities calculated by Yang *et al.*, as can be seen in figure 6.4 and table 6.2.

## 6.2 Percolation Analysis

The atomic structure of GO (shown in figures 6.1b and 6.6) may have important implications for its physical properties and interactions with other molecules, including other GO sheets. It is valuable to have an accurate way to generate this structure, but the arrangement of alcohol and epoxy groups within an oxidised region does not itself appear to form a discernible pattern. The structure and evolution of these oxidised regions, however, is of great importance. Here, we design a model to study the properties of this two-phase system.

It is obvious from figure 6.6 that GO could be approximated to a two phase system, namely a purely graphitic phase and a graphene oxide phase which increases in size. To study the mesoscale evolution of a graphene sheet undergoing oxidisation we constructed a continuum model [9].

The reactions requiring consideration are:

$$\mathrm{MnO_4^- + G \xrightarrow{k_n} G_O}, \tag{6.6}$$

$$\mathrm{MnO_4^- + G_r \xrightarrow{k_{rx}} G_O}, \tag{6.7}$$

where $G$ denotes a graphitic site, the subscript $r$ a reactive site (i.e. near to an already oxidised graphitic carbon), and $O$ an oxygenated site. $k_n$ and $k_{rx}$ are the rate constants of the nucleation and catalysed reactions respectively; as discussed above $k_{rx} \gg k_n$. After a graphene sheet is nucleated by an oxygen site, $G_O$, we consider the oxidation reaction as a propagating circle around that nucleation site, a reasonable approximation as one can see from the shape of the island in figure 6.6 (approximating the boundary of an island is discussed in more detail in section 6.2.2.1).

Considering a propagating oxidised island of radius $r$ on a very large graphene

sheet, the area of oxidised graphene is $A_O = \pi r^2$. The reactive area of graphene $A_r$ is defined as the narrow strip, of width $w$, around the circumference of the oxidised island:

$$A_r = \pi \left[ (r+w)^2 - r^2) \right] \tag{6.8}$$

$$\approx 2\pi r w, \text{ when } w < r, \tag{6.9}$$

where $w$ approximately the length of a carbon bond. Rearranging the equation for the area of oxidised graphene, $r = \sqrt{A_O/\pi}$, and substituting into the above equation, we find:

$$A_r = 2w\sqrt{\pi A_O}. \tag{6.10}$$

As discussed above, the oxidation of graphene is limited by the formation of the graphene-$MnO_4^-$ intermediate structure; we assume the reaction is elementary and construct the rate law for equation 6.7:

$$\frac{\mathrm{d}A_O}{\mathrm{d}t} = k_{rx}[\mathrm{MnO_4^-}]\mathrm{A_r} = \mathrm{k_{rx}}[\mathrm{MnO_4^-}] \cdot 2\mathrm{w}\sqrt{\pi \mathrm{A_O}} \tag{6.11}$$

Assuming that the concentration of $MnO_4^-$ remains constant (in experiment it is added in excess), we have that:

$$A_O = \left(k_{rx}[\mathrm{MnO_4^-}] \cdot \mathrm{w}\sqrt{\pi}\mathrm{t}\right)^2 \tag{6.12}$$

$$A_r = k_{rx}[\mathrm{MnO_4^-}] \cdot 2\mathrm{w}^2\pi\mathrm{t} \tag{6.13}$$

Recalling $A_O = \pi r^2$, we can see from equation 6.12 the radius of an oxidised island grows at a constant rate.

$$r = k_{rx}[\mathrm{MnO_4^-}] \cdot \mathrm{wt} \tag{6.14}$$

$$\frac{\mathrm{d}r}{\mathrm{d}t} = k_{rx}[\mathrm{MnO_4^-}] \cdot \mathrm{w} \tag{6.15}$$

I model a graphene sheet as a square, the oxygenation is nucleated at a random point and the oxidised island's radius increases at a constant rate. The primary interest is identifying the percolation threshold: past the tipping point where there is no continuous area of conjugated aromatic carbons one can expect its electrical and mechanical properties to steeply degrade. The percolation threshold of this system is defined as occurring when there is no continuously connected path in physical terms that connects opposite edges of the square via unoxidised regions of graphene i.e. it cannot conduct electrically from one edge to another. This is a special case of an established problem in mathematics of finding the 2D continuum percolation threshold with fully penetrable disks [212, 213, 214].

For the case where the rate of nucleation, $k_n$ is insignificant compared to $k_{rx}$, there will be only one oxidised island present. By observing atomic precision images of GO [201, 202, 203, 204] it is clear that nucleation of oxidised regions happens at more than one point on a graphene sheet. While oxidation may be vastly ($10^{20}$ times) faster near oxidised sites than pristine graphene, we know that most samples of graphene are not pristine and contain many defects. These defects could feasibly encourage nucleation, raising $k_n$, the rate of nucleation. We can then predict the effect of the ratio of $k_{rx}$ and $k_n$ on the resulting material. From now, we absorb $[\mathrm{MnO_4^-}]$ into the rate constant for clarity:

$$\frac{\mathrm{d}A_O}{\mathrm{d}t} = k_{rx}A_r \ , \ \frac{\mathrm{d}N_{\mathrm{islands}}}{\mathrm{d}t} = k_n(A - A_O), \tag{6.16}$$

where $A$ is the total area of graphene, and $A_r$ is the strip of graphene of width $w$ adjacent to all oxidised graphene sites. With the possibility of several nucleation sites, so that propagating islands can overlap, this problem must be approached numerically.

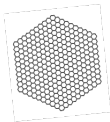

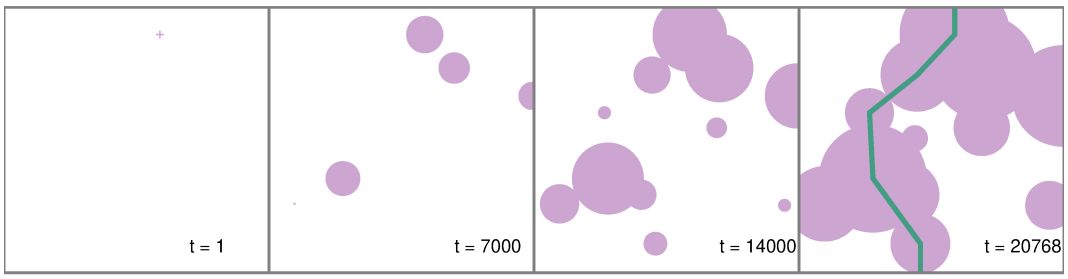


**Figure 6.7:** Typical example of a graphene oxidation simulation. The algorithm is described in the main text. A node is added at step 0. Pink regions represent oxidised regions. The simulation is stopped at step 20768, when a path can be made from one edge to its opposite, shown by the green line. $\delta t = 10^{-5} s^{-1}$

### 6.2.1 Program

It can be seen that all sets of systems that satisfy $Ak_n/k_{rx} = \chi$ behave identically when considering the fractional coverage of the system, $\phi$, where $\chi$ is a dimensionless constant that characterises the system. For example, a larger system which has a slower nucleation rate would reach its percolation threshold at the same fractional coverage. We use a unit area sheet and $k_{rx} = 1$ s$^{-1}$ for simplicity; we also assume that $k_{rx}$ is independent of $k_n$ and we use different values of $k_n$ to assess all possible systems.

The algorithm advances as follows:

1. A nucleation site (node) is added to a square cell
2. The island centred on each node has its radius increased by $\delta r$
3. New island nodes are added
4. Repeat steps 2-4 until no continuous unoxidised region exists.

In step 2, $\delta r$ is proportional to $\delta t k_{rx} \sqrt{A}$. Step 3 is achieved by adding a number of new nodes drawn from a Poisson distribution with mean $k_n A \delta t$, only accepting nodes that fall in unoxidised regions. The procedure terminates when a path can be made from one edge to its opposite with overlapping islands (see figure 6.7). For periodic 2D systems this has been postulated many times to be equivalent to the

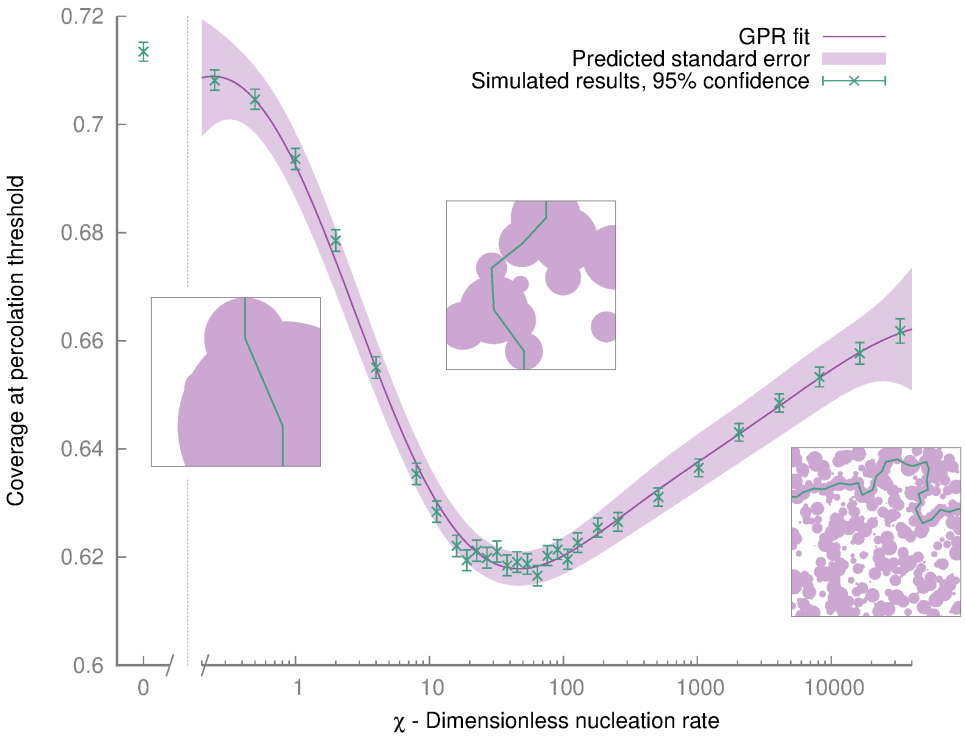


**Figure 6.8:** Variation of percolation threshold by varying $\chi = Ak_n/k_{rx}$. By varying $\chi$ all unique systems can be tested. At low values, the system tends to a percolation threshold of 0.71, a minimum is reached at $\chi = 45$, before rising logarithmically. Data from our simulations are shown in green, a curve fit and confidence interval are shown in purple, fit via a gaussian process regression. Snapshots of typical simulations at the percolation threshold with different values of $\chi$ are shown in boxes; purple areas correspond to oxidised islands; the green lines indicate the first path that can be made from one edge to its opposite, i.e. there is no longer a continuously connected unoxidised region.

percolation threshold [212, 213, 214]. Here we apply it to a non-periodic system as graphene sheets have edges.

### 6.2.2 Results

The fraction of graphene that has been oxidised at time $t$ is $\phi(t)$. The critical time at which the percolation threshold is reached is denoted $t_c$. If the algorithm reports that a path can be made between two opposite edges with oxidised regions at time $t'$, we know that $t_c$ lies between $t'$ and $t' - \delta t$. We therefore report the cell coverage at the percolation threshold as:

$$\phi(t_c) = \phi(t') - \frac{\phi(t') - \phi(t' - \delta t)}{2} \tag{6.17}$$

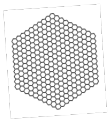

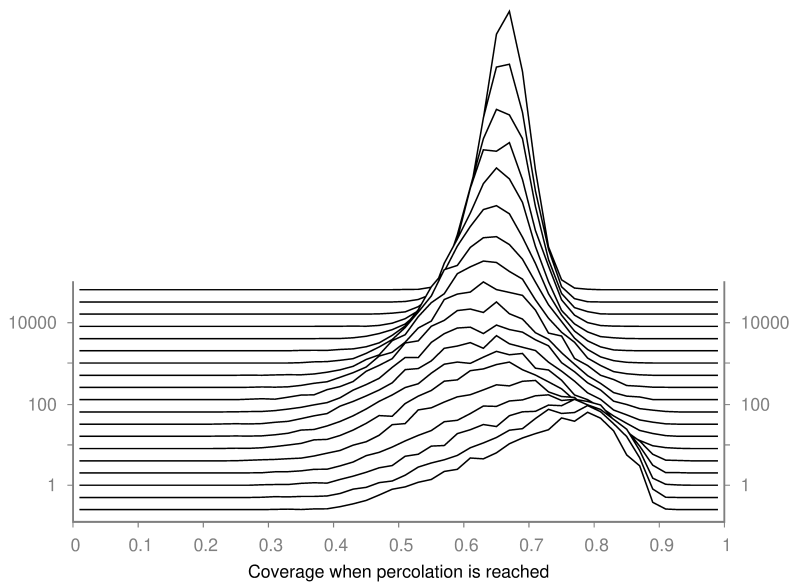


**Figure 6.9:** Distribution of fractional coverage at percolation thresholds at different values of $\chi$. The distribution becomes narrower as the nucleation rate increases. Each line represents a different value for $\chi$, noted by where the value meets the vertical axis, the area under each line is equal to one.

The fractional coverage in figure 6.8 is reported as the average of 10,000 simulation runs. Error bars are the 95% confidence interval based on a bootstrap analysis on the simulation runs plus the algorithmic error, taken as the average value of equation 6.17. Reducing $\delta t$ increases the accuracy of each run but the additional computational cost means less simulations can be run so the confidence interval increases. Results are within error for different values of $\delta t$ showing that our results are independent of the variable $\delta t$.

From figure 6.8 we can see that a minimum percolation threshold exists when $\chi = 45$. Below this value the model has fewer islands and more coverage is required to reach the percolation threshold, tending asymptotically to a value of 0.71 for a system where no additional nucleation is permitted ($\chi = 0$). Above $\chi = 45$, the percolation threshold rises logarithmically; we did not simulate higher values of $\chi$ as precision errors in the model become more pronounced and no new behaviour is observed. What mechanism underlies this relationship between $\chi$ and $\phi(t_c)$ is not known but the competing mechanisms are interesting. Asymmetries, similar to this case, in the percolation threshold of circles with different radii have been

observed before [213], but the origin of this phenomenon has not been explained. The distribution of coverage at the percolation threshold is shown in figure 6.9.

Recalling that $\chi = Ak_n/k_{rx}$, a real sample of graphene flakes will have surface areas that cross several orders of magnitude. This is inevitable by any of today's current methods of synthesising graphene. Therefore, a sample of graphene will not have a unique value of $\chi$ that can describe it. It is clear, however, that there is minimum average value that $\phi(t_c)$ takes and within one standard deviation it can be shown from the data in figure 6.9 that $\phi(t_c)$ will always be greater than 0.58. Similarly, we can say that, when $\phi(t_c) > 0.43$, at least 95% of the sheets will be below the percolation threshold, whatever the flake size distribution. It would be pertinent to know what C/O ratio this corresponds to.

#### 6.2.2.1 Calculating the C/O ratio of an oxidised domain

In this continuum model, oxygenated regions are approximated to a circle with even density whereas in reality they are more irregularly shaped and have unoxidised islands that must be accounted for. Calculating where the effective boundary of this region lies (to approximate it as a circle) is not a trivial task.

Using the RF algorithm described in the previous section, a sheet was nucleated only once and 10,000 oxygen atoms were added according to the calculated reactivities to analyse the resulting oxidised region. The values discussed are taken from an ensemble of 15 such simulations to estimate the density of oxygen atoms in a propagating graphene oxide phase.

Figure 6.10 plots the O/C ratio as a radial distribution of oxygenated sites from the centre of mass of the oxidised region (the O/C ratio is used here so the value does not tend to infinity at large radii). The radial distribution forms a plateau near the centre of the island, which can be taken as the O/C ratio of a fully oxygenated region. The average O/C ratio between $0 < r < 5$ nm is 0.53. This cut-off is somewhat arbitrary but is necessary so that we can derive an effective area of the island. We can then say that the radius at which one is equally likely to find oxidised and unoxidised region has a O/C ratio of 0.26: reading from figure 6.10, this is at $r = 12.0 \pm 0.7$ nm.

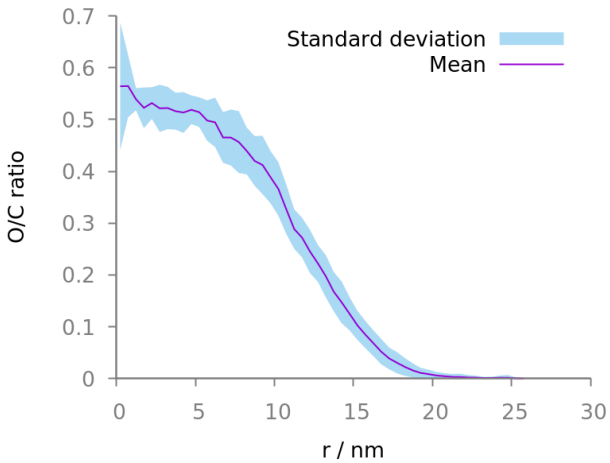


**Figure 6.10:** Oxygen density as measured from the centre of an oxidised island. The island has a higher oxygen density near the centre. We use an O/C ratio (instead of the usual C/O) because it does not rise to infinity at large radial distance. It is not immediately obvious what the effective radius of this island is.

However, this constitutes a lower bound to the effective radius of an oxidised island. We are concerned with the point at which these islands overlap, so we want to know about the island's boundary. The offshoots that grow from the edges constitute the long tail in figure 6.10 and are more important than the density at the centre.

One way to measure the area of the oxidised region is to find the area of the convex hull that encloses all the oxygenated sites. The convex hull can be thought of as the shape enclosed by an elastic band stretched around all oxygenated sites. The effective radius of this island would be the radius that gives a circle with the same area as this convex hull. Over the same 15 simulations, the radius calculated by this method is $15.1 \pm 0.6$ nm. This method captures the area around the irregular boundary. Some examples are shown in figure 6.11.

In reality, the effective radius of these islands probably lies between the two values calculated. Any value in between can be arrived at by constructing the concave hull of the island with a given shape parameter. A concave hull is a polygon which contains all the oxygenated sites but has less area than a convex hull; several algorithms exist which can be used to construct such a polygon. Here, we constructed a Delaunay triangulation using the coordinates of all the oxygen sites, then

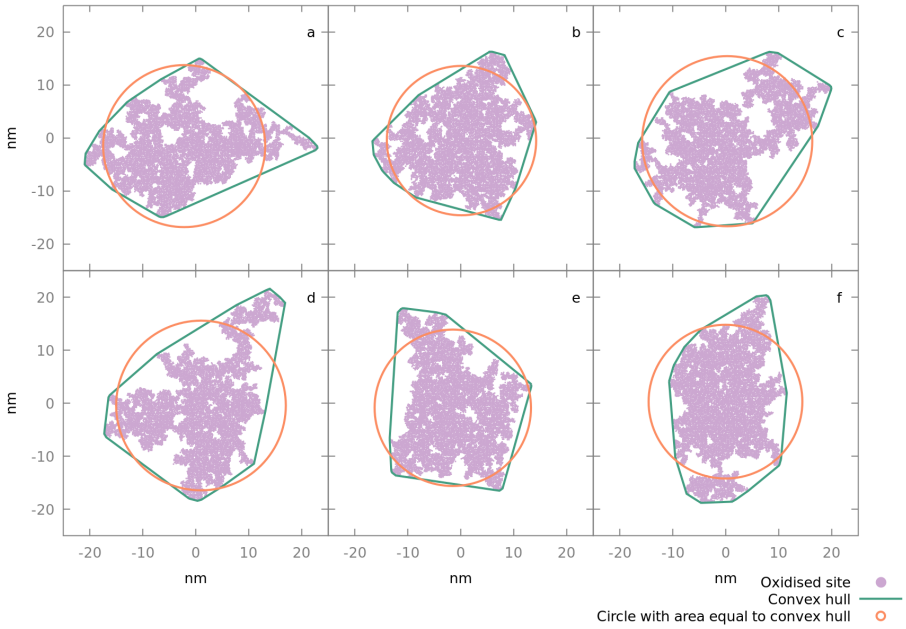


**Figure 6.11:** This figure shows 6 examples, (a)-(f), of oxidised domains. We are trying to calculate the effective radius of these islands to map to the continuum model described in the main text. To do this we draw the convex hull of the island (green) and select a radius which gives the same area, which is also shown in each case (orange).

removed the triangles with edges larger than $\alpha$ [215]. We find the outer edges of the resulting mesh and use this boundary to compute the effective area of the island. Examples are shown in figure 6.12. The code used for this is given in appendix C.1. By sweeping through different values of $\alpha$, the relationship in figure 6.13 is found.

The carbon to oxygen ratio is then calculated by dividing the number of carbon atoms that fall within the effective area of the island by the number of oxygen atoms added (10,000). The density of carbon atoms on a graphene sheet is 38.46 carbons per nm$^2$. Using the first method described (half the radial distribution plateau), the C/O ratio within a propagating oxidised island is $1.94 \pm 0.20$. Using the area of the convex hull, C/O $= 2.76 \pm 0.22$. In the main text we use a ratio of 2.76 which is the most conservative estimate for calculating the percolation threshold of GO. The true value is likely to be smaller.

The C/O ratio of a propagating oxidised region is at most 2.76. All this means,

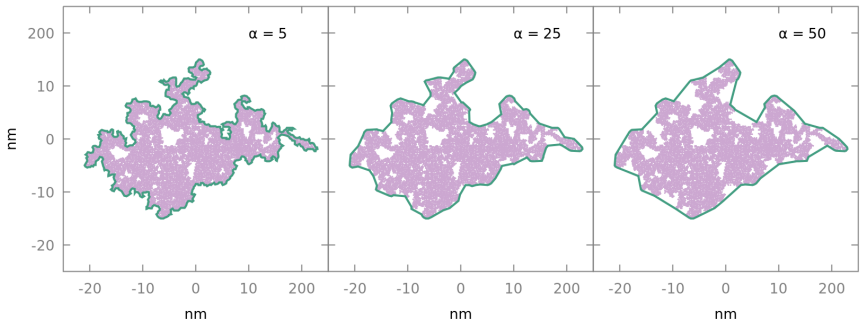


**Figure 6.12:** The convex hull drawn in figure 6.11a probably overestimates an island's effective radius. Here we show the same island but draw the concave hull with different $\alpha$ values. The exact algorithm to produce these boundaries is given in appendix C.

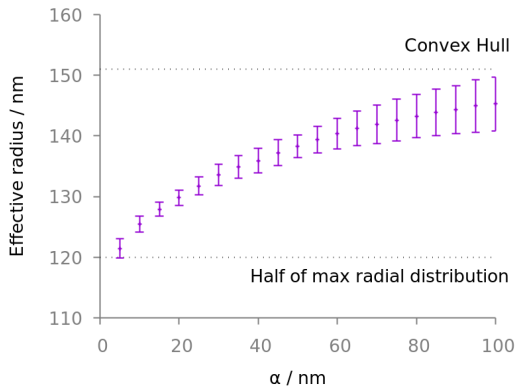


**Figure 6.13:** The effective radius of an oxidised domain with 10,000 oxygen atoms calculated by different means. The convex hull method provides a maximum, and the half maximum of the radial distribution of oxygen atoms is a minimum. The real value is expected to lie between these values.

by a conservative estimate, that the percolation threshold arises for C/O ratios no greater than $2.76/0.46 = 6.0$. It can then be concluded that, if the formation reaction of graphene oxide could be quenched before this point, i.e. if the C/O ratio exceeds 6.0, many of graphene's mechanical and electrical properties could be preserved.

This prediction is also borne out by our atomistic model. Using different nucleation rates, $k_n$, that spanned several orders of magnitude, the percolation threshold was reached at an average C/O ratio of 4.2, and never exceeded 4.5.

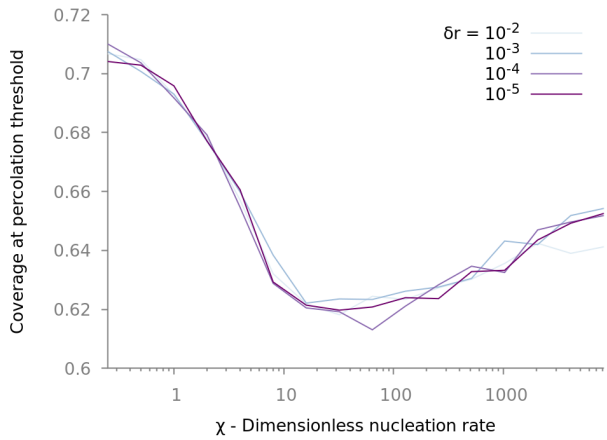


**Figure 6.14:** The fractional coverage of a sheet in oxidised islands against different values of $\chi$, for diminishing discretisation. These are mean values of 1000 runs at each point; the mean is essentially unchanged for different precisions but the error in these values is correlated, as shown figure 6.15.

#### 6.2.2.2 Percolation model sensitivity

The main results gained from these two GO models have been discussed above. Below I will discuss some of the numerical errors associated with the continuum model.

The percolation model has three input parameters: $\chi$, the relative nucleation rate; $\delta t$, the timestep used to increment the islands' radii and how often the algorithm attempts to add new islands; and the number of Monte Carlo points used to estimate the coverage area once the percolation threshold has been reached.

Recall that $\delta r = \delta t k_{rx}\sqrt{A}$. We can show that the qualitative outcome of the simulations, taken as the average coverage at the percolation threshold for different values of $\chi$, is independent of the timestep used; see the plot in figure 6.14. However, the error in the result due to the change in coverage between timesteps is pronounced; see figure 6.15.

The error reported in our method of calculating the coverage for a given system state is straightforward to explain. See figure 6.16, in which the errors shown are relative to the best estimate we had, namely that using $10^6$ points.

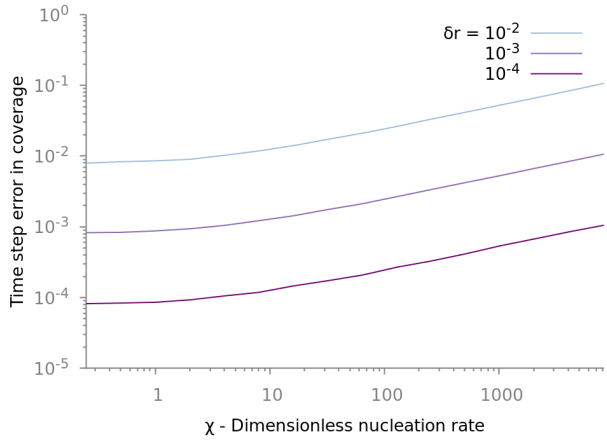


**Figure 6.15:** Average uncertainty in the coverage over 1000 runs for different precisions vs. the nucleation rate $\chi$. The percolation threshold is reached at time $c$, therefore $\phi(t-\delta t) < \phi(c) < \phi(t)$. The uncertainty in the result is reported as $[\phi(t) - \phi(t-\delta t)]/2$.

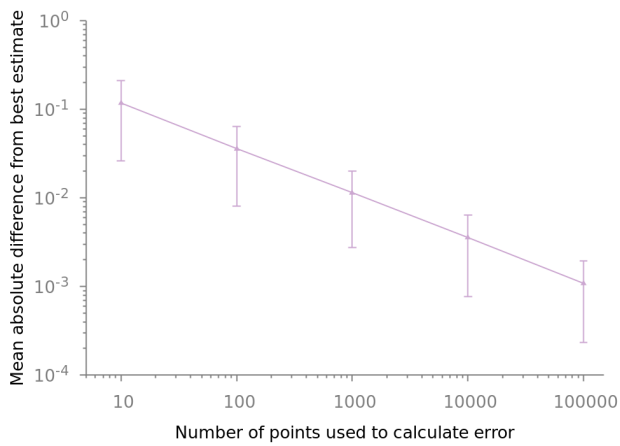


**Figure 6.16:** Average difference in the estimated coverage of the same set of 10000 systems, compared to an estimate using $10^6$ points.

## 6.3 Conclusion

A systematic method to build accurate GO structures has been provided, and by using this nanoscale knowledge one can understand its macroscale structure. This method is encapsulated into a program released alongside this manuscript [8, 9], offering a significant improvement to the Lerf-Klinowski model commonly used in constructing GO structures. To our knowledge, this is the first analysis of the percolation threshold in a graphene oxide synthesis reaction. It is important that GO models have two distinct domains present on the nanoscale, rather than a homogeneous distribution of functional groups. Models that generate random amorphous

regions of oxidised graphene will not vary significantly in their results; by contrast, structures that have large separate aromatic and oxidised domains will drastically affect properties such as aggregation, exfoliation, solvation and adsorption, since the two domains have very different long range interaction characteristics.

In particular, by keeping the C/O ratio above 6, a continuous domain of conjugated carbon atoms will exist, improving the mechanical and electronic properties of GO. We hope that this will serve to inform experimentalists as well as modellers, and help predict the characteristic behaviour of GO, while improving the consistency with which GO can be synthesised. At a minimum the work presented in this chapter provides an improved way to simulate GO on an atomic scale.

The atomic models of GO shown here have been used in separate projects with more applied scientific goals. GO structures built in this way have bene included in a multiscale model of epoxy nanocomposites, which will be described in the next chapter, and a coarse-grained model for polyethylene glycol (PEG)-GO nanocomposites, developed by James Suter and will be discussed in chapter 7.

# Chapter 7

# Future Research

The work presented in this thesis has led to several investigations that are currently being planned or ongoing. The ultimate goal of creating a virtual lab to predict the mechanical properties of a material, knowing only its atomic composition, is an exciting one, and one that our lab believes warrants further investment. The efforts to achieve this are detailed in the next seciton.

## 7.1 Nanocomposites

The design and manufacture of materials like nanocomposites and other advanced multi-functional materials remains a slow, uncertain, expensive and time consuming process [216, 217]. It can take 20 or more years to move a material from initial discovery to the market [23]. Ideally computational models could offer an opportunity to access a wide range of materials configurations rapidly, posing a serious alternative to the traditional experimental trial-and-error process currently used in industry [218, 219] and will surely increase in importance as computational power increases.

Many materials have multiple characteristic length scales which need to be considered to fully understand their behaviour. A 'virtual lab' that could predict a material's properties before needing to synthesise it would need to encompass all these length scales.

In this chapter I present the work I have contributed to a project that hopes to do just this, by linking two simulation methods together in a heterogeneous mul-

tiscale model (HMM). The basic principle of our HMM is to use finite element analysis (FEA) to model an object undergoing an engineering test. FEA handles the complex, evolving boundaries of the problem and calculates the evolving strain-stress response of the object. Instead of using constitutive equations to predict the material's stress-strain response at each finite element, MD simulations predict this relationship by applying the local strain present in a finite element to an MD simulation of the material.

The separation of scales is crucial for this approach to work. The microscale model, in this case the the MD simulations, must correspond to the scale over which the system can be treated as effectively homogeneous. For thermoset polymers the characteristic length scale at the molecular level is the distance between cross-links.

The continuum system is discretised using a generally unstructured mesh on which the FEA is used to solve the linear momentum conservation equation for a locally homogeneous solid:

$$\nabla \cdot \sigma(x) + \rho \mathbf{f}(x) = \rho \frac{\mathrm{d}^2 x}{\mathrm{d}t^2} \tag{7.1}$$

where $\rho$ is the continuum density and $\mathbf{f}$ the volumetric forces applied to the system. I did not contribute to the specific FEA implementation; for the interested reader, see references [4, 131].

### 7.1.1 Epoxy Composites

We quickly realised that running all atom MD with graphene dispersed into thermoplastics was not a feasible system to model. The characteristic length of a thermoplastic is the distance between entanglements, this will vary between polymers but is far larger than the distance between cross-links in thermosets. Another drawback to using thermoplastic is their long dynamic processes such as reptation or crystallisation which can take on the order of seconds. This negates the separation of scales in time and is impossible to fully account for with MD.

The graphene flakes were built using the program described in section 2.5.2

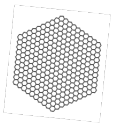

[8] and parameterised using the OPLS force-field [140]. Graphite carbons not at the edge were parameterised as neutral aromatic carbons, as is typical for graphene simulation [118].

To make the polymer simulations, monomer and cross-linker molecules (figure 7.1) were randomly arranged in a periodic cube using a script contained within PACKMOL [220]. This was then compressed and equilibrated for 200 ps until the volume had stabilised.

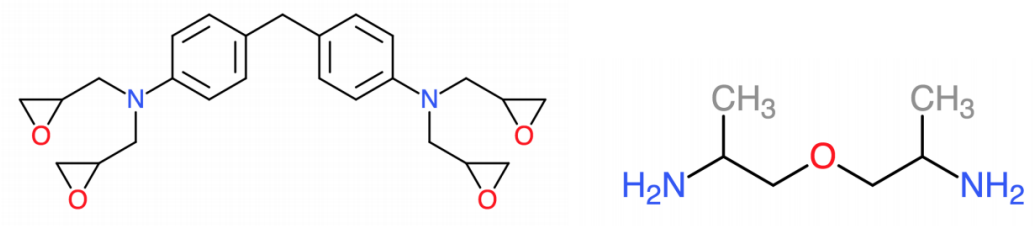


(a) tetraglycidyl methylene dianiline (b) polyetheramine

**Figure 7.1:** Molecular structure of (a) the precursor epoxy monomer and of (b) the crosslinking agent used during the cure of the resin

To cross-link the monomers I wrote a new program so that we could build any polymer system to our specifications [7]. The program proceeds as shown in figure 7.2. The program takes compatible functional groups (in this case an amine and epoxy group) that are within a given radius and forms a bond between them. Neighbouring atoms are deleted or have their labels and charges changed as required. The atom positions are not changed by the program, so after the bond is made it is in a very high energy state. To get rid of this high energy state we run an MD minimization within LAMMPS.

Running a LAMMPS minimization after each bond is made makes this process very costly and also scales as the cube of the simulation's length scale. To reduce the cost, we allow several bonds to be made between minimizations, as long as they are not very close together. This criteria is described in figure 7.2 as the masked region. Once a bond has been made, no bonds within a given mask radius will be allowed to be formed until a minimization has been performed. In this way we avoid the most unphysical situations and allow the code to scale linearly with system size. We chose this mask radius to be 20 Å, chosen to be larger than the length of a

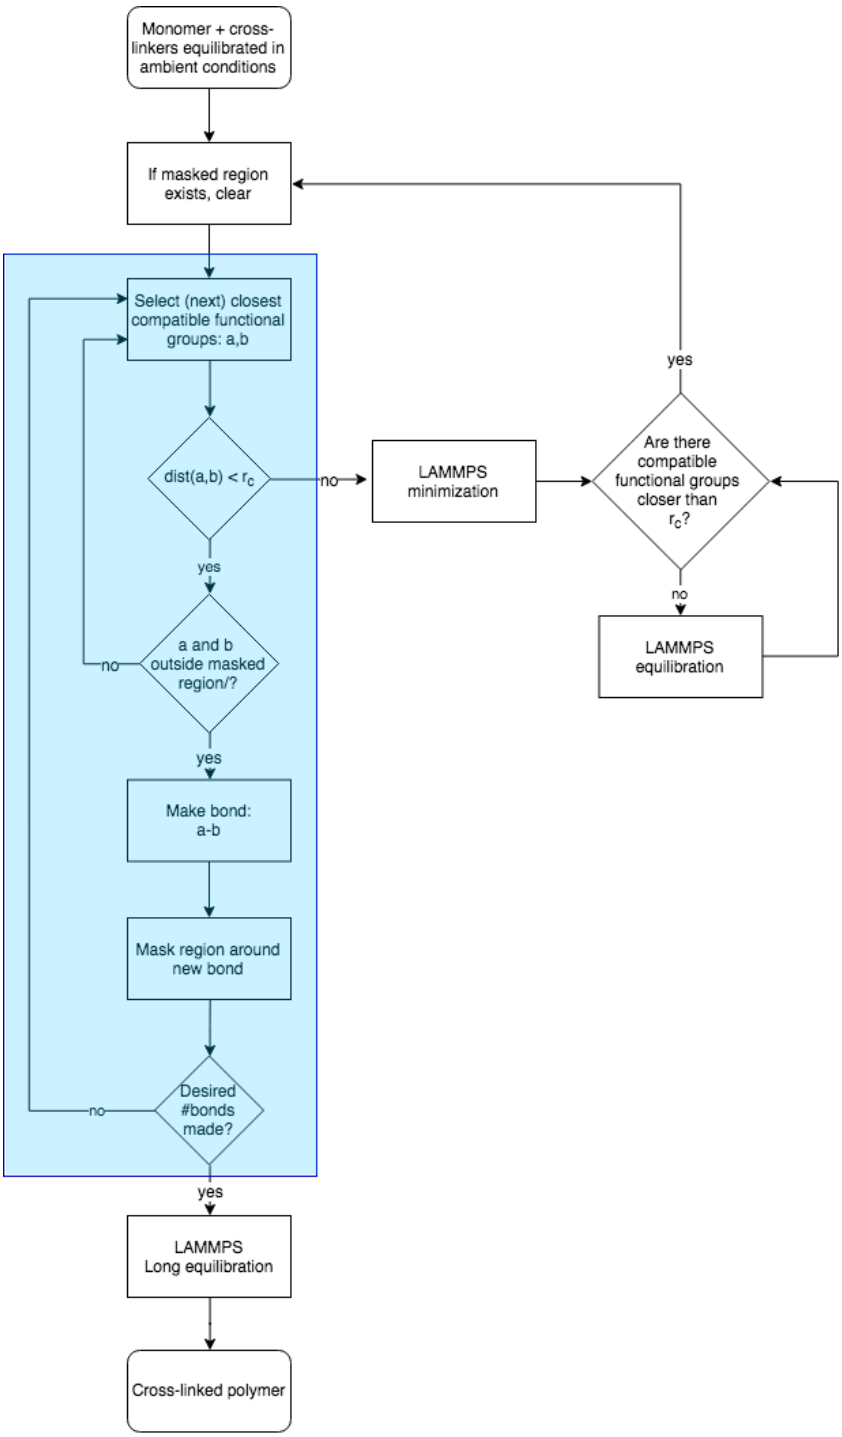


**Figure 7.2:** Flowchart for the program to cross-link an epoxy polymer. When a bond is created, no more bonds are allowed to be made within a given 'mask radius' until high energy configurations are removed with a LAMMPS minimization. The cut-off radius, $r_c$, was 6 nm. The blue box highlights the main loop, which is implemented by one python script. The overall program is controlled from another script which can execute the parallel LAMMPS simulations. This flow chart ignores the difference between primary and secondary amine reactions for simplicity; primary amines were reacted preferably to secondary as they are more reactive in general.

monomer unit which was 15 Å.

The overall logic and the cross-linking section are written in python, this requires LAMMPS data files to be read in and written out periodically to carry out the various minimization and equilibration runs. After much optimisation of the python code, the I/O is now the most time consuming part of the program, at least for our system sizes of 70,000 atoms.

We used the epoxy TGMDA cured with PEA, shown in figure 7.1. For algorithmic reasons, we started with the rings in TGMADA open (we close the epoxy rings of unused branches after we finish curing).

Referring to figure 7.2, the equilibration runs at ambient conditions for 10 ps, the long equilibration runs for 200 ps. The shorter equilibration runs are simply to give functional groups a chance to diffuse towards each other. Real epoxy curing experiments take on the order of seconds or even minutes. This time scale is of course impossible to simulate using MD, therefore this method attempts to mimic the network generation. Building epoxy networks in this way is common place and has been shown to give useful results on a given polymers mechanical properties [221, 222, 130].

In the literature there are several methods for creating atomic models of epoxy resins. Some create all possible bonds in one step [221, 222], another performs a minimization after making 16 bonds [130]; Jang *et al.* compared the two approaches and found minimal differences in the polymers created [223]. How the algorithm affects the polymer network is somewhat unclear; it is also not obvious how to map the cross-link density in the simulation to that measured in experiment. The most common experimental technique for estimating the degree of cross-linking is to measure a sample's equilibrium swelling [224]. This is often erroneous and gives minimal insight into the nanostructure of the networks; there is also no obvious way to map this metric to simulation. The nature of epoxy network's is an area of interest and ongoing research in our research group, including nanocomposites effects a network's formation and the implications for a materials properties are not fully understood.

### 7.1.2 DeaLAMMPS Development

DeaLAMMPS is a program developed collaboratively within our lab. Previously the majority of the development had be carried out by Maxime Vassaux. Contributors have included Robin Richardson, Alex Patronis and myself. Since Maxime's departure from the lab at the start of 2019, I have taken over as lead developer. Contributions can be tracked on the GitHub repository [6], and details of the implementation are given in Ref. [131].

The developments I have made have been centred on refactoring the code base to make it more reliable, general and parallel. This has been necessary to apply the software to different applications.

The main hurdle for many of DeaLAMMPS potential applications is its enormous computational cost. Imagine a simple dogbone test (a column of epoxy composite strained along its long axis), a reasonable minimum mesh size we would need to hope to capture accurate boundary conditions for a cuboid object would be $3 \times 3 \times 8$ cubic cells. Each cell would have 8 quadrature points (for order 2 polynomial shape functions) and at lease 6 replicas for each simulation [131]. We have already reached 3,500 individual MD simulations per macroscale timestep.

To simulate fracture we may need to induce a strain of 2.5 [130], applying the separation of scales intrinsic to DeaLAMMPS [131], a strain rate of $10^{-11}s^{-1}$, and some hefty simplifications for a rough estimate, each quadrature point will simulate for roughly 10 ps. Each macroscale step will also sample the stress at each point for 100 fs. Therefore, 500 macroscale timesteps, need 50 ps sampling. Simulating 70,000 atoms for 60 ps for a uniaxial strain takes 20 cpuh (using one node on ARCHER (`http://www.archer.ac.uk/`)). Multiplied by the 3,500 separate MD runs required brings the cost to 70,000 cpuh.

For such high strains we would like to use a reactive forcefield, as this would allow us to observe fracture and more complex plastic deformation. A typical penalty for ReaxFF [92] is 50 slower than a standard forcefield like OPLS, raising the cost for a single calculation to 3.5 million cpuh.

These numbers are rough estimates but assume that the finite element step is

negligible in computational cost and that the MD jobs are distributed across the compute resources with perfect efficiency; this of course will never be true. It is also desirable to use longer sampling times and lower strain rates, this would again increase the simulation cost.

Given all of the above, we have tried to reduce the number of MD simulations called at each macroscale timestep. The strain at each quadrature point is described by a two dimensional strain matrix, this is a symmetric matrix so the strain at any point is described by six values. We can then define each quadrature point by its strain history: six series of each unique value from the strain matrix from every previous macroscale timestep. To reduce the number of MD simulations necessary, we fit a spline to each history and use these to calculate a difference threshold. If the difference between two quadrature points is small enough we can then run just one MD simulation for the next step. One can also run fewer simulations if these strains are similar when rotated or inverted; the stress-strain response is trivially identical by applying the same symmetry operator. The validity of these approaches, and the value that this threshold will take, are matters of ongoing research.

This method shows some interesting results such as: a different fracture toughness between a pure MD uniaxial strain and a heterogeneous compact tension test [131]; and strain energy restoration in graphene epoxy nanocomposites, only resolved when the complex shear forces considered in a heterogeneous multiscale model [4]. I will not go into the details of these simulations here as I was not the one who performed those specific runs. These examples are, however, illustrative, that this methodology can furnish new insight into materials that would have been missed by single scale methods.

Testing is underway on reducing the number of MD simulations required for a DeaLAMMPS run. Using a dogbone setup as a test case, we will run one simulation with all quadrature points simulated independently by their own MD representation, then steadily increase the difference threshold described in the previous chapter which compares the similarity of quadrature points' histories. It is unnecessary to explicitly simulate every point in a dogbone test where many points are stretching

uniaxially and the small difference between their local environment is negligible compared to other sources of uncertainty in the model. Instead, we would like to only explicitly simulate areas with significantly different impulses; for example a 'necking' region, a feature that occurs once the rod has started deforming plastically. With the expected reduction in cost, we hope to simulate with reactive forcefields and test more materials, increasing our understanding of how dynamic processes like fracture propagation, hysteresis and toughness are influenced by the material's atomic composition.

## 7.2 Graphene Oxide

With my new model of GO, described in chapter 6, we hope to probe GO in polymer composites. James Suter is developing a coarse grained (CG) model of GO in thermoplastic polymers. So far he has parameterised it with PEG and polyvinyl alcohol. CG-MD allows us to investigate much larger systems: importantly, the polymers can have many more monomer units, increasing the number of entanglements. With CG-MD we can also simulate longer timescales, allowing the study of GO/graphene dispersion/aggregation, and polymer reptation and diffusion. Suter has also developed an algorithm to map a CG representation of a system to an atomistic structure. This allows one to diffuse and equilibrate a large system with CG, then convert the system to its atomic representation to test its mechanical responses.

The new model of GO will be used in future materials tested by DeaLAMMPS. Switching graphene to GO in the current setup will be trivial, but allows us to study many more materials, some of which have much better experimental data to validate against and greater industrial promise.

## 7.3 Epoxy nanosctructure

As part of my research into developing an algorithm to build epoxy polymer networks (described in section 7.1.1) it became apparent that there is no consensus on how to characterise a network on the nanoscale which allows one to predict its macroscale properties. I would like to investigate this with a large ensemble of network polymers to uncover the critical parameters of a network that improve its

mechanical response. If we could know how to improve a network's characteristics in this way, it would be possible to design monomers on the molecular level that give improved material properties.

## Chapter 8

# Conclusions

The research conducted for this doctorate has been wide ranging, encompassing many modelling techniques (both established codes and newly developed ones), multiple materials, forcefield parameterisation, machine learning, equilibrium and non-equilibrium dynamics, and *ab initio*/experimental validation. As such, I have been able to shown that theoretical methods hold much potential in several areas for understanding the properties of graphene and graphene nanocomposites.

I have shown how to simulate graphene in a superlubric state and resolved the dynamics of this unusual behaviour. The process occurs when matching 2D lattices become unaligned, in this incommensurate state there is a near uniform potential energy surface acting between them. Unless a graphene flake is closely aligned to the substrate's lattice, it rotates and translates freely across the surface, encountering nearly zero energy barrier and losing very little energy to the substrate, until it realigns with the substrate. A smooth substrate increases the distance a flake can travel, however, periodic undulations in the surface could be used as a means of separating graphene flakes by size. These findings have applications for graphene as a lubricant and for the design and processing of graphene nanostructures.

I have examined graphene exfoliation in two different scenarios: first, in vacuum pulled by a virtual spring; and second facilitated by polymer adhesives. In chapter 4, I show that peeling, not shearing, has the lower energy barrier; in chapter 5, I show that viscous polymers facilitate graphene exfoliation via a similar mechanism. This insight could aid the design of future graphene exfoliation methods.

Using DFT calculated reaction energy barriers in GO, I have developed a program to build the realistic atomistic structures of GO. It is clear from the analysis presented, that the surface of GO consists of oxidised and pristine graphene phases. This leads to significantly different behaviour compared with an even distribution of oxygen containing functional groups. I investigated the relationship between these two phases and found that the percolation threshold for the oxidised region will, with 95% certainty, not occur when the C/O ratio is greater than 6. This presents a clear goal for experimentalists synthesising GO to preserve the electronic and mechanical properties of graphene.

The difficulty of capturing graphene's enormous aspect ratio can be in part addressed with multiscale simulation. I present a HMM capable of capturing both macro and micro scale dynamics of graphene-epoxy composites, leading to better predictions and a wider scope for simulations. This tool is still in development but will hopefully provide a means of seeing the interaction of nanoscale mechanisms on a materials macroscale properties.

To study a complex system such as nanocomposites with computational models requires accurate descriptions of many components at many disparate length and time scales. I have developed methods to describe graphene's interactions, GO's structure, epoxy polymer networks and hierarchical simulation setups, all of which are open source. These tools, and the understanding that comes with them, will bring closer the ability to predict a material's properties from first principles.

# Appendix A

# Simulation Videos

Below is a summary of the videos contained in the CD accompanying this thesis. The videos are also available to view online, see references [1] and [3].

## A.1 Superlubricity

Video `superlubricity.avi`: Trajectory of a flake pushed over a graphite substrate at 1 K. This is the same trajectory that is shown in figure 4.2a. The flake is 10 nm in diameter; the trajectory's duration is 418 ps.

Video `surfing.avi`: Trajectory of a flake 'surfing' in the trough of a travelling wave in a compressed graphene substrate; see section 4.2.3. The substrate is coloured by its $z$ coordinate: black represents valleys, white represents peaks. Any drift has been removed from the video, i.e. the coordinates have been shifted so the substrate does not move in the $xy$ plane. The flake is 10 nm in diameter; the trajectory's duration is 684 ps.

Video `collision.avi`: Flakes colliding, see section 4.2.4. Drift has been removed form the video as above. Flakes are 10 nm in diameter; the trajectory is 526 ps in duration.

## A.2 Micromechanical Exfoliation

Two videos are provided showing the exfoliation of 7 layer 12 nm diameter graphite by PMMA (video `pmma.avi`) and PDMS (video `pdms.avi`). The videos show the same simulations from which the snapshots are taken in figure 5.3. The periodic

box begins at $20 \times 20 \times 10$ nm$^3$ and the simulation box is increased in the z direction by 10 ms$^{-1}$. The PMMA exfoliation lasts 1.4 ns seconds, the PDMS exfoliation lasts 1.1 ns.

The simulations differ in the mechanism of exfoliation. The PDMS is the more fluid polymer and rearranges to stay in contact with the graphite as the tapes are pulled apart, giving rise to a shearing mechanism. The PMMA does not stay in contact with the graphite for as long, and the outer layers are peeled off from the graphite.

# Appendix B

# Software Developed

## B.1 GraFF

- `github.com/velocirobbie/GraFF`

Files necessary for running a LAMMPS simulation with the GraFF forcefield. `pair_hbond_graphene.h` contains the implementation of the forcefield described in section 3, adapted from a hydrogen bond potential distributed in the vanilla LAMMPS library.

## B.2 DeaLAMMPS

- `github.com/UCL-CCS/DeaLAMMPS`

Coupling library for multiscale simulation using LAMMPS and Deal.II. Currently only implemented on ARCHER. Detailed in section 7.1.2.

## B.3 VECMAtk

- `vecma-toolkit.eu`

During my project I participated in the EU H2020 project VECMA (Verified Exascale Computing for Multiscale Applications), including developing a toolkit for validation, verification and uncertainty quantification (VVUQ) in simulations.

- `github.com/djgroen/FabSim3`

FabSim is a tool I have contributed to that automated submission of jobs to multiple remote computing resources.

- `github.com/UCL-CCS/FabMD`

FabMD is a plugin to FabSim that focuses on tools specifically for molecular dynamics.

## B.4 Make Graphitics

- `github.com/velocirobbie/make-graphitics`

Scripts to generate LAMMPS structure files for various graphitic structures. Structures that can be made include: graphene 'flakes' (hexagonal shaped graphene, edges terminated with hydrogen); graphene 'sheets' (layer of graphene extending through periodic boundaries); graphite; graphene/graphite oxide (oxidised as described in chapter 6).

## B.5 Epoxy Polymerisation

- `github.com/velocirobbie/epoxy_polymerisation`

Program to systematically crosslink polymers. Implementation described in section 7.1.1. Currently only includes inputs for TGMDA cured with PEA; input format would allow for different epoxies to be generated.

## B.6 Graphene Oxide Percolation

- `github.com/velocirobbie/graphene-oxide-percolation`

Graphene oxide percolation model described in chapter 6.

# Appendix C

# Code fragments

## C.1 Concave Hull

Below is a snippet of python code used to calculate the concave hull of an graphene-oxide island, modified from the code found in Ref.[215].

```
def alpha_shape(points, alpha):

    def add_edge(edges, edge_points, coords, i, j):
            if (i, j) in edges or (j, i) in edges:
                return
            edges.add( (i, j) )
            edge_points.append(coords[ [i, j] ])

    tri = Delaunay(points)
    edges = set()
    edge_points = []
    # loop over triangles:
    # ia, ib, ic = indices of corner points of a triangle
    for ia, ib, ic in tri.vertices:
        pa = coords[ia]
        pb = coords[ib]
        pc = coords[ic]
        # Lengths of sides of triangle
        a = math.sqrt((pa[0]-pb[0])**2 + (pa[1]-pb[1])**2)
        b = math.sqrt((pb[0]-pc[0])**2 + (pb[1]-pc[1])**2)
        c = math.sqrt((pc[0]-pa[0])**2 + (pc[1]-pa[1])**2)

        if a > alpha or b > alpha or c > alpha:
            pass
        else:
            add_edge(edges, edge_points, points, ia, ib)
            add_edge(edges, edge_points, points, ib, ic)
            add_edge(edges, edge_points, points, ic, ia)
    m = geometry.MultiLineString(edge_points)
    triangles = list(polygonize(m))
    concave_hull = cascaded_union(triangles), edge_points
    return concave_hull
```

# Appendix D

# Datasets

Here are some of the large datasets created during my thesis that may be useful to the interested reader or future research. The are stored on the CCS network, a private network owned by Peter Coveney's lab. Data is available on request.

## D.1 Superlubricity

- `/net/dirac/mnt/store7/robert/superlubricity`

Simulation trajectories presented in chapter 4.

## D.2 Exfoliation

- `/net/dirac/mnt/store7/robert/exfoliation`

Simulation trajectories presented in chapter 5.

## D.3 Epoxy Polymers

- `/cs/materials/robert/epoxy`

Creation of epoxy and graphene-epoxy simulations, used for DeaLAMMPS simulations presented in chapter 7.1.

## Appendix E

# Colophon

This document was set in the Times Roman typeface using LaTeX and BibTeX, composed within TexShop.

Flipbook functionality provided by `https://github.com/nilsleiffischer/latex-flipbook`.

Figures were generated using the following packages:

- GNUPLOT 5.2 — graphs and schematic plotting
- VMD 1.9.3 — MD simulation visualisation
- GIMP 2.10 — raster graphics editor
- Inkscape 2.0 — vector graphics editor
- `draw.io` — flowchart software
- Chemdoodle — drawing chemical structures

# Bibliography


[1] Robert C Sinclair, James L Suter, and Peter V Coveney. Graphene–graphene interactions: friction, superlubricity, and exfoliation. *Advanced Materials*, 30(13):1705791, 2018.

[2] Robert C. Sinclair and Peter V. Coveney. Modelling nanostructure in graphene oxide: inhomogeneity and the percolation threshold. *Journal of Chemical Information and Modeling*, 59(6):2741, 2019.

[3] Robert C Sinclair, James L Suter, and Peter V Coveney. Micromechanical exfoliation of graphene on the atomistic scale. *Physical Chemistry Chemical Physics*, 21(10):5716–5722, 2019.

[4] Maxime Vassaux, Robert C Sinclair, Robin A Richardson, James L Suter, and Peter V Coveney. The role of graphene in enhancing the material properties of thermosetting polymers. *Advanced Theory and Simulations*, 2(5):1800168, 2019.

[5] Derek Groen, Robin A Richardson, David W Wright, Vytautas Jancauskas, Robert Sinclair, Paul Karlshoefer, Maxime Vassaux, Hamid Arabnejad, Tomasz Piontek, Piotr Kopta, et al. Introducing vecmatk-verification, validation and uncertainty quantification for multiscale and hpc simulations. In *International Conference on Computational Science*, pages 479–492. Springer, 2019.

[6] Robert C. Sinclair, Maxime Vassaux, Robin A. Richardson, and Alex Patronis. DeaLAMMPS repository. https://github.com/UCL-CCS/DeaLAMMPS, 2019. Accessed: 2019-08-29.

[7] Robert C. Sinclair. Epoxy polymerisation code. https://github.com/velocirobbie/epoxy_polymerisation, 2019. Accessed: 2019-08-29.

[8] Robert. C. Sinclair. make-graphitics, 2019. 10.5281/zenodo.2548538.

[9] Robert C. Sinclair. Graphene oxide percolation, 2019. 10.5281/zenodo.2554086.

[10] Robert C. Sinclair. GraFF. 10.5281/zenodo.2421114, 2018.

[11] Kostya S Novoselov, Andre K Geim, Sergei V Morozov, D Jiang, Y_Zhang, Sergey V Dubonos, Irina V Grigorieva, and Alexandr A Firsov. Electric field effect in atomically thin carbon films. *Science*, 306(5696):666–669, 2004.

[12] N David Mermin. Crystalline order in two dimensions. *Physical Review*, 176(1):250, 1968.

[13] Andre K Geim and Konstantin S Novoselov. The rise of graphene. In *Nanoscience and Technology: A Collection of Reviews from Nature Journals*, pages 11–19. World Scientific, 2010.

[14] Changgu Lee, Xiaoding Wei, Jeffrey W Kysar, and James Hone. Measurement of the elastic properties and intrinsic strength of monolayer graphene. *Science*, 321(5887):385–388, 2008.

[15] Alexander A. Balandin, Suchismita Ghosh, Wenzhong Bao, Irene Calizo, Desalegne Teweldebrhan, Feng Miao, and Chun Ning Lau. Superior Thermal Conductivity of Single-Layer Graphene. *Nano Letters*, 8(3):902–907, 2008.

[16] Rahul Raveendran Nair, Peter Blake, Alexander N Grigorenko, Konstantin S Novoselov, Tim J Booth, Tobias Stauber, Nuno MR Peres, and Andre K Geim. Fine structure constant defines visual transparency of graphene. *Science*, 320(5881):1308–1308, 2008.

[17] Kostya S Novoselov, Andre K Geim, SVb Morozov, Da Jiang, Michail I Katsnelson, IVa Grigorieva, and SVb Dubonos. Two-dimensional gas of massless dirac fermions in graphene. *Nature*, 438(7065):197, 2005.

[18] Robert J. Young, Ian A. Kinloch, Lei Gong, and Kostya S. Novoselov. The mechanics of graphene nanocomposites: A review. *Composites Science and Technology*, 72(12):1459–1476, 2012.

[19] Milan Orlita, Clement Faugeras, Paulina Plochocka, Petr Neugebauer, Gerard Martinez, Duncan K Maude, A-L Barra, Mike Sprinkle, Claire Berger, Walter A De Heer, et al. Approaching the dirac point in high-mobility multilayer epitaxial graphene. *Physical Review Letters*, 101(26):267601, 2008.

[20] Xiaozhi Xu, Zhihong Zhang, Jichen Dong, Ding Yi, Jingjing Niu, Muhong Wu, Li Lin, Rongkang Yin, Mingqiang Li, Jingyuan Zhou, et al. Ultrafast epitaxial growth of metre-sized single-crystal graphene on industrial cu foil. *Science bulletin*, 62(15):1074–1080, 2017.

[21] Alain Peigney, Christophe Laurent, Emmanuel Flahaut, Revathi Bacsa, and Abel Rousset. Specific surface area of carbon nanotubes and bundles of carbon nanotubes. *Carbon*, 39(4):507–514, 2001.

[22] Garcia switches to callaway golf balls. `https://blog.lostgolfballs.com/garcia-switches-to-a-new-ball-as-part-of-move-to-callaway-from-taylormade`. Accessed: 2019-06-05.

[23] Materials Genome Initiative. `https://www.mgi.gov/`. Accessed: 2019-06-05.

[24] Ammon J Salter and Ben R Martin. The economic benefits of publicly funded basic research: a critical review. *Research policy*, 30(3):509–532, 2001.

[25] Security tags using graphene based inks. `https://www.graphene-info.com/new-security-tags-built-using-vorbecks-graphene-based-inks-start-shipping-q1-2012`. Accessed: 2019-06-05.

[26] Graphene-assisted high temperature li-ion batteries. `https://www.huawei.com/en/press-events/news/2016/12/Graphene-Assisted-Li-ion-Batteries`. Accessed: 2019-06-05.

[27] Sasha Stankovich, Dmitriy A. Dikin, Geoffrey H. B. Dommett, Kevin M. Kohlhaas, Eric J. Zimney, Eric A. Stach, Richard D. Piner, SonBinh T. Nguyen, and Rodney S. Ruoff. Graphene-based composite materials. *Nature*, 442(7100):282–286, 2006.

[28] Leon Mishnaevsky Jr. and Gaoming Dai. Hybrid and hierarchical nanoreinforced polymer composites: Computational modelling of structure–properties relationships. *Composite Structures*, 117:156–168, 2014.

[29] Hyunwoo Kim and Christopher W. Macosko. Processing-property relationships of polycarbonate/graphene composites. *Polymer*, 50(15):3797–3809, 2009.

[30] Dimitrios G Papageorgiou, Ian A Kinloch, and Robert J Young. Mechanical properties of graphene and graphene-based nanocomposites. *Progress in Materials Science*, 90:75–127, 2017.

[31] Dimitrios G. Papageorgiou, Ian A. Kinloch, and Robert J. Young. Graphene/elastomer nanocomposites. *Carbon*, 95:460–484, 2015.

[32] Jeffrey R. Potts, Daniel R. Dreyer, Christopher W. Bielawski, and Rodney S. Ruoff. Graphene-based polymer nanocomposites. *Polymer*, 52(1):5–25, 2011.

[33] Biqiong Chen, Julian R. G. Evans, H. Christopher Greenwell, Pascal Boulet, Peter V. Coveney, Allen A. Bowden, and Andrew Whiting. A critical appraisal of polymer-clay nanocomposites. *Chemical Society Reviews*, 37(3):568–594, 2008.

[34] H. Chris Greenwell, William Jones, Peter V. Coveney, and Stephen Stackhouse. On the application of computer simulation techniques to anionic and cationic clays: A materials chemistry perspective. *Journal of Materials Chemistry*, 16(8):708–723, 2006.

[35] Youngkun Ahn, Jungyoon Kim, Shraddha Ganorkar, Young-Hwan Kim, and Seong-Il Kim. Thermal annealing of graphene to remove polymer residues. *Materials Express*, 6(1):69–76, 2016.

[36] Yi Zhang, Luyao Zhang, and Chongwu Zhou. Review of Chemical Vapor Deposition of Graphene and Related Applications. *Accounts of Chemical Research*, 46(10):2329–2339, 2013.

[37] Qingkai Yu, Jie Lian, Sujitra Siriponglert, Hao Li, Yong P. Chen, and Shin-Shem Pei. Graphene segregated on Ni surfaces and transferred to insulators. *Applied Physics Letters*, 93(11):113103, 2008.

[38] Xuesong Li, Weiwei Cai, Jinho An, Seyoung Kim, Junghyo Nah, Dongxing Yang, Richard Piner, Aruna Velamakanni, Inhwa Jung, Emanuel Tutuc, Sanjay K. Banerjee, Luigi Colombo, and Rodney S. Ruoff. Large-Area Synthesis of High-Quality and Uniform Graphene Films on Copper Foils. *Science*, 324(5932):1312–1314, 2009.

[39] Ji Won Suk, Alexander Kitt, Carl W. Magnuson, Yufeng Hao, Samir Ahmed, Jinho An, Anna K. Swan, Bennett B. Goldberg, and Rodney S. Ruoff. Transfer of CVD-Grown Monolayer Graphene onto Arbitrary Substrates. *ACS Nano*, 5(9):6916–6924, 2011.

[40] Keith R Paton, Eswaraiah Varrla, Claudia Backes, Ronan J Smith, Umar Khan, Arlene O'Neill, Conor Boland, Mustafa Lotya, Oana M Istrate, Paul

King, et al. Scalable production of large quantities of defect-free few-layer graphene by shear exfoliation in liquids. *Nature Materials*, 13(6):624, 2014.

[41] Jonathan N. Coleman. Liquid-Phase Exfoliation of Nanotubes and Graphene. *Advanced Functional Materials*, 19(23):3680–3695, 2009.

[42] Minzhen Cai, Daniel Thorpe, Douglas H. Adamson, and Hannes C. Schniepp. Methods of graphite exfoliation. *Journal of Materials Chemistry*, 22(48):24992–25002, 2012.

[43] Umar Khan, Arlene O'Neill, Mustafa Lotya, Sukanta De, and Jonathan N. Coleman. High-Concentration Solvent Exfoliation of Graphene. *Small*, 6(7):864–871, 2010.

[44] Deepthi Konatham and Alberto Striolo. Molecular Design of Stable Graphene Nanosheets Dispersions. *Nano Letters*, 8(12):4630–4641, 2008.

[45] Cheng Lv, Qingzhong Xue, Dan Xia, Ming Ma, Jie Xie, and Huijuan Chen. Effect of Chemisorption on the Interfacial Bonding Characteristics of Graphene−Polymer Composites. *The Journal of Physical Chemistry C*, 114(14):6588–6594, 2010.

[46] Xiaolin Li, Guangyu Zhang, Xuedong Bai, Xiaoming Sun, Xinran Wang, Enge Wang, and Hongjie Dai. Highly conducting graphene sheets and Langmuir–Blodgett films. *Nature Nanotechnology*, 3(9):538–542, 2008.

[47] Jens Baringhaus, Ming Ruan, Frederik Edler, Antonio Tejeda, Muriel Sicot, Amina Taleb-Ibrahimi, An-Ping Li, Zhigang Jiang, Edward H Conrad, Claire Berger, et al. Exceptional ballistic transport in epitaxial graphene nanoribbons. *Nature*, 506(7488):349, 2014.

[48] Yuan Cao, Valla Fatemi, Shiang Fang, Kenji Watanabe, Takashi Taniguchi, Efthimios Kaxiras, and Pablo Jarillo-Herrero. Unconventional superconductivity in magic-angle graphene superlattices. *Nature*, 556(7699):43, 2018.

[49] Yuan Cao, Valla Fatemi, Ahmet Demir, Shiang Fang, Spencer L Tomarken, Jason Y Luo, Javier D Sanchez-Yamagishi, Kenji Watanabe, Takashi Taniguchi, Efthimios Kaxiras, et al. Correlated insulator behaviour at half-filling in magic-angle graphene superlattices. *Nature*, 556(7699):80, 2018.

[50] Elizabeth Gibney. Howmagic angle'graphene is stirring up physics. *Nature*, 565:15–18, 2019.

[51] Matthew Yankowitz, Shaowen Chen, Hryhoriy Polshyn, Yuxuan Zhang, K Watanabe, T Taniguchi, David Graf, Andrea F Young, and Cory R Dean. Tuning superconductivity in twisted bilayer graphene. *Science*, 363(6431):1059–1064, 2019.

[52] E Suárez Morell, JD Correa, P Vargas, M Pacheco, and Z Barticevic. Flat bands in slightly twisted bilayer graphene: Tight-binding calculations. *Physical Review B*, 82(12):121407, 2010.

[53] B. C. Brodie. On the Atomic Weight of Graphite. *Philosophical Transactions of the Royal Society of London*, 149:249–259, 1859.

[54] Kris Erickson, Rolf Erni, Zonghoon Lee, Nasim Alem, Will Gannett, and Alex Zettl. Determination of the Local Chemical Structure of Graphene Oxide and Reduced Graphene Oxide. *Advanced Materials*, 22(40):4467–4472, 2010.

[55] Anton Lerf, Heyong He, Michael Forster, and Jacek Klinowski. Structure of Graphite Oxide Revisited. *The Journal of Physical Chemistry B*, 102(23):4477–4482, 1998.

[56] Wei Gao, Lawrence B. Alemany, Lijie Ci, and Pulickel M. Ajayan. New insights into the structure and reduction of graphite oxide. *Nature Chemistry*, 1(5):403–408, 2009.

[57] K. Andre Mkhoyan, Alexander W. Contryman, John Silcox, Derek A. Stewart, Goki Eda, Cecilia Mattevi, Steve Miller, and Manish Chhowalla. Atomic

and Electronic Structure of Graphene-Oxide. *Nano Letters*, 9(3):1058–1063, 2009.

[58] Jinrong Yang, Guosheng Shi, Yusong Tu, and Haiping Fang. High Correlation between Oxidation Loci on Graphene Oxide. *Angewandte Chemie International Edition*, 53(38):10190–10194, 2014.

[59] Karthikeyan Krishnamoorthy, Murugan Veerapandian, Kyusik Yun, and S-J Kim. The chemical and structural analysis of graphene oxide with different degrees of oxidation. *Carbon*, 53:38–49, 2013.

[60] Peter Wick, Anna E. Louw-Gaume, Melanie Kucki, Harald F. Krug, Kostas Kostarelos, Bengt Fadeel, Kenneth A. Dawson, Anna Salvati, Ester Vázquez, Laura Ballerini, Mauro Tretiach, Fabio Benfenati, Emmanuel Flahaut, Laury Gauthier, Maurizio Prato, and Alberto Bianco. Classification Framework for Graphene-Based Materials. *Angewandte Chemie International Edition*, 53(30):7714–7718, 2014.

[61] William S. Hummers and Richard E. Offeman. Preparation of Graphitic Oxide. *Journal of the American Chemical Society*, 80(6):1339–1339, 1958.

[62] Cristina Gómez-Navarro, Marko Burghard, and Klaus Kern. Elastic Properties of Chemically Derived Single Graphene Sheets. *Nano Letters*, 8(7):2045–2049, 2008.

[63] Dmitriy A. Dikin, Sasha Stankovich, Eric J. Zimney, Richard D. Piner, Geoffrey H. B. Dommett, Guennadi Evmenenko, SonBinh T. Nguyen, and Rodney S. Ruoff. Preparation and characterization of graphene oxide paper. *Nature*, 448(7152):457–460, 2007.

[64] Sasha Stankovich, Dmitriy A. Dikin, Richard D. Piner, Kevin A. Kohlhaas, Alfred Kleinhammes, Yuanyuan Jia, Yue Wu, SonBinh T. Nguyen, and Rodney S. Ruoff. Synthesis of graphene-based nanosheets via chemical reduction of exfoliated graphite oxide. *Carbon*, 45(7):1558–1565, 2007.

[65] Jonathan P. Rourke, Priyanka A. Pandey, Joseph J. Moore, Matthew Bates, Ian A. Kinloch, Robert J. Young, and Neil R. Wilson. The Real Graphene Oxide Revealed: Stripping the Oxidative Debris from the Graphene-like Sheets. *Angewandte Chemie*, 123(14):3231–3235, 2011.

[66] Alberto Bianco, Hui-Ming Cheng, Toshiaki Enoki, Yury Gogotsi, Robert H. Hurt, Nikhil Koratkar, Takashi Kyotani, Marc Monthioux, Chong Rae Park, Juan M. D. Tascon, and Jin Zhang. All in the graphene family – A recommended nomenclature for two-dimensional carbon materials. *Carbon*, 65:1–6, 2013.

[67] John P. A. Ioannidis. Why Most Published Research Findings Are False. *PLOS Medicine*, 2(8):e124, 2005.

[68] Composite industry overview. `https://acmanet.org/composites-industry-overview/`. Accessed: 2019-06-06.

[69] Sarah O'Donnell, Kevin Sprong, and Brennan Haltli. Potential impact of carbon nanotube reinforced polymer composite on commercial aircraft performance and economics. In *AIAA 4th Aviation Technology, Integration and Operations (ATIO) Forum*, page 6402, 2004.

[70] Ivan Vlassiouk, Georgios Polizos, Ryan Cooper, Ilia Ivanov, Jong Kahk Keum, Felix Paulauskas, Panos Datskos, and Sergei Smirnov. Strong and electrically conductive graphene-based composite fibers and laminates. *ACS Applied Materials & Interfaces*, 7(20):10702–10709, 2015.

[71] Lei Gong, Ian A. Kinloch, Robert J. Young, Ibtsam Riaz, Rashid Jalil, and Konstantin S. Novoselov. Interfacial Stress Transfer in a Graphene Monolayer Nanocomposite. *Advanced Materials*, 22(24):2694–2697, 2010.

[72] Cristina Vallés, Amr M. Abdelkader, Robert J. Young, and Ian A. Kinloch. Few layer graphene–polypropylene nanocomposites: the role of flake diameter. *Faraday Discuss.*, 2014.

[73] Lei Gong, Robert J Young, Ian A Kinloch, Ibtsam Riaz, Rashid Jalil, and Kostya S Novoselov. Optimizing the reinforcement of polymer-based nanocomposites by graphene. *ACS nano*, 6(3):2086–2095, 2012.

[74] Hao-Bin Zhang, Wen-Ge Zheng, Qing Yan, Yong Yang, Ji-Wen Wang, Zhao-Hui Lu, Guo-Ying Ji, and Zhong-Zhen Yu. Electrically conductive polyethylene terephthalate/graphene nanocomposites prepared by melt compounding. *Polymer*, 51(5):1191–1196, 2010.

[75] Chun-Teh Chen and Grace X Gu. Effect of constituent materials on composite performance: Exploring design strategies via machine learning. *Advanced Theory and Simulations*, 2019.

[76] Zheling Li, Robert J. Young, Neil R. Wilson, Ian A. Kinloch, Cristina Vallés, and Zheng Li. Effect of the orientation of graphene-based nanoplatelets upon the Young's modulus of nanocomposites. *Composites Science and Technology*, 123:125–133, 2016.

[77] Gad Marom and H. Daniel Wagner. Should polymer nanocomposites be regarded as molecular composites? *Journal of Materials Science*, 52(14):8357–8361, 2017.

[78] Robert J. Young, Mufeng Liu, Ian A. Kinloch, Suhao Li, Xin Zhao, Cristina Vallés, and Dimitrios G. Papageorgiou. The mechanics of reinforcement of polymers by graphene nanoplatelets. *Composites Science and Technology*, 154:110–116, 2018.

[79] Zheling Li, Robert J Young, Ian A Kinloch, Neil R Wilson, Alexander J Marsden, and Arun Prakash Aranga Raju. Quantitative determination of the spatial orientation of graphene by polarized raman spectroscopy. *Carbon*, 88:215–224, 2015.

[80] Bin Shen, Wentao Zhai, Cao Chen, Dingding Lu, Jing Wang, and Wenge Zheng. Melt Blending In situ Enhances the Interaction between Polystyrene

and Graphene through pi-pi Stacking. *ACS Applied Material Interfaces*, 3(8):3103–3109, 2011.

[81] Cristina Vallés, Fabian Beckert, Laura Burk, Rolf Mülhaupt, Robert J. Young, and Ian A. Kinloch. Effect of the C/O ratio in graphene oxide materials on the reinforcement of epoxy-based nanocomposites. *Journal of Polymer Science Part B*, 54(2):281–291, 2016.

[82] Emmanuel N. Skountzos, Alexandros Anastassiou, Vlasis G. Mavrantzas, and Doros N. Theodorou. Determination of the Mechanical Properties of a Poly(methyl methacrylate) Nanocomposite with Functionalized Graphene Sheets through Detailed Atomistic Simulations. *Macromolecules*, 47(22):8072–8088, 2014.

[83] Jaewoon Jung, Wataru Nishima, Marcus Daniels, Gavin Bascom, Chigusa Kobayashi, Adetokunbo Adedoyin, Michael Wall, Anna Lappala, Dominic Phillips, William Fischer, et al. Scaling molecular dynamics beyond 100,000 processor cores for large-scale biophysical simulations. *Journal of Computational Chemistry*, 2019.

[84] Alfons G Hoekstra, Bastien Chopard, David Coster, Simon Portegies Zwart, and Peter V Coveney. Multiscale computing for science and engineering in the era of exascale performance. *Philosophical Transactions of the Royal Society A*, 377(2142):20180144, 2019.

[85] Guillaume Houzeaux, Marta Garcia, Juan Carlos Cajas, Antoni Artigues, Edgar Olivares, Jesús Labarta, and Mariano Vázquez. Dynamic load balance applied to particle transport in fluids. *International Journal of Computational Fluid Dynamics*, 30(6):408–418, 2016.

[86] M Pütz, K Kremer, and G. S Grest. What is the entanglement length in a polymer melt? *Europhysics Letters*, 49(6):735–741, 2000.

[87] Maxime Vassaux, Robert C Sinclair, Robin A Richardson, James L Suter, and Peter V Coveney. Toward high fidelity materials property prediction

from multiscale modeling and simulation. *Advanced Theory and Simulations*, 3(1):1900122, 2020.

[88] Martin Pykal, Petr Jurečka, František Karlický, and Michal Otyepka. Modelling of graphene functionalization. *Physical Chemistry Chemical Physics*, 2015.

[89] Vivekanand V. Gobre and Alexandre Tkatchenko. Scaling laws for van der Waals interactions in nanostructured materials. *Nature Communications*, 4:2341, 2013.

[90] Annalisa Fasolino, JH Los, and Mikhail I Katsnelson. Intrinsic ripples in graphene. *Nature Materials*, 6(11):858, 2007.

[91] Mikuláš Kocman, Martin Pykal, and Petr Jurečka. Electric quadrupole moment of graphene and its effect on intermolecular interactions. *Physical Chemistry Chemical Physics*, 16(7):3144–3152, 2014.

[92] Adri CT Van Duin, Siddharth Dasgupta, Francois Lorant, and William A Goddard. Reaxff: a reactive force field for hydrocarbons. *The Journal of Physical Chemistry A*, 105(41):9396–9409, 2001.

[93] Steven J Stuart, Alan B Tutein, and Judith A Harrison. A reactive potential for hydrocarbons with intermolecular interactions. *The Journal of Chemical Physics*, 112(14):6472–6486, 2000.

[94] Zhaohui Wang, Sverre M. Selbach, and Tor Grande. Van der Waals density functional study of the energetics of alkali metal intercalation in graphite. *RSC Advances*, 4(8):3973–3983, 2014.

[95] Chun Hung Lui and Tony F Heinz. Measurement of layer breathing mode vibrations in few-layer graphene. *Physical Review B*, 87(12):121404, 2013.

[96] E Mostaani, ND Drummond, and VI Fal'Ko. Quantum monte carlo calculation of the binding energy of bilayer graphene. *Physical review letters*, 115(11):115501, 2015.

[97] John F. Dobson, Tim Gould, and Giovanni Vignale. How Many-Body Effects Modify the van der Waals Interaction between Graphene Sheets. *Physical Reviews X*, 4(2):021040, 2014.

[98] Chih-Jen Shih, Shangchao Lin, Michael S. Strano, and Daniel Blankschtein. Understanding the Stabilization of Liquid-Phase-Exfoliated Graphene in Polar Solvents: Molecular Dynamics Simulations and Kinetic Theory of Colloid Aggregation. *Journal of the American Chemical Society*, 132(41):14638–14648, 2010.

[99] Cameron M Hadden, Danielle René Klimek-McDonald, Evan J Pineda, Julia A King, Alex M Reichanadter, Ibrahim Miskioglu, S Gowtham, and Gregory M Odegard. Mechanical properties of graphene nanoplatelet/carbon fiber/epoxy hybrid composites: Multiscale modeling and experiments. *Carbon*, 95:100–112, 2015.

[100] Cuili Fu and Xiaoning Yang. Molecular simulation of interfacial mechanics for solvent exfoliation of graphene from graphite. *Carbon*, 55:350–360, 2013.

[101] Haoyue Sun and Xiaoning Yang. Molecular simulation of self-assembly structure and interfacial interaction for SDBS adsorption on graphene. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 462:82–89, 2014.

[102] Jie Yang, Xiaoning Yang, and Yaping Li. Molecular simulation perspective of liquid-phase exfoliation, dispersion, and stabilization for graphene. *Current Opinion in Colloid & Interface Science*, 20(5-6):339–345, 2015.

[103] Talal Al-Mulla, Zhao Qin, and Markus J. Buehler. Crumpling deformation regimes of monolayer graphene on substrate: a molecular mechanics study. *Journal of Physics: Condensed Matter*, 27(34):345401, 2015.

[104] A Lohrasebi, M Neek-Amal, and MR Ejtehadi. Directed motion of $C_{60}$ on a graphene sheet subjected to a temperature gradient. *Physical Review E*, 83(4):042601, 2011.

[105] Rafał Podeszwa. Interactions of graphene sheets deduced from properties of polycyclic aromatic hydrocarbons. *The Journal of Chemical Physics*, 132(4):044704, 2010.

[106] Yan Zhao and Donald G. Truhlar. A Prototype for Graphene Material Simulation: Structures and Interaction Potentials of Coronene Dimers. *Journal of Physical Chemistry C*, 112(11):4061–4067, 2008.

[107] Nikhil V. Medhekar, Ashwin Ramasubramaniam, Rodney S. Ruoff, and Vivek B. Shenoy. Hydrogen Bond Networks in Graphene Oxide Composite Paper: Structure and Mechanical Properties. *ACS Nano*, 4(4):2300–2306, 2010.

[108] Qing-Xiang Pei, Yong-Wei Zhang, and Vivek B. Shenoy. A molecular dynamics study of the mechanical properties of hydrogen functionalized graphene. *Carbon*, 48(3):898–904, 2010.

[109] Teng Zhang, Xiaoyan Li, Sara Kadkhodaei, and Huajian Gao. Flaw Insensitive Fracture in Nanocrystalline Graphene. *Nano Letters*, 12(9):4605–4610, 2012.

[110] GangSeob Jung, Zhao Qin, and Markus J Buehler. Molecular mechanics of polycrystalline graphene with enhanced fracture toughness. *Extreme Mechanics Letters*, 2:52–59, 2015.

[111] Ahmed R. Alian, Nuwan Dewapriya, and Shaker A. Meguid. Molecular dynamics study of the reinforcement effect of graphene in multilayered polymer nanocomposites. *Materials & Design*, 124:47–57, 2017.

[112] Feng Liu, Ning Hu, Huiming Ning, Yaolu Liu, Yuan Li, and Liangke Wu. Molecular dynamics simulation on interfacial mechanical properties of poly-

mer nanocomposites with wrinkled graphene. *Computational Materials Science*, 108:160–167, 2015.

[113] Songul Güryel, Martin Walker, P Geerlings, F De Proft, and Mark R. Wilson. Molecular dynamics simulations of the structure and the morphology of graphene/polymer nanocomposites. *Physical Chemistry Chemical Physics*, 19(20):12959–12969, 2017.

[114] Anastassia Rissanou, Albert Power, and Vagelis Harmandaris. Structural and dynamical properties of polyethylene/graphene nanocomposites through molecular dynamics simulations. *Polymers*, 7(3):390–417, 2015.

[115] Jun Liu, Jianxiang Shen, Zijian Zheng, Youping Wu, and Liqun Zhang. Revealing the toughening mechanism of graphene–polymer nanocomposite through molecular dynamics simulation. *Nanotechnology*, 26(29):291003, 2015.

[116] Chang-Tsan Lu, Asanka Weerasinghe, Dimitrios Maroudas, and Ashwin Ramasubramaniam. A comparison of the elastic properties of graphene-and fullerene-reinforced polymer composites: the role of filler morphology and size. *Scientific Reports*, 6:31735, 2016.

[117] Doros N. Theodorou and Ulrich W. Suter. Atomistic modeling of mechanical properties of polymeric glasses. *Macromolecules*, 19(1):139–154, 1986.

[118] Yunlong Li, Shijie Wang, and Quan Wang. A molecular dynamics simulation study on enhancement of mechanical and tribological properties of polymer composites by introduction of graphene. *Carbon*, 111:538–545, 2017.

[119] Çağlar Ö Girit, Jannik C. Meyer, Rolf Erni, Marta D. Rossell, C. Kisielowski, Li Yang, Cheol-Hwan Park, M. F. Crommie, Marvin L. Cohen, Steven G. Louie, and A. Zettl. Graphene at the Edge: Stability and Dynamics. *Science*, 323(5922):1705–1708, 2009.

[120] Zheng Liu, Kazu Suenaga, Peter J. F. Harris, and Sumio Iijima. Open and Closed Edges of Graphene Layers. *Physical Review Letters*, 102(1):015501, 2009.

[121] Karen Johnston and Vagelis Harmandaris. Hierarchical simulations of hybrid polymer–solid materials. *Soft Matter*, 9(29):6696–6710, 2013.

[122] Luis Ruiz, Wenjie Xia, Zhaoxu Meng, and Sinan Keten. A coarse-grained model for the mechanical behavior of multi-layer graphene. *Carbon*, 82:103–115, 2015.

[123] Hossein Eslami, Hossein Ali Karimi-Varzaneh, and Florian Müller-Plathe. Coarse-Grained Computer Simulation of Nanoconfined Polyamide-6,6. *Macromolecules*, 44(8):3117–3128, 2011.

[124] Feng Lin, Yang Xiang, and Hui-Shen Shen. Temperature dependent mechanical properties of graphene reinforced polymer nanocomposites–a molecular dynamics simulation. *Composites Part B: Engineering*, 111:261–269, 2017.

[125] Janghyuk Moon, Seunghwa Yang, and Maenghyo Cho. Interfacial strengthening between graphene and polymer through stone-thrower-wales defects: Ab initio and molecular dynamics simulations. *Carbon*, 118:66–77, 2017.

[126] Chanwook Park and Gun Jin Yun. Characterization of interfacial properties of graphene-reinforced polymer nanocomposites by molecular dynamics-shear deformation model. *Journal of Applied Mechanics*, 85(9):091007, 2018.

[127] Anastassia N. Rissanou and Vagelis Harmandaris. A Molecular Dynamics Study of Polymer/Graphene Nanocomposites. *Macromolecular Symposia*, 331-332(1):43–49, 2013.

[128] Anastassia N. Rissanou and Vagelis Harmandaris. Structure and dynamics of poly(methyl methacrylate)/graphene systems through atomistic molecular dynamics simulations. *Journal of Nanoparticle Research*, 15(5):1–14, 2013.

[129] James L. Suter, Derek Groen, and Peter V. Coveney. Chemically Specific Multiscale Modeling of Clay–Polymer Nanocomposites Reveals Intercalation Dynamics, Tactoid Self-Assembly and Emergent Materials Properties. *Advanced Materials*, 27(6):966–984, 2015.

[130] Zhaoxu Meng, Miguel A Bessa, Wenjie Xia, Wing Kam Liu, and Sinan Keten. Predicting the macroscopic fracture energy of epoxy resins from atomistic molecular simulations. *Macromolecules*, 49(24):9474–9483, 2016.

[131] Maxime Vassaux, Robin A. Richardson, and Peter V. Coveney. The heterogeneous multiscale method applied to inelastic polymer mechanics. *Philosophical Transactions of the Royal Society A*, 377(2142):20180150, 2019.

[132] Konstantinos N. Spanos, Stelios K. Georgantzinos, and Nick K. Anifantis. Mechanical properties of graphene nanocomposites: a multiscale finite element prediction. *Composite Structures*, 132:536–544, 2015.

[133] Chih-Jen Shih, Michael S. Strano, and Daniel Blankschtein. Wetting translucency of graphene. *Nature Materials*, 12(10):866–869, 2013.

[134] Yi Lin and John W Connell. Advances in 2d boron nitride nanostructures: nanosheets, nanoribbons, nanomeshes, and hybrids with graphene. *Nanoscale*, 4(22):6908–6939, 2012.

[135] Michael Naguib, Vadym N Mochalin, Michel W Barsoum, and Yury Gogotsi. 25th anniversary article: Mxenes: a new family of two-dimensional materials. *Advanced Materials*, 26(7):992–1005, 2014.

[136] Michael Naguib, Olha Mashtalir, Joshua Carle, Volker Presser, Jun Lu, Lars Hultman, Yury Gogotsi, and Michel W Barsoum. Two-dimensional transition metal carbides. *ACS nano*, 6(2):1322–1331, 2012.

[137] Michael Naguib, Murat Kurtoglu, Volker Presser, Jun Lu, Junjie Niu, Min Heon, Lars Hultman, Yury Gogotsi, and Michel W Barsoum. Two-

dimensional nanocrystals produced by exfoliation of ti3alc2. *Advanced Materials*, 23(37):4248–4253, 2011.

[138] Steve Plimpton. Grapheneviasonication assisted liquid-phase exfoliation. *Journal of Computational Physics*, 43(1):381–398, 2014.

[139] D.A. Case, R.M. Betz, W. Botello-Smith, D.S. Cerutti, T.E. Cheatham, III, T.A. Darden, R.E. Duke, T.J. Giese, H. Gohlke, A.W. Goetz, N. Homeyer, S. Izadi, P. Janowski, J. Kaus, A. Kovalenko, T.S. Lee, S. LeGrand, P. Li, C. Lin, T. Luchko, R. Luo, B. Madej, D. Mermelstein, K.M. Merz, G. Monard, H. Nguyen, H.T. Nguyen, I. Omelyan, A. Onufriev, D.R. Roe, A. Roitberg, C. Sagui, C.L. Simmerling, J. Swails, R.C. Walker, J. Wang, R.M. Wolf, X. Wu, L. Xiao, D.M. York, and P.A. Kollman. *AMBER 2016*. University of California and San Francisco, 2016.

[140] William L. Jorgensen, David S. Maxwell, and Julian Tirado-Rives. Development and Testing of the OPLS All-Atom Force Field on Conformational Energetics and Properties of Organic Liquids. *Journal of the American Chemical Society*, 118(45):11225–11236, 1996.

[141] Stephen L. Mayo, Barry D. Olafson, and William A. Goddard. DREIDING: a generic force field for molecular simulations. *The Journal of Physical Chemistry*, 94(26):8897–8909, 1990.

[142] Huai Sun. COMPASS: An ab Initio Force-Field Optimized for Condensed-Phase ApplicationsOverview with Details on Alkane and Benzene Compounds. *The Journal of Physical Chemistry B*, 102(38):7338–7364, 1998.

[143] Vitaly V Chaban, Eudes Eterno Fileti, and Oleg V Prezhdo. Exfoliation of graphene in ionic liquids: pyridinium versus pyrrolidinium. *The Journal of Physical Chemistry C*, 121(1):911–917, 2017.

[144] Martin Fitzner, Laurent Joly, Ming Ma, Gabriele C Sosso, Andrea Zen, and Angelos Michaelides. Communication: Truncated non-bonded potentials can

yield unphysical behavior in molecular dynamics simulations of interfaces. *The Journal of Chemical Physics*, 147(12):121102, 2017.

[145] Rodger W. Hockney and James W. Eastwood. *Computer Simulation Using Particles*. Adam Hilger, NY, 1989.

[146] Shūichi Nosé. A molecular dynamics method for simulations in the canonical ensemble. *Molecular Physics*, 52(2):255–268, 1984.

[147] William G. Hoover. Canonical dynamics: Equilibrium phase-space distributions. *Physical Reviews A*, 31(3):1695–1697, 1985.

[148] E Polak and G Ribiere. Note sur la convergence de directions conjugées. rev. francaise informat. *Recherche Opertionelle, 3e année*, 16:35–43, 1969.

[149] Denis J Evans and Brad Lee Holian. The nose–hoover thermostat. *The Journal of chemical physics*, 83(8):4069–4074, 1985.

[150] Wang Gao and Alexandre Tkatchenko. Sliding Mechanisms in Multilayered Hexagonal Boron Nitride and Graphene: The Effects of Directionality, Thickness, and Sliding Constraints. *Physical Review Letters*, 114(9):096101, 2015.

[151] Jiří Klimeš, David R Bowler, and Angelos Michaelides. A critical assessment of theoretical methods for finding reaction pathways and transition states of surface processes. *Journal of Physics: Condensed Matter*, 22(7):074203, 2010.

[152] Kyuho Lee, Éamonn D Murray, Lingzhu Kong, Bengt I Lundqvist, and David C Langreth. Higher-accuracy van der waals density functional. *Physical Review B*, 82(8):081101, 2010.

[153] Wolfgang Bangerth, Ralf Hartmann, and Guido Kanschat. deal. ii—a general-purpose object-oriented finite element library. *ACM Transactions on Mathematical Software (TOMS)*, 33(4):24, 2007.

[154] Marcus D Hanwell, Donald E Curtis, David C Lonie, Tim Vandermeersch, Eva Zurek, Geoffrey R Hutchison, Avogadro: An advanced semantic chemical editor visualization, and analysis platform". Avogadro: an advanced semantic chemical editor, visualization, and analysis platform. *Journal of Cheminformatics*, 4(1):17, 2012.

[155] William Humphrey, Andrew Dalke, and Klaus Schulten. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics*, 14:33–38, 1996.

[156] Bradley Efron and Robert J Tibshirani. *An introduction to the bootstrap*. CRC press, 1994.

[157] James L. Suter, Derek Groen, and Peter V. Coveney. Mechanism of Exfoliation and Prediction of Materials Properties of Clay–Polymer Nanocomposites from Multiscale Modeling. *Nano Letters*, 15, 2015.

[158] Renju Zacharia, Hendrik Ulbricht, and Tobias Hertel. Interlayer cohesive energy of graphite from thermal desorption of polyaromatic hydrocarbons. *Physical Review B*, 69(15):155406, 2004.

[159] Ze Liu, Jefferson Zhe Liu, Yao Cheng, Zhihong Li, Li Wang, and Quanshui Zheng. Interlayer binding energy of graphite: A mesoscopic determination from deformation. *Physical Review B*, 85(20):205418, 2012.

[160] Lorin X Benedict, Nasreen G Chopra, Marvin L Cohen, A Zettl, Steven G Louie, and Vincent H Crespi. Microscopic determination of the interlayer binding energy in graphite. *Chemical Physics Letters*, 286(5–6):490–496, 1998.

[161] Xiaofeng Feng, Sangku Kwon, Jeong Young Park, and Miquel Salmeron. Superlubric Sliding of Graphene Nanoflakes on Graphene. *ACS Nano*, 7(2):1718–1724, 2013.

[162] Alberto Ambrosetti, Anthony M. Reilly, Robert A. DiStasio, and Alexandre Tkatchenko. Long-range correlation energy calculated from coupled atomic

response functions. *The Journal of Chemical Physics*, 140(18):18A508, 2014.

[163] Alexandre Tkatchenko, Robert A. DiStasio, Roberto Car, and Matthias Scheffler. Accurate and Efficient Method for Many-Body van der Waals Interactions. *Physical Review Letters*, 108(23):236402, 2012.

[164] Tod A Pascal. LAMMPS Dreiding H-bond potential, 2003. [Online; accessed 29/08/2019].

[165] Zhaohui Wang, Sverre M. Selbach, and Tor Grande. Van der Waals density functional study of the energetics of alkali metal intercalation in graphite. *RSC Advances*, 4(8):3973–3983, 2014.

[166] Paul J. Bryant, Paul L. Gutshall, and Lyle H. Taylor. A study of mechanisms of graphite friction and wear. *Wear*, 7(1):118–126, 1964.

[167] Elad Koren, Emanuel Lörtscher, Colin Rawlings, Armin W Knoll, and Urs Duerig. Adhesion and friction in mesoscopic graphite contacts. *Science*, 348(6235):679–683, 2015.

[168] Suzhi Li, Qunyang Li, Robert W Carpick, Peter Gumbsch, Xin Z Liu, Xiangdong Ding, Jun Sun, and Ju Li. The evolving quality of frictional contact with graphene. *Nature*, 539(7630):541, 2016.

[169] Roberto Guerra, Ugo Tartaglino, Andrea Vanossi, and Erio Tosatti. Ballistic nanofriction. *Nature Materials*, 9(8):634, 2010.

[170] Irina V Lebedeva, Andrey A Knizhnik, Andrey M Popov, Olga V Ershova, Yurii E Lozovik, and Boris V Potapkin. Fast diffusion of a graphene flake on a graphene layer. *Physical Review B*, 82(15):155460, 2010.

[171] Chih-Jen Shih, Shangchao Lin, Michael S. Strano, and Daniel Blankschtein. Understanding the Stabilization of Liquid-Phase-Exfoliated Graphene in Polar Solvents: Molecular Dynamics Simulations and Kinetic Theory of Col-

loid Aggregation. *Journal of the American Chemical Society*, 132(41):14638, 2010.

[172] John H. Gaddum. Lognormal distributions. *Nature*, 156(3964):463, 1945.

[173] Jannik C Meyer, A K Geim, M I Katsnelson, K S Novoselov, T J Booth, and S Roth. The structure of suspended graphene sheets. *Nature*, 446(7131):60, 2007.

[174] Ming Ma, Gabriele Tocci, Angelos Michaelides, and Gabriel Aeppli. Fast diffusion of water nanodroplets on graphene. *Nature Materials*, 15(1):66, 2016.

[175] Sang Jin Kim, Teajun Choi, Bora Lee, Sunwoo Lee, Kyoungjun Choi, Jong Bo Park, Je Min Yoo, Yong Seok Choi, Jaechul Ryu, Philip Kim, et al. Ultraclean patterned transfer of single-layer graphene by recyclable pressure sensitive adhesive films. *Nano letters*, 15(5):3236, 2015.

[176] Alan Grossfield. WHAM: the weighted histogram analysis method, 2014. `http://membrane.urmc.rochester.edu/content/wham` Accessed: 2019-08-30.

[177] Min Yi and Zhigang Shen. A review on mechanical exfoliation for the scalable production of graphene. *Journal of Materials Chemistry A*, 3(22):11700–11715, 2015.

[178] Artur Ciesielski and Paolo Samorì. Grapheneviasonication assisted liquid-phase exfoliation. *Chem. Soc. Rev.*, 43(1):381–398, 2014.

[179] Da Chen, Hongbin Feng, and Jinghong Li. Graphene Oxide: Preparation, Functionalization, and Electrochemical Applications. *Chemical Reviews*, 112(11):6027–6053, 2012.

[180] Kirill I. Bolotin, Kenneth J. Sikes, Zhigang Jiang, M Klima, G Fudenberg, James Hone, Philip Kim, and Horst L. Stormer. Ultrahigh electron mobility

in suspended graphene. *Solid State Communications*, 146(9-10):351–355, 2008.

[181] Min Yi, Zhigang Shen, and Jinyang Zhu. A fluid dynamics route for producing graphene and its analogues. *Chinese Science Bulletin*, 59(16):1794–1799, 2014.

[182] Zhigang Shen, Jinzhi Li, Min Yi, Xiaojing Zhang, and Shulin Ma. Preparation of graphene by jet cavitation. *Nanotechnology*, 22(36):365306, 2011.

[183] Chih-Jen Shih, Shangchao Lin, Michael S Strano, and Daniel Blankschtein. Understanding the stabilization of liquid-phase-exfoliated graphene in polar solvents: molecular dynamics simulations and kinetic theory of colloid aggregation. *Journal of the American Chemical Society*, 132(41):14638–14648, 2010.

[184] Vitaly V Chaban, Eudes Eterno Fileti, and Oleg V Prezhdo. Exfoliation of graphene in ionic liquids: pyridinium versus pyrrolidinium. *The Journal of Physical Chemistry C*, 121(1):911–917, 2017.

[185] Vivekanand V Gobre and Alexandre Tkatchenko. Scaling laws for van der waals interactions in nanostructured materials. *Nature Communications*, 4:2341, 2013.

[186] Buddhika Jayasena and Shreyes N. Melkote. An Investigation of PDMS Stamp Assisted Mechanical Exfoliation of Large Area Graphene. *Procedia Manufacturing*, 1:840–853, 2015.

[187] Petr Lazar, František Karlický, Petr Jurečka, Mikuláš Kocman, Eva Otyepková, Klára Šafářová, and Michal Otyepka. Adsorption of small organic molecules on graphene. *Journal of the American Chemical Society*, 135(16):6372–6377, 2013.

[188] Anastassia N Rissanou and Vagelis Harmandaris. Dynamics of various polymer–graphene interfacial systems through atomistic molecular dynamics simulations. *Soft Matter*, 10(16):2876–2888, 2014.

[189] RR Nair, HA Wu, PN Jayaram, IV Grigorieva, and AK Geim. Unimpeded permeation of water through helium-leak–tight graphene-based membranes. *Science*, 335(6067):442–444, 2012.

[190] Peter V Coveney and Shunzhou Wan. On the calculation of equilibrium thermodynamic properties from molecular dynamics. *Physical Chemistry Chemical Physics*, 18(44):30236–30240, 2016.

[191] Alister J Page, Tetsushi Isomoto, Jan M Knaup, Stephan Irle, and Keiji Morokuma. Effects of molecular dynamics thermostats on descriptions of chemical nonequilibrium. *Journal of chemical theory and computation*, 8(11):4019–4028, 2012.

[192] Jan Wedekind, David Reguera, and Reinhard Strey. Influence of thermostats and carrier gas on simulations of nucleation. *The Journal of Chemical Physics*, 127(6):064501, 2007.

[193] Rafael I. González, Felipe J. Valencia, José Rogan, Juan Alejandro Valdivia, Jorge Sofo, Miguel Kiwi, and Francisco Munoz. Bending energy of 2D materials: graphene, $MoS_2$ and imogolite. *RSC Advances*, 8(9):4577–4583, 2018.

[194] Steven Cranford, Dipanjan Sen, and Markus J. Buehler. Meso-origami: Folding multilayer graphene sheets. *Applied Physics Letters*, 95(12):123121, 2009.

[195] Carl P Dettmann and Orestis Georgiou. Product of n independent uniform random variables. *Statistics & probability letters*, 79(24):2501–2503, 2009.

[196] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, volume 55. Courier Corporation, 1965.

[197] Daniel R. Dreyer, Sungjin Park, Christopher W. Bielawski, and Rodney S. Ruoff. The chemistry of graphene oxide. *Chem. Soc. Rev.*, 39(1):228–240, 2010.

[198] Chih-Jen Shih, Shangchao Lin, Richa Sharma, Michael S. Strano, and Daniel Blankschtein. Understanding the pH-Dependent Behavior of Graphene Oxide Aqueous Solutions: A Comparative Experimental and Molecular Dynamics Simulation Study. *Langmuir*, 28(1):235–241, 2012.

[199] Haiwei Dai, Zhijun Xu, and Xiaoning Yang. Water Permeation and Ion Rejection in Layer-by-Layer Stacked Graphene Oxide Nanochannels: A Molecular Dynamics Simulation. *The Journal of Physical Chemistry C*, 120(39):22585–22596, 2016.

[200] Martin Trömel and Manuel Russ. Dimanganheptoxid zur selektiven Oxidation organischer Substrate. *Angewandte Chemie*, 99(10):1037–1038, 1987.

[201] D. Pacilé, J.C. Meyer, A. Fraile Rodríguez, M. Papagno, C. Gómez-Navarro, R.S. Sundaram, M. Burghard, K. Kern, C. Carbone, and U. Kaiser. Electronic properties and atomic structure of graphene oxide membranes. *Carbon*, 49(3):966–972, 2011.

[202] Weiwei Cai, Richard D Piner, Frank J Stadermann, Sungjin Park, Medhat A Shaibat, Yoshitaka Ishii, Dongxing Yang, Aruna Velamakanni, Sung Jin An, Meryl Stoller, Jinho An, Dongmin Chen, and Rodney S Ruoff. Synthesis and solid-state NMR structural characterization of 13C-labeled graphite oxide. *Science (New York, N.Y.)*, 321(5897):1815–7, 2008.

[203] Sumit Saxena, Trevor A. Tyson, and Ezana Negusse. Investigation of the Local Structure of Graphene Oxide. *The Journal of Physical Chemistry Letters*, 1(24):3433–3437, 2010.

[204] Kris Erickson, Rolf Erni, Zonghoon Lee, Nasim Alem, Will Gannett, and Alex Zettl. Determination of the Local Chemical Structure of Graphene Oxide and Reduced Graphene Oxide. *Advanced Materials*, 22(40):4467–4472, 2010.

[205] LG Cançado, K Takai, T Enoki, M Endo, YA Kim, H Mizusaki, A Jorio, LN Coelho, R Magalhaes-Paniago, and MA Pimenta. General equation for the determination of the crystallite size $l_a$ of nanographite by raman spectroscopy. *Applied Physics Letters*, 88(16):163106, 2006.

[206] William S Hummers Jr and Richard E Offeman. Preparation of graphitic oxide. *Journal of the american chemical society*, 80(6):1339–1339, 1958.

[207] Daniela C Marcano, Dmitry V Kosynkin, Jacob M Berlin, Alexander Sinitskii, Zhengzong Sun, Alexander Slesarev, Lawrence B Alemany, Wei Lu, and James M Tour. Improved synthesis of graphene oxide. *ACS nano*, 4(8):4806–4814, 2010.

[208] Peter Wick, Anna E. Louw-Gaume, Melanie Kucki, Harald F. Krug, Kostas Kostarelos, Bengt Fadeel, Kenneth A. Dawson, Anna Salvati, Ester Vázquez, Laura Ballerini, Mauro Tretiach, Fabio Benfenati, Emmanuel Flahaut, Laury Gauthier, Maurizio Prato, and Alberto Bianco. Classification Framework for Graphene-Based Materials. *Angewandte Chemie International Edition*, 53(30):7714–7718, 2014.

[209] S Das, Sabita Patel, and Bijay K Mishra. Oxidation by permanganate: synthetic and mechanistic aspects. *Tetrahedron*, 2009.

[210] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed 26/08/2019].

[211] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos,

D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python . *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[212] Stephan Mertens and Cristopher Moore. Continuum percolation thresholds in two dimensions. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 86(6), 2012.

[213] J Quintanilla, S Torquato, and R M Ziff. Efficient measurement of the percolation threshold for fully penetrable discs. *Journal of Physics A*, 33(3300):399–407, 2000.

[214] John Quintanilla. Measurement of the percolation threshold for fully penetrable disks of different radii. *Physical Review E*, 63(6):061108, 2001.

[215] Kevin Dwyer. Drawing boundaries in python. `http://blog.thehumangeo.com/2014/05/12/drawing-boundaries-in-python/` Accessed: 2019-01-9.

[216] Michael Farries Ashby. Technology of the 1990s: advanced materials and predictive design. *Phil. Trans. R. Soc. Lond. A*, 322(1567):393–407, 1987.

[217] Tim Palucka and Bernadette Bensaude-Vincent. History of composites - Overview, 2015.

[218] G. B. Olson. Computational Design of Hierarchically Structured Materials. *Science*, 277(5330):1237–1242, 1997.

[219] James L. Suter, Peter V. Coveney, Richard L. Anderson, Hugh C. Greenwell, and S Cliffe. Rule based design of clay-swelling inhibitors. *Energy & Environmental Science*, 4(11):4572–4586, 2011.

[220] Leandro Martínez, Ricardo Andrade, Ernesto G Birgin, and José Mario Martínez. Packmol: a package for building initial configurations for molecular dynamics simulations. *Journal of Computational Chemistry*, 30(13):2157–2164, 2009.

[221] Shaorui Yang and Jianmin Qu. Computing thermomechanical properties of crosslinked epoxy by molecular dynamic simulations. *Polymer*, 53(21):4806–4817, 2012.

[222] Jihua Gou, Bob Minaie, Ben Wang, Zhiyong Liang, and Chuck Zhang. Computational and experimental study of interfacial bonding of single-walled nanotube reinforced composites. *Computational Materials Science*, 31(3-4):225–236, 2004.

[223] Changwoon Jang, Timothy W Sirk, Jan W Andzelm, and Cameron F Abrams. Comparison of crosslinking algorithms in molecular dynamics simulation of thermosetting polymers. *Macromolecular Theory and Simulations*, 24(3):260–270, 2015.

[224] JL Valentín, J Carretero-González, I Mora-Barrantes, W Chassé, and K Saalwachter. Uncertainties in the determination of cross-link density by equilibrium swelling experiments in natural rubber. *Macromolecules*, 41(13):4717–4729, 2008.