Structures, Vibrational Modes and Charge Transport of Organic Semiconductors

Smith, Alexander

Award date:
2020

Awarding institution:
University of Bath

Structures, Vibrational Modes and Charge Transport of Organic Semiconductors

submitted by

Alexander R. Smith

for the degree of Doctor of Philosophy

of the

University of Bath

Department of Physics

September 2019

DECLARATION

The material presented here for examination for the award of a higher degree by research has not been incorporated into a submission for another degree.

In memory of my mother, Melanie.

Acknowledgements

Producing this document has only been possible due to the help and goodwill of others. I would therefore like to begin by giving due thanks to those that have lent support, assistance and advice that lead towards its production and credit those who contributed directly to this project.

I would like to begin by thanking Professor Alison Walker for creating and supervising this project and providing me with the opportunity to pursue it. I would particularly like to thank her for proofreading the many unfinished drafts of this document.

Thanks must also be given to the many people who directly contributed to this work by providing measurements, calculations and pieces of code. I would like to thank: Ian Thompson, who wrote the optimisation that produces the minimum volume enclosing ellipsoids and pointed this effort towards the realm of glass physics and structural relaxation; Thomas McManus who's implementation of the template rotation and fitting algorithm features heavily in the implementation used to generate all of the results in this book; Dibyajyoti Ghosh who performed all of the DFT-level optimisations of the molecular structures and Yoann Oliver who performed the transfer integral calculations for the IDT-BT morphology; Krishna Feron who was my host during a visit to the University of Newcastle, New South Wales and who provided information on previous P3HT nanoparticle simulations as well as motivating and enabling the ones presented here; and Levi Tegg who performed the XRD measurements on the P3HT nanoparticles. I would also like to thank David Bossanyi, Joshua McDonald, Enrico Da Como and Jenny Clark for further engagement on excitons in P3HT.

I would also like to thank Otello Roscioni and Claudio Zannoni for providing us with access to results on several of their coarse-grained molecular dynamics simulations and Vincent Lemaur and David Beljonne for providing us with atomistic results. I shall also thank them again, along with Peter Bobbert, Chad Risko, Roland Faller, Wolfgang Wenzel, Burag Yaglioglu (and many others) for useful discussions across several EXTMOS project and engagement meetings. I'd also like to thank former students Jacob Peach and Robin Winchcombe for their work on CGMD P3HT simulations in LAMMPS. Further thanks should be extended to Tom Underwood for discussions on SAMSEN and the energy landscape and to Mike Allen for feedback on the method.

Author's Note

Chapters 3 and 5 both draw heavily and include figures and results from previously published work (Smith et al., J. Chem. Phys. 150(16) 2019). Chapter 3 presents SAMSEN from a different perspective than that previously published and expands on its applicability, giving particular regard to polymers and long-time dynamics of multi-component systems. Chapter 5 adds consideration to the distribution of potential minima and how it affects structural and charge transport measurements.

Abstract

Organic semiconductors are quietly revolutionising the modern world, finding application in light-emitting diodes, thin-film transistors and photovoltaics. However their short device lifetimes and reduced power-conversion efficiencies have limited their commercial uptake. Both experimental and computational work to find new materials and architectures and study of their charge transport properties continues.

Computational studies of the charge transport properties of organic semiconductors depend heavily on the structure/morphology. Structures are often assumed to be a regular lattice, even in the case of amorphous materials, with some groups using molecular dynamics (MD) and Monte Carlo (MC) methods to capture the disorder. However these approaches have significant drawbacks, some of which are technical (MD and MC can be computationally arduous at large system size) and some of which are physical (below the glass transition, the energy landscape makes it difficult to sample a large number of states). Other approaches such as coarse-grained and basin-hopping methods attempt to overcome these problems but often create new problems such as the need to develop new force-fields, having to regain full-atomic representation of molecules afterwards, and still being hindered by the landscape.

In this work, a new method, Simulation of Atomistic Molecular Structures using an Elastic Network (SAMSEN), is proposed and applied to molecular and polymeric systems. SAMSEN contains both a structural and dynamical model that individually attempt to overcome the technical and physical problems of current methods. The structural model requires that molecules are split into rigid sections, which retain their atomistic representation and are constructed to interlock with their neighbouring sections, and restrict the maximum displacement of atoms from their locally optimised positions. Atomic collision rules, limiting minimum separations, are also enforced. The combination allows SAMSEN to recreate the short-range structure of weakly-interacting non-polar small molecules. The dynamical model creates an elastic network between the rigid sections and displaces them along the low-frequency vibrational modes to achieve collective large-scale motion and computationally-fast structural relaxation.

SAMSEN is applied to systems of spheres to study the structural and dynamical parameters and a regime is found where a band of collective low-frequency modes can be found, sampling rate can be increased without altering structure and the mean overlap of atoms can be controlled without altering sampling rate. The structural parameters are determined entirely by the class of system being studied, leaving the dynamical pa-

rameters to be chosen to maximise sampling rate. SAMSEN is then applied to systems of small molecules and is shown to be widely applicable, providing good approximations to the short-range structures produced by full atomistic force-field methods. SAMSEN phenyl-C_{61}-butyric acid methyl ester (PCBM) states are found to have structures close to those of the higher energy minima in a potential energy landscape of an all-atom potential and the distribution of these potential minima is found to be near-Gaussian. SAMSEN outputs are then used as inputs of MD simulations, recreating the full-MD structures and quickly finding the lower-energy minima. This method provides a useful pathway for those interested in sampling the distribution of morphologies at equilibrium. The diffusion of electrons is simulated for each state in the distribution of MD and SAMSEN inherent structures and it is found that the diffusion coefficient is near-independent of the potential of the minima, despite short-range structural differences.

Turning to polymers, the coarse-graining into rigid sections now controls the relative structure between repeat units. The structure and vibrational modes of poly(3-hexylthiophene) (P3HT) is studied in the pure amorphous phase and also when blended with PCBM. Increased persistence of the P3HT backbone is found upon mixing with PCBM but the density of neighbouring chains is unaffected up to the miscibility limit. Studying the long time (low frequency) modes, the rate of relaxation of the P3HT backbone is slowed by the addition of PCBM due to a shift in frequencies, rather than a change of collective behaviour of the individual modes. The excitonic transport properties of P3HT are then studied in the amorphous and crystalline phase in a transfer matrix approach, finding the exciton diffusion length in the low charge limit and comparing to experimental and previous computational work on P3HT nanoparticles. Low disorder polymer, Indacenodithiophene-co-benzothiadiazole (IDT-BT), is also studied, with a morphology generated, the inter-molecular transfer integrals calculated using a quantum chemistry package and the diffusion constant for holes determined for a range of intra-molecular transport rates. Comparing the hole diffusion constant to experimental and computational work, an estimate of the intra-molecular transport rate is made.

This work serves as a framework for researchers interested in determining the structures, long-time vibrations and charge transport properties of weakly-interacting amorphous organic semiconductors where sampling of states is required. It can achieve this without molecule-specific parameters or forces and does so at modest computational cost and therefore opens the possibility of sampling a large number of states or simulating bigger system sizes (approaching device scale) without requiring access to high-performance computing resources. Further opportunity exists to constrain the dihedral angles, con-

sider electrostatic interactions, as well as to study the charge transport properties of polymer systems incorporating an accurate model for intra-chain charge-transport.

Contents

Chapter 1

Introduction

Organic semiconductors are re-shaping modern electronics. These molecular-based semiconducting materials hold a significant advantage over traditional solid-state electronic materials: they are cheap to manufacture. Unlike silicon-based technology, these materials can be solution processed at room temperature, driving down production costs for companies willing to make the transition. With an expansive range of molecules to choose from, a manufacturer can pick an organic semiconductor (or modify an existing molecule) that has a property specific to its needs; for instance, the band gap can be selected to provide a certain colour or a material can be used which allows greater mechanical flexibility. Benefits such as these have begun to make organic semiconductors (OSCs) commercially competitive and they have already found a home in a wide-range of devices including televisions and mobile-phone displays as organic light-emitting diodes (OLEDs), wearable devices and biological sensors as organic field-effect transistors (OFETs) and solar cells in organic photovoltaics (OPV) [1].

Research into organic semiconductors and device manufacturing began in earnest after the award of the 1992 Nobel Prize in Chemistry to Rudolph A. Marcus whose reaction rate equation is used to describe the transfer of charge between two neighbouring molecules [2] and the 2000 prize to a group "for the discovery and development of conductive polymers" [3]. Since then, different device architectures were designed and tested. Experimental groups began work on optimising parameters such as layer thicknesses, processing temperatures and solvent evaporation rates [4]. Computational work was carried out on modelling charge mobility through master equation and, later, kinetic Monte Carlo approaches [5, 6, 7, 8]. And new theories, such as the semi-band-like

transport were created to account for transport in materials with more delocalised charges both within and between molecules [9, 10, 11].

These efforts have led to the selection of better combinations of molecules such that devices now produce higher operating voltages and currents, higher absorption/emission efficiency, and exhibit lower degradation rates [12, 13, 14]. This has been in attempt to improve the three properties that determine if a material is suitable for commercial use: device lifetime, power-conversion efficiency and charge mobility.

The problem with OSCs

In the silicon solar panel industry, it is common to guarantee a given output - usually 80% of initial performance - for 25 to 30 years after purchase [15]. This allows buyers to calculate the break-even point on the purchase and installation of a solar panel. For silicon this usually occurs within the guaranteed device lifetime (the exact time depends on the level of subsidies and the electricity tariffs in the given jurisdiction). In organic photovoltaics (OPV), device lifetime is still a significant factor for investors to consider, with most commercial devices lasting around 2-3 years [16, 13, 17]. A shorter lifetime than silicon is acceptable due to the lower unit cost of OPV modules, however the break-even time needs to be lowered before OPV is widely adopted in capital-rich economies. Degradation pathways include delamination, chemical reactions with contaminants and mechanical damage caused by repeat flexing or impact. Investigation into reducing degradation and/or the lowering of the unit cost by molecular-selection or alternate processing methods is therefore a current topic of research [17, 15, 12, 18].

Another strategy for improving the commercial viability of OPV is to maximise the amount of power generated and in the case of OLEDs the aim is to maximise the reverse process - light emitted from a given power input. The device's effectiveness in these processes is known as the power conversion efficiency. This accounts for the energy that could have been absorbed and passed to the connecting circuitry (or received from the circuit emitted and released as light, in the case of OLEDs) but was not. In OPV this accounts for a range of losses such as low absorption coefficient, thermalisation and charge recombination. These losses are again dependent on the material, device architecture and processing methods. In particular, the device morphology—the structure and arrangement of molecules, including the structure and arrangement of domains of molecules in blends—has a strong effect as it directly affects a charge's ability to move through an organic semiconductor film and either reach an electrode to do useful work or recombine with an opposite charge to release light [19, 5, 20].

For all organic devices, however, the charge mobility is seen as key to their commercial advancement. This quantity determines how a charge responds to an applied field, and therefore influences the electrical characteristics of the device, as well as the internal quantum efficiency. In OFETs, the key metrics for applications such as communications and displays is switching time and the on/off voltage ratio [21, 22, 23]. A high device mobility means a stronger or faster response to an applied field which reduces the switching time, allowing for faster electronics. This property is also dependent on the material and the morphology.

Computer simulation as a strategy to optimise OSC performance

As has already been demonstrated, the choice of organic semiconductor matters significantly and as a result, a large amount of experimental and computational work is being performed to find new, better-suited organic materials. It is also clear that the morphology of the organic layers has a strong effect on device performance and a significant amount of computational work is being performed to determine how the micro-structure affects the macro-scale properties [5, 20, 24, 25, 26, 27]. Computer simulations to study these two factors will be the primary concern of this work.

Computer simulations of organic semiconductors can, for the most part, be split into three categories: molecular and electronic structure prediction, morphology prediction, and charge transport simulations. All three are, in principle, required to assess the suitability of a new molecule for use in an OSC device.

The first is usually accomplished by density-functional theory (DFT) methods and provides the optimised unexcited molecular structure and energy, its associated excited/reduced energy levels, and other properties such as its internal vibrational modes. This quantum-level approach is burdensome to compute and is therefore often limited to single molecules or oligomers in the gas phase (e.g. PCPDTBT [28]) or a few unit cells for molecular crystals such as pentacene [29], ruberene [30] or perylene [31]. However, the results are essential as they determine the parameters used in the abstraction to larger scale simulation.

The second category is usually performed by less computationally intensive methods such as molecular dynamics (MD), coarse-grained (CG) molecular dynamics or Monte Carlo methods (MC) such as Metropolis Monte Carlo (MMC). These rely on classical (or semi-classical) force-fields, parameterised from lower-level simulations such as DFT or ab-initio dynamics, to determine the energy (and control other thermodynamic

properties) of a simulated state [32]. This provides the basis for evolving one state into future states. In MD, particles step forwards in time according to their calculated velocities and accelerations often using the velocity-verlet algorithm [33]. In MMC, a new state is proposed and then accepted if the energy is less than the current state or with an Arrhenius-like probability if there is an increase in energy [34]. These methods are computationally simpler to perform than DFT methods, but can get expensive as system sizes increase. Most simulations are limited to roughly a thousand molecules, as this can take CPU weeks or months to perform at high densities - where the slowing of dynamics limits the number of states that can be sampled [35]. Coarse-grained methods have been developed to reduce the number of calculations required by parameterising the force-field further by grouping atoms and their interactions together, reducing the number of calculations needed to be performed [35, 36, 37]. This has an impact on the overall accuracy but increases the system sizes that can be studied.

The third category, charge transport simulations, is then performed on a generated device morphology. These methods can be broken down into continuous approaches, such as drift-diffusion (DD) and direct master equation (ME) methods, and discrete event approaches such as kinetic Monte Carlo (KMC). In the steady-state limit using on-lattice models, the three approaches are equivalent, however the level of detail and convergence time vary. The DD and ME approaches smear charges into densities which reside on given sites (such as a molecule), the probability of charges flowing to a neighbouring site is calculated and the densities updated accordingly [38, 39, 40, 41]. In principle, the steady state can be solved for exactly (a stable solution is guaranteed) [42] unlike in KMC where the steady-state properties are only approximated from samples in the long-time limit. In a KMC, particles hop around the morphology with calculated rates modified by randomly picking from a log-uniform distribution [19]. This means that microscopic events in the system such as charge extraction and recombination, as well as correlations and interactions with other particles, can also be recorded.

In all cases of charge transport simulation, finite size effects and a limited number of sample morphologies can lead to significant inaccuracies in the prediction [43, 44]. In lattice ME approaches, for instance, systems of millions of sites (or 100 nm dimension) are required before finite size effects become unobservable [43, 45]. Similarly, the set of morphologies which are used for simulation must also be distributed to reflect the real-world distribution before it can be said that a charge transport prediction predicts the real-world device characteristics. However, producing morphologies on this scale is a computationally expensive task. Most studies are limited to a handful of atomistic states of approximately a thousand molecules for this reason. If charge transport

predictions on non-crystalline morphologies are to become systematically reliable, the obstacle of limited system sizes and limited states needs to be addressed.

The purpose and outline of this thesis

This work is primarily concerned with addressing those two problems: improving system size and increasing the number of sampled morphologies. This is achieved by creating a new coarse-grained method, Simulation of Atomistic Molecular Structures using an Elastic Network (SAMSEN) [46], which retains the atomistic representation of a molecule by preserving the strongly repulsive region of atomistic interactions and the stiffness of the bonded potentials but removes the longer range interactions. This allows larger systems to be studied through a reduction in the computational work. The method is also molecule-agnostic as it doesn't require a molecule-specific force-field and applies well to non-polar amorphous small molecules and polymers.

However, the main benefit of SAMSEN is it overcomes the problem of long-time relaxation (experienced by MD and MC) and therefore poor sampling of the potential energy surface by instead calculating the dynamics between coarse-grained units and evolving the system along the low-frequency (long-timescale) modes. It is possible to use this method to generate morphologies of small molecules and polymers in the amorphous phase with structures similar to those from classical simulation methods with a distribution of potential energies that is only very weakly biased. It is then also possible to recover the finer details by using the output states as inputs in a very brief MD simulation. By performing MD (or at least minimising) in an atomistic-force field across a large sample of SAMSEN states it is then possible to attempt to study the material in thermodynamic equilibrium.

I shall demonstrate that using this method it is therefore possible to reach a large number of states or system sizes in CPU weeks and therefore possible to produce charge transport predictions without significant sampling errors or finite size effects. I shall also demonstrate the broad applicability of this method to organic semiconducting systems, including polymers, and discuss the prospects for improving the structural model.

The rest of this work is therefore laid out as follows. In chapter 2, I will present the background material and theory underpinning this work, describing organic semiconductors from the level of molecules to devices and considering their charge transport, the physical effects that limit dynamical simulations and how they might be overcome.

In chapter 3, I shall then establish the structural and dynamical model and the SAMSEN algorithm used to generate and sample morphologies. I will also describe the transfer matrix approach that will be used for charge transport prediction. I will go on, in chapter 4 to describe the effect of changing simulation parameters by studying systems of hard spheres. The structural and dynamical model will be tested and independent control over the structures and sampling rate will be established. In chapter 5 I will then proceed to demonstrate the validity of the structural model by studying three small molecules and comparing their structures to those of other simulation techniques. I will also show that by following the long-time dynamics of the low-frequency vibrational modes, it is possible to relax the structure in short compute times. I will also discuss the distribution of the potential energy of the configurational states generated using SAMSEN and show one can use MD, starting from the SAMSEN states (minimised in an atomistic force-field) to obtain the lower energy amorphous states. From chapter 6 onwards, the model is used investigate the structure, low frequency modes and charge transport in typical OPV and OLED materials. I present a study of polymer-fullerene blends in chapter 6 and how the mixing ratio alters their dynamics and the conformations of the polymers. In chapter 7, I also study exciton diffusion in crystal and amorphous morphologies of poly(3-hexylthiophene) (P3HT) and compare to experimental measurements of P3HT nanoparticles. Consideration is also given to the finite size effects on the charge transport predictions. In chapter 8 I then proceed to study hole diffusion in a donor-acceptor co-polymer and its dependence on the intra-chain charge transport rate. In chapter 9 I will then summarise my findings and comment on the suitability of the models employed here to simulate organic morphologies and predict charge transport properties of organic semiconductors.

Chapter 2

Background

2.1 Organic Semiconductors

In many ways, organic semiconductors (OSCs) can be considered similar to conventional inorganic semiconductors: electrons fill distinct energy levels in the material and these energy levels are such that there is a gap between a valence state (the bonding band which in OSCs is the highest occupied molecular orbital or HOMO) and the conduction state (the lowest unoccupied molecular orbital, the LUMO). Electrons in the LUMO and the absence of electrons (holes) in the HOMO can then move throughout the material and respond to an applied electric field. This similarity has allowed OSCs to begin to compete and, in some cases, displace silicon in semiconductor applications such as field-effect transistors and light-emitting diode displays. In other ways, OSCs have some distinct differences from the conventional materials which can provide some advantages and disadvantages in their charge transport, mechanical and structural properties. This chapter shall be dedicated to describing these key intrinsic differences which also modify how organic semiconductor devices are designed and fabricated and alter how they need to be both theoretically and computationally modelled.

2.1.1 Molecules, HOMO, LUMO and the Band Gap

The conducting and semiconducting properties of organic molecules is determined by the bonding arrangement of its constituent atoms. A simple alkane chain consists of carbon atoms bonded to two other carbon atoms (one at the ends of the chain) and also to two hydrogen atoms (three at the ends of the chain). This alkane, take hexane as an

example, can be drawn as a Lewis structure that has only single bonds connecting any carbon to another [47]. In other molecular structures, such as benzene, it is possible to draw alternating double and single bonds between two carbon atoms (there are multiple resonant structures) [48]. This alternating double and single bond pattern allows the electron p-orbitals to overlap across the σ-bond (the single bonds in this case), forming a conjugated system. π-orbitals, which are delocalised orbitals across the aligned set of p-orbitals, then form. In the case of benzene, two π-orbitals form and lie above and below the plane of the molecule and are delocalised across the diameter of the molecule.

The delocalised molecular orbitals are of chief concern in organic semiconductors as they contain the states that are the equivalent of those in the valence band and conduction band in inorganic semiconductors [49]. These two electronic states/orbitals, referred to as the highest-occupied molecular orbital (HOMO) and lowest-unoccupied molecular orbital (LUMO), represent the ionisation potential of the neural molecule and the electron affinity (the energy change for adding an electron to the neutral molecule) respectively [50, 51, 49]. These are sometimes taken to be the energies at which, after absorption of a photon, the electron vacancy (the hole) and the excited electron exist - but due to an interaction between the hole and electron, this is not strictly the case [49]. They must both be calculated at the quantum chemistry/density-functional theory level and the precise methodology used can alter the result [50].

Between the HOMO and LUMO, a gap in energy levels exists: the HOMO-LUMO gap (sometimes the energy gap), which is often referred to as the band gap or the optical gap or the fundamental gap. Strictly speaking, each of these has a different and precise meaning described well by Bredas [50] and Kahn [49]. In this work, we shall follow that convention and 'HOMO-LUMO gap' shall be used to describe the difference in the calculated HOMO and LUMO energy in vacuum and 'energy gap' the difference between the two levels for a specific molecule in the bulk.

It is, however, difficult to consider a small molecule or polymer organic semiconductor having HOMO and LUMO bands in the same way that inorganic semiconductors have valence and conduction bands. Due to disorder, the HOMO and LUMO electronic wavefunction is often confined to an individual molecule (or conjugated region) and lacks the extend bonded structure of inorganic crystals required to exhibit band-like transport or to even draw a band-diagram [50, 26]. Instead, many individual molecules with their valence and conduction states are positioned in proximity to one another to form a bulk material. The HOMO and LUMO energies of a given molecule is then altered by interactions with neighbouring molecules. Small scale structural changes of the molecule due to thermal fluctuations or interactions with surfaces or electric fields

also alter the energy levels [50, 49]. This creates a distribution of energy levels of the HOMO and LUMO within the bulk material - a 'band' of HOMO and LUMO levels is then sometimes referred to and the band-gap represents the difference between the lowest LUMO energy and the highest HOMO energy in the system [50].

It should be noted that in some cases (particularly molecular crystals with low disorder), inter-molecular delocalisation also occurs and charge transport can also sometimes be described as band-like [24, 25, 52, 53, 50, 54, 55], however this work shall focus on amorphous organic semiconductors and inter-molecular delocalisation shall not be considered in further detail here.

2.1.2 Polymers and Charge Delocalisation

Across a larger molecule, such as a polymer, with an extended conjugated system, these orbitals can, in principle, be delocalised over the length of the polymer. However, the amount of delocalisation depends strongly on the given conformation of the polymer and its chemical structure [56]. In the case of the former, it is understood to be important to consider how twists (change in the torsional angles between polymer segments) and bends (change in the bonding angle) alter both the delocalisation of the electronic states and, if a polymer consists of multiple delocalised states, the rate of charge transfer between the delocalised states [56, 11, 9].

The extent of intra-molecular delocalisation therefore depends strongly on the torsional potential that biases the distribution of torsional angles. In some polymers, such as P3HT, this torsional potential is very weak and it is common to find 180° rotations away from the minimum-potential configuration [57]. Therefore, before considering charge transport in a polymer, strict consideration of the extent of delocalisation should be made on a per-polymer basis (the regions over which the conjugation is unbroken should be identified). These regions should be re-identified whenever the polymer moves. Then on a per delocalised orbital basis, the amount of overlap between the orbitals and the rate of electron transfer between the orbitals can then be calculated. Methods such as the flexible surface hopping or a method such as that of Fornari et al. provide a framework for incorporating these effects [9, 11].

2.1.3 Excitons

As in inorganic semiconductors when light is absorbed by the material, an electron, under the correct conditions, can absorb the energy of the photon and elevate itself into an excited state. This transition is from the valence band to a higher energy state, ideally at or above the bottom of the conduction band. In organics this transition is from the HOMO to the LUMO (or LUMO+1, LUMO+2, etc.) and the change in energy is known as the optical gap which also accounts for the interaction between the excited electron and the electron vacancy/hole [50, 51, 58, 54]. These charges are unable to separate at room temperature, as they might in inorganic semiconductors, due to the low dielectric constant of OSCs, often given as $\epsilon_r \approx 3.5$ [19, 51, 54]. This creates an attraction - the binding energy - calculated by the Coulomb potential with a distance of approximately the size of the molecule/conjugated region, that is larger than the thermal energy at room temperature (≈ 0.026 meV). The electron and hole, unable to separate, are then treated as a charge-neutral quasi-particle: an exciton [51, 59, 54, 60].

The exciton will be delocalised over some region, defined by the wavefunction and often over large sections of the molecule or several repeat units in a polymer. Further coupling of the exciton to low-frequency vibrations, torsional changes to the structure of the molecule (or the polymer) and the influence of the local environment can all increase the localisation of the excited state [54, 56, 61]. The excited electron is then either able to transfer to neighbouring conjugated regions in the system or transition to lower energy states by releasing another photon.

Excitons can be either spin singlets or spin triplets with the singlets being the emissive species in π-conjugated molecules [62, 63]. The ratio of formation of singlet and triplet excitons can vary between molecule/polymer [64, 65] but is usually taken to be 1:3 [66]. The singlet exciton is usually modelled as transporting via a resonance energy transfer process (see section 2.2.1) and the triplet via a Dexter-type transfer [60, 67]. We will not, however, consider the distinguishing characteristics in further detail and only consider singlets in this work.

In order for an organic photovoltaic device to do useful work, the exciton must find an additional amount of energy so that the electron and hole can separate. This is usually achieved by constructing a device with an interface between two materials which have been selected such that the total energy difference when one of the species transfers into the other material is larger than the exciton binding energy [68, 19, 51, 58]. If an exciton reaches an interface and the charges separate, the electron and hole then move in opposite directions under the influence of an external electric field. The charges

continue to attract one another however and can recombine and re-form an exciton [68, 19, 5]. If separation doesn't occur quickly enough, the exciton can relax back to the ground state and release that energy as a photon.

2.2 Charge Transport Models

The process of charge transfer in organic semiconductors is different to inorganic semiconductors, as already noted: the localisation of electronic states to a molecule or conjugated region means band-like transport with a momentum vector is a poor description of charge transport in organics [60, 25, 52]. Instead, transport is often described as a series of hops between discrete sites - taken as the centre of a molecule or the centre of the wavefunction of the HOMO or LUMO. A model for hopping, often dependent on the particle species or structure of the material, must be chosen before simulation can begin. Throughout the rest of this work the following three transport models will be used.

2.2.1 Hopping Models

Marcus Theory

Perhaps the choice charge hopping model in the organic semiconductor literature, Marcus theory attempts to describe the rate at which a pair of redox reactions involving two neighbouring molecules will take place [2, 60]. The idea is that there exists an energy surface which describes the energy of many possible states of the system (the molecules involved in the reaction and surrounding solvent molecules). This landscape includes the two states (or range of states) which are the reactant and products of the reaction. To reach the products, the reactants must gain enough energy to cross the barriers in the landscape and reach the product states. Between the two states a saddle-point region may exist known as the transition state(s) which can be used to determine reaction rates from statistical mechanics [2]. The increase in energy must describe the rearrangement of the nuclear coordinates of the molecules and solvent as well as the change in electronic energy. To calculate the rates of reaction, Marcus assumes the potential energy surface around the reactants and products can be approximated as two parabolas, neglecting electronic coupling, which intersect. An electron transfer, now calculated including electronic coupling, can occur at this point of intersection [2, 69, 70].

The standard free energy of reaction is ΔG^0 and is the difference between the minima of reactants and products, ΔG^* represents the activation barrier for the reaction and is related to the reorganisation energy λ, which includes both solvent and vibrational components and describes the transition from reactant to product state (in the geometry of the reactant state) [2, 71]. The Marcus rate expression is of the form of a chemical reaction rate with barrier ΔG^* which can, due to the parabolic approximation, be described by the free energy of the minimum structures and is often written as

$$k_{ij} = \frac{|J_{ij}|^2}{\hbar}\sqrt{\frac{\pi}{\lambda_{ij}k_BT}}\exp\left(-\frac{(\Delta G^0_{ij}-\lambda_{ij})^2}{4\lambda_{ij}k_BT}\right) \tag{2.1}$$

Determination of the Marcus rate of charge transfer between two molecules therefore requires knowledge of three quantities: the electronic coupling (sometimes electronic overlap') between two molecules/segments, J_{ij}, the molecular reorganisation energy for the transfer λ_{ij} and the relative energy levels of the product and reactant states (including the difference in HOMO or LUMO, as appropriate), ΔG^0_{ij} and k_B is the Boltzmann constant, T is the temperature. The calculation of these values can become a serious undertaking. They, strictly, require quantum-level calculation of electronic wavefunctions of the two molecules, i and j, in their local environment before and after transfer. Approximations are often made, such as the HOMO, LUMO and re-organisation energy can be calculated in a vacuum (with the HOMO and LUMO modified to mimic the local disorder). However, alterations to the molecular or external geometry are usually taken to not persist or alter the environment of other molecules for subsequent electrons transfers. For simplicity, we will also follow that convention here.

Miller-Abrahams

Another expression for charge transport often found in the literature is the Miller-Abrahams equation [72], given by

$$\nu_{ij} = \nu_0\exp\left(-2\alpha r_{ij}\right)\times\begin{cases}\exp\left(-\Delta E/k_BT\right), & \Delta E>0\\ 1 & \Delta E\leq 0\,.\end{cases} \tag{2.2}$$

where α is the inverse localisation constant, ΔE is the difference in site-occupation energy for the relevant charge, ν_0 is a rate prefactor constant and ν_{ij} is the rate of the charge transfer between two sites. The Miller-Abrahams equation is therefore a repre-

sentation of a polaron (a charge that distorts the local structure) which, due to disorder and a lack of band-structure, undergoes a tunnelling process in low-temperature conditions and has been described as a low-temperature version of the Marcus expression [60]. It has therefore found widespread use in organic semiconductor charge-transport simulations [73, 74, 75, 36, 5] and has been shown to produce similar results to Marcus theory in some instances [76]. For instance, considering inter-chain hops, the Miller-Abrahams expression (excluding the energetic term) was once described to be a good representation of the electronic couplings $|J|$ of the Marcus expression [19] (although it should be noted that $|J|$ strictly depends on both the distance and orientation of neighbouring molecules while equation 2.2 does not). It has also been used to describe charge injection from an electrode into the organic layers of OSC devices [77].

We will use the Miller-Abrahams expression here in situations where we do not have the electronic couplings between molecules and use an appropriate value of α taken from the literature.

Förster Resonance Energy Transfer

For the transfer of excitons an alternate model is required. Förster resonance energy transfer (FRET) is a quantum-mechanical description of energy transfer of the excited state. In the original paper published in 1948 [78], Förster describes a point-to-point model of energy transfer between excited states in organic molecules. In that work, the formal quantum-mechanical treatment is presented, which describes the interaction energy between two dipoles (the bound electron and hole in OSCs) at a given relative orientation and separation. The result was that for two dipoles in separate molecules, i and j, the rate of transfer between them can be described by a rate

$$\nu_{ij} = \frac{1}{\tau}\left(\frac{R_0}{r_{ij}}\right)^6 \times \begin{cases} \exp\left(-\dfrac{\epsilon_j - \epsilon_i}{k_b T}\right), & \Delta E > 0 \\ 1, & \Delta E \leq 0 \end{cases} \tag{2.3}$$

where R_0 is the Förster radius (which describes the localisation of the exciton) and ϵ represents the exciton energy at that site (strictly, the $S_0 \rightarrow S_1$ optical gap for singlet excitons) [78, 79]. Excitons are also able to recombine (sometimes 'decay' as it describes the de-excitation of the electron to the HOMO) and do so at a rate with a decay constant equal to the reciprocal of the characteristic exciton lifetime, τ, as [78]

$$\nu_{rec} = \frac{1}{\tau}. \tag{2.4}$$

The Förster radius, R_0, is strictly defined in terms of τ - it is the distance at which there is a 50% chance of the exciton transferring if one waited over a period of time, τ [73]. This property, combined with equation 2.3, means τ does not alter the statistical distance an exciton might travel before decaying [80]. In the 1948 paper it was also written that ${R_0}^6 \propto \kappa^2$ where κ was the dipole orientation factor which is given as

$$\kappa = \hat{\boldsymbol{u}}_D \cdot \hat{\boldsymbol{u}}_A - 3(r_{DA} \cdot \hat{\boldsymbol{u}}_D)(r \cdot \hat{\boldsymbol{u}}_A) \tag{2.5}$$

where $\hat{\boldsymbol{u}}_D$ and $\hat{\boldsymbol{u}}_A$ are the unit vectors for the donor and acceptor transition dipoles respectively (molecules i and j, previously) and r_{DA} is the separation between their centres [73, 78]. In situations where the the molecules can reorient under rotational Brownian motion at a rate much faster than the rate of energy transfer, the number of energy transfers per unit time can simply become an average across all relative orientations of the pair of molecules at that separation. However, in more viscous materials, the theory strictly requires that the transfer rate between two molecules is determined for their specific relative orientations. For a dynamic system that can reorient fast and freely, $\kappa^2 = 2/3$ for all pairs. For a static system of molecules in random orientations, $\langle\kappa^2\rangle = 2/3$, and for a completely aligned system, $\langle\kappa^2\rangle = 4$.

In some work [81, 82, 83], the effect of incorporating the relative orientations into the transfer rate has been studied. This static transfer rate is then compared to the dynamic (fast rotation) transfer rates. Other studies [84, 85] used the system average relative dipole orientations and applied it to all pairs in the system and, in their polymer film, reported it as being less than the unconstrained dynamical time average. For the truly dynamical model, $\kappa_{ij}^2 = 2/3$ while, for the polymer system, $\langle\kappa^2\rangle = (0.845\sqrt{2/3})^2$. The strict application of the Förster theory requires the rate to be determined from individual pairs of dipoles, however [78].

2.2.2 Intra-chain Transport

Charge transport in polymers is usually described by a different mechanism to the charge hopping and resonance transfer described in the previous section. For instance, with excitons, it is common to employ some form of Dexter transfer [67, 73, 86] between conjugated regions where the torsion within the polymer has not broken the conjuga-

tion of π-orbitals [56]. For electrons and holes, several methods exist with their basis in quantum-mechanical approaches that describe how the wavefunction moves throughout. One such approach is the flexible surface hopping approach discussed by Wang et al. [9]. Another method is described by Fornari and Troisi [11, 87]. Other research has taken a simpler approach and, assuming charges travel much faster intra-chain, use a high-valued fixed rate constant as an intra-chain hopping rate or, at a given time-step in a KMC simulation, giving a particle a random position within a polymer chain (which is equivalent to a very fast rate) [19].

In some research it is still common, however, to treat charge transport as hopping within the polymers. Some lattice-models perform this implicitly or do not make the distinction between intra- and inter-chain transport [80, 18, 88]. As the study of the mechanism of intra-chain transport is outside the scope of this thesis, charge transport within polymers will here be described as hopping using a combination of the Marcus theory rates and a constant internal transport rate. For exciton transport, only Förster transfer rates, described previously, will be considered. However, it should be noted that a more accurate representation exists and could be utilised in principle.

2.3 Disorder in OSC Simulations

The hopping models discussed in the previous section all explicitly describe some degree of disorder in the charge-transfer efficiency between any two given sites. The Marcus equation attempts to describe the vibrational/geometric disorder as the limiting process for charge transfer, with the reorganisation energy and the energy of the reactant and product states for defined for a single transfer process between two molecules [2]. This, however, can vary across the system and between pairs of molecules involved in potential transfer processes. The Förster rate for exciton transfer also puts great important on the alignment of the molecular dipoles involved in the transfer process which can, again, vary greatly between all points of molecules in the system and over time [78].

For the Marcus theory, Miller-Abrahams and FRET rate equations the most obvious source of disorder, however, is in the energy of the before and after states. In the Marcus equation, the energies of the reactants and product states as well as the reorganisation energy, while in Miller-Abrahams and FRET it is the change in electronic energy before and after transfer. The dependence of the rate equations on energy change acts to reduce the probability of a charge increasing in energy and, in the case of the Marcus

equation, increase the probability of a charge hopping to a site with a lower potential. The energetic disorder therefore has a large influence on the charge dynamics within a simulation.

This disorder substantially increases the computational work required to mimic charge-transfer of an organic semiconductor system in a computer simulation. Each of the rates depends on the relative position (Marcus: $|J|^2$; Miller-Abrahams: $\exp(-\alpha r_{ij})$; FRET: $(R_0/r_{ij})^6$), the relative orientation (Marcus: $|J|^2$; FRET: ${R_0}^6 \propto \kappa^2$) and the energy of the states before and after transfer (which themselves strictly depend on the relative positions of other molecules' position and orientation). Each system of molecules studied therefore has to be parameterised in isolation: quantum chemistry packages such as the Versatile Object-oriented Toolkit for Coarse-graining Applications (VOTCA) [89] or Amsterdam Density Functional (ADF) [90] should be used to find the HOMO and LUMO energies for each molecule/conjugated region in their environment (and also to obtain $|J|^2$ for the Marcus rate) and each charge transport simulation would then represent that one state. If one takes Marcus theory literally, the minimised state of the potential landscape that contains the reactants and products for that process (with the ion on one molecule and another) would need to be found and used to calculate ΔG^0 before every transfer process. This is perhaps too computationally and practically arduous to perform before each charge-transfer process and, instead, some approximations and more tractable disorder models are often applied.

2.3.1 Static Morphology and Dynamic Dipole Approximation

The first assumption often made is that the structure, on which the charge transport simulation is being performed, is completely static throughout the simulation. This is similar to the Born-Oppenheimer approximation, where the nuclear motion within the molecule can be treated separately to the electronic motion. Marcus' theory implies that a geometric relaxation after transfer should occur, however this approximation drastically speeds up the computation. As the molecules are often residing close to a stable (or metastable) potential minima and charges hop on the femtosecond timescale, this adiabatic picture of charge transfer is usually retained in the literature [91, 92, 71], although non-adiabatic methods exist and are described in recent review articles [71, 92]. In this approximation where the charges move instantaneously compared to the molecular motion, there is no need to consider the effect of molecular ions on the surrounding geometry and a single minimised state or an average of thermalised states can be used to find the electronic couplings and the HOMO and LUMO levels of each

molecule [91]. It is also often assumed that the reorganisation energy for the transfer processes is the same within a single material: $\lambda_{ij} \rightarrow \lambda$. These approximations will also be made in this work.

A consequence of the static nuclei assumption is that the molecules are not able to undergo any dynamics during the charge transport simulation. The translational and rotational vibrations, representing the thermal energy of the molecules, no longer occur. In simulations where charges move on longer timescales (i.e. exciton transport under FRET) the dynamical processes are no longer included in the simulation. This means that the dipole alignments, which govern the efficiency of exciton transfer between molecular dipole pairs, become static. To represent the motion of the dipoles, the dynamical dipole approximation is made which assumes that the molecules can rotate freely and on a timescale faster than the transfer processes, such that the average dipole orientation factor, κ, can be used between all pairs of dipoles. This means setting $\kappa^2 = 2/3$. This approximation (or one assuming perfect alignment) is taken widely in the literature [78, 93, 80, 94]. Other researchers have taken to using the system-average dipole orientation [84, 85]. However, the true representation requires the relative dipole orientation between the pair undergoing transfer to be considered, the effect of this model compared to a system-averaged value has been discussed in recent work [81, 83]. This latter model is more appropriate for a system where rotational dynamics is hindered and the system-average model appropriate for a freely-rotating system (a non-viscous liquid) or a highly-ordered lattice. The dynamic dipole approximation will be made in this work and compared to a static dipole model with complete alignment in chapter 7.

2.3.2 Gaussian Disorder Model

Due to the computational intensity or calculating the electron wavefunctions (or approximating them through parameterisation) of each conjugated region in an organic morphology and the difficulty in determining the effect of the thermal disorder on such a property, the HOMO and LUMO levels for each molecule are often approximated by a Gaussian distribution. The Gaussian disorder model [95] takes the energy of each conjugated region from a Gaussian distribution of width, σ around a mean HOMO/LUMO level (with the mean value calculated in vacuum conditions at the DFT level) for that material. The width of the Gaussian density of states must be either assumed, taken from literature or calculated beforehand (using quantum-chemistry methods or inferred from experiment) and a combination of these approaches has been previously used in

simulations using all three types of hopping models considered here [96, 74, 19, 80, 6, 97].

Use of a Gaussian disorder model, therefore can allow the HOMO and LUMO levels of a molecule to be determined for a given simulation without quantum-chemistry calculations. This has enabled disorder to be captured in on-lattice charge transport simulations. Of course, it is important to average over many simulations, each assigning energies from the same Gaussian distribution for an average to be obtained. Other versions of this disorder model such as the correlated Gaussian disorder model [98, 77] which introduce correlations in neighbouring HOMO/LUMO energies and the extended correlated Gaussian disorder model [99, 75] also exist but will not be considered further. In this work, the energetic disorder will only be captured by taking a value of σ from the literature for the relevant material and no correlations will be introduced.

2.4 OSC Devices

As we have seen already, organic semiconductor devices differ in terms of their electronic states in a significant way compared to inorganic semiconductors: charges do not travel through the device in band-like motion and do so instead by thermally activated hopping (or some 'semi-band-like' transport in polymers and crystals) [11, 26]; an electron, excited to a higher state, cannot do useful work until it is separated from its associated hole; the energetics of the OSC material is much more varied than an inorganic crystal. These differences have all contributed to changes in the typical device architecture of OSCs compared to regular semiconductors.

The most notable architectural change is the emergence of the bulk hetero-junction. This architecture involves blending two materials together (with an offset in HOMO and LUMO energies) such that exciton separation into free holes and electrons is enhanced. As an exciton only exists for a relatively short lifetime, τ, covering a diffusion length, L_D, in that time, it is important that an exciton is formed close to an interface with a neighbouring material. This places a restriction on material domain sizes within the mixed film, forcing them to be smaller. However, once the charges separate it is also important there is a pathway of their respective lower energy states (a contiguous route through a single material) to the device electrodes so that useful work can be done on an external system. If a charge cannot escape it then hinders motion of other charges or undergoes recombination with an opposite charge to form an exciton - which can then decay and release the absorbed energy. This limits the fineness of domains and places a lower limit on domain size. These two constraints need balancing to optimise

the efficiency of an OPV or OLED device [5, 19].

2.5 Structure of Organic Films

The morphology of organic semiconductors can vary depending on the processing methods and techniques as well as the molecules and blends of molecules used. A spray deposition of Phenyl-C_{61}-butyric acid methyl ester (PCBM) can result in an ordered, even structure which, as an electron-transporting layer, has enhanced efficiencies of perovskite solar cells [100]. A solution deposition process such as spin-coating, followed by a high temperature annealing stage (in P3HT-PCBM blends for instance), can also improve transport performance (compared to non-annealed films) by allowing molecules to re-arrange into a crystal or semi-crystalline structure where charge delocalisation (alluded to in section 2.1.2) increases charge mobility [24, 101, 102]. However, processes with a high solvent evaporation rate without annealling can produce an amorphous structure in many OSC materials and annealling, while enabling crystallisation, can also induce phase segregation and lower performance [103, 104, 105].

In early work on organic semiconductors, crystalline structures were seen as essential for finding charge mobilities that could compete with silicon, for instance. However, in recent years, organic semiconductor devices which are amorphous in structure have been created that have charge mobilities that can compete with amorphous silicon (some crystalline materials can exceed amorphous silicon) [1] and low-disorder polymers such as IDT-BT (indacenodithiophene-co-benzothiadiazole) and derivatives, have shown high charge mobility thanks to fast intra-molecular transport [106, 96]

Polymers can also introduce an extra degree of complexity into the characterisation of the structure of organic semiconductors. Polymers can exist in the amorphous phase (where chains have uncorrelated positions and orientations with their neighbours), the crystalline phase (where all chains are strongly co-aligned and regularly spaced (such as in cellulose, and sometimes P3HT, which often forms fibrils)[105, 107, 108]) and also a semi-crystalline phase where a small number of neighbours will be aligned and the rest may still be highly disordered [57, 109, 110]. If one were to hope to simulate the charge-transport properties of organic semiconductors that are polymeric or polymer-fullerene blends, for instance, one needs to capture the right degree of crystalline-to-amorphous content and their interfaces as well as capturing semi-crystallinity (often characterised by the paracrystalinity in a given plane - for instance, the degree of π-π stacking [111, 112, 113, 114]) and the relevant length scales of and between the different phases

[110].

2.6 Simulations of OSC Transport

Before a charge transport simulation can begin, the morphology on which it is to be performed must be created. In many charge-transport simulations these are often represented as regular cubic lattices, including in simulations designed to study amorphous/semi-crystalline systems, with periodic boundary conditions in all dimensions or all but one dimension which represents the contact electrodes [80, 18, 6, 97, 5, 19, 88, 74, 115, 45, 75, 93] . The transport in these systems is commonly modelled using one of the hopping models described in section 2.2.1 (often Miller-Abrahams or Förster, as there is often no atomistic molecular morphology from which to calculate electron orbital overlaps), between nearest neighbours or neighbours within a cut-off radius. The value of α (or R_o in FRET), the spacing between sites and the distribution in site energies and, in some cases, the combination of materials, then determine the layer/device transport properties. In some cases, an atomistic morphology is generated using methods such as molecular dynamics, coarse-grained molecular dynamics or Monte Carlo methods [57, 113, 116, 35, 109, 36, 117, 118, 119, 120, 96, 106]. The morphology can then be used to calculate electronic orbital overlaps and the parameters for the Marcus Theory rates (ΔG^0 and $|J|$) calculated in quantum-chemistry packages. Other charge-transport methods relying on atomistic morphologies include quantum-mechanical approaches such as Flexible Surface Hopping (as previously discussed) and even methods involving non-adiabatic transfer processes have been used [92, 9, 11, 87].

When performing charge-transport simulations there are a few quantities of interest which are commonly calculated and provided to those attempting to model the performance of organic devices. These are generally the diffusivity/diffusion constant, D, which describes how quickly charges can navigate the organic film, the charge-mobility, μ, which describes the charge response to the electric field, and the density of states of the electronic energy levels. The latter can be provided through the quantum chemistry calculations performed on an atomistic morphology and the diffusion constant and charge mobility need to be inferred from simulation of charge motion through the organic morphology.

Simulations of hopping charge-transport processes in organic semiconductors are very often performed by one of two methods. The first is the kinetic Monte Carlo (KMC) approach (performed in one of two numerically equivalent flavours) which considers

the individual hops of individual charges [5, 19] and a master-equation approach (ME) which considers the flow of a concentration of particles [75, 45]. Both methods are able to calculate properties such as the steady-state distribution of charges and responses to electric fields, however the KMC method requires sampling of the carrier motion until approximate convergence to obtain some of this information while the ME approach can calculate the exact steady-state for a given transport model. KMC is perhaps better suited to situations where consideration of individual charge trajectories becomes important while the ME approaches are more suited for quickly measuring properties such as the diffusion coefficient and charge mobility. For this reason an ME-type approach will be used in this work and the full description of the method is provided in section 3.4.

There are some further constraints on the charge-transport simulations (in-part, already alluded to) which must be considered before accepting a predicted value of the charge-transport properties produced by either of these approaches. The first is finite size effects. One such effect is that if the system is very small and under periodic boundary conditions, it is much more likely that a percolation pathway (a route of very high reaction rates and therefore probability of hopping) exists that spans the simulation volume, increasing the predicted charge diffusivity/mobility [45]. If the system is too small there is then also a poor representation of the Gaussian distribution of the density of states of the HOMO/LUMO (it is likely that low-energy tail states, which charges are slow to hop away from due to the Boltzmann factor in the Miller-Abrahams and Förster rates and the ΔG^0 term in the Marcus rate, are missing from the charge transport simulation) which also increases the predicted rate [43]. Beyond these limitations, one also needs to perform a sufficient sample of transport simulations (each using a new set of energies from the same Gaussian density of states) and, in some instances, may also need to perform these transport simulations on a sufficient number of atomistic morphologies to obtain an accurate picture of charge-transport of a given material [40].

2.7 Simulation of OSC Materials

Moving away from lattice-based charge-transport simulations towards atomistic amorphous or semi-crystalline morphologies brings obvious benefits: it allows the simulations to better capture the structural disorder and explore different length-scales where order may exist. Capturing the balance of order and disorder, something which can only be predicted through generation of atomistic morphologies, then allows the charge-transport simulations to capture the fast straight intra-chain and crystalline transport

and the slower amorphous regions that connect them [110] in a way a lattice model cannot.

There are already some conventional ways of creating an atomistic morphology of OPV materials *in silico*. The, 'go-to' method, is to perform a molecular dynamics or Monte Carlo simulation. These take a set of atoms, a set of forces that describe inter- and intra-molecular interactions and calculate moves based on an energy condition as in Metropolis Monte Carlo (MMC) or by using force-velocity information to calculate subsequent steps in the particle trajectories as in the velocity verlet algorithm implemented in molecular dynamics packages. By setting a temperature (and pressure or volume), the move accept/reject condition (in MMC) or the particle velocities (in MD) are adjusted, biasing the probability of visiting different states during a simulation. The force-field for an atomistic (sometimes all-atom) MD/MC simulation is often parameterised from DFT-level calculations (which involves observing the energetic response when an atom is moved) or taken from a generalised force-field for the relevant class of material, such as the General Amber Force-Field (GAFF) [121].

Evolving the system along a trajectory in MD or determining the next Monte Carlo move to accept/reject can be a computationally expensive process, especially as the size of the system being studied increases [35]. Not only does it take longer to reach equilibrium, but it takes longer to find a new state that represents a new configuration (in another basin in the energy landscape) rather than just thermal fluctuations [35]. In some simulations it can take tens of millions of time-steps to obtain one new statistically independent state in a system of 50,000 atoms and this increases with system size [35]. This makes it difficult to sample the equilibrium distribution of states. It is therefore common for simulations of OPV and OLED materials to be approximately 1,000 molecules (or repeat units) in size and only a small sample of states considered or little consideration given to the ensemble [24, 118, 122, 116, 123].

To overcome this computational limitation on system size, researchers have turned to coarse-grained simulations, including coarse-grained molecular dynamics (CGMD) [109, 57, 116, 113, 36, 123, 124]. CGMD lowers the number of particles in the system by instead representing a group of atoms (and the forces acting between them) by a single large particle, or bead, interacting with other representations of groups of atoms. This effectively reduces the number of atom velocities and accelerations that need to be calculated at each time step, lowering the computational load or enabling the study of longer polymer chains and larger systems [124, 35]. While in an atomistic MD simulation, the time-step might be advanced by a femtosecond (to prevent instability) in coarse-grained simulation this time-step can be much larger, sometimes as high as

a picosecond [125, 35]. This coarse-graining, however, also alters the dynamics of the simulation and measures such as the density, melting temperature, glass transition temperature and other observables have been known to change between the atomistic and the corresponding coarse-grained approximation [126, 127] and the quality of the coarse-graining is therefore measured in comparison to the ability to recreate some physical property, such as density [109].

2.8 Glassy Dynamics

When performing atomistic or coarse-grained simulations of organic semiconductor morphologies, some physical effects which are present in an experiment will also be found in the computer experiments. One of these, with a major consequence for the required simulation time, is dynamical slowdown [40]. This is perhaps best explained by considering the inter-atomic potential and the associated potential energy landscape of a simulation beginning in the gas phase.

The potential energy landscape describes all of the possible configurations of the system in $3N$-dimensional space (where N is the number of atoms), the associated potential of each state, and can be used to describe transitions between states [128]. In the gas-phase this landscape is smoothly varying with atomic coordinates, with high-energy states occurring only when atoms are in very close proximity to others (with separations in the repulsive region of the chosen force-field). If the force-field has an attractive region (e.g. pairwise Lennard-Jones potential where the attractive region is negative in energy at $r_{ij} > \sigma_{ij}$) then low energy states will occur when neighbouring atoms are in this region. It should be recalled that, at a given thermal energy, the chance of visiting a state that is high in energy is generally lower than the probability of finding a low energy configuration and (in a smooth landscape) if the system has been left for a long enough period of time the time average will represent the ensemble average. This statement describes the thermodynamic equilibrium of an ergodic system [129]. As the system gets denser (at a given pressure and temperature the box may reduce in volume), the number of states at high energy (say, above the thermal energy) will increase in-proportion to those low in energy. To get between the most likely-observed states (the minima in this energy landscape), the system must now pass through higher energy states (the barriers or saddle-points) more often, and in liquids, the rate at which this happens is often well described as being Arrhenius [130]. At very high density, these barriers dominate and, at temperatures below the glass transition temperature, the relaxation time (measured through decay in polarisation or spatial correlation,

for instance) becomes longer than is experimentally observable. For all intents and purposes, the structure may be considered solid [131, 130, 132]. If the crystalline phase is avoided (by cooling the system rapidly enough) and the system has become a viscous liquid, the relation time can show sub-Arrhenius temperature dependence (the activation barrier itself becomes dependent on temperature) [130].

The properties of a glass are known to depend on the exact processing method (for instance, the glass transition temperature depends on the initial cooling-rate [130]). This is because, during glass formation, the time the system has to sample new configurations (or states in a given potential energy landscape) depends on the temperature and therefore the cooling rate. In a very slow cooling system beginning at a temperature above T_g, the system should have enough time to establish thermodynamic equilibrium, however, at much faster cooling rates the time provided to find equilibrium is diminished and a non-equilibrium state may be found with different properties. As described by Debenedetti & Stillinger [130]

> "This falling out of equilibrium occurs across a narrow transformation range where the characteristic molecular relaxation time becomes of the order of 100 seconds, and the rate of change of volume or enthalpy with respect to temperature decreases abruptly (but continuously) to a value comparable to that of a crystalline solid. The resulting material is a glass."

With many commonly used organic semiconductors having a glass transition temperature at or above room temperature (e.g. PCBM: 118.3 to 132.2°C; P3HT: 40°C; PCPDTBT: 130°C; α-NPD: 88.9°C; TCTA: 144.9°C), relaxation times, inaccessible to a computer, let alone experiment, can also occur in organic semiconductor simulations [133, 134, 135, 136, 137]. This means that an atomistic molecular dynamics simulation, starting from the gas phase and attempting to find a morphology at thin-film density, will not only struggle to reach the an equilibrium density, but will likely find an out-of-equilibrium state and struggle to undergo relaxation and sample further configurations in other basins - doing so on longer and longer timescales. Given that molecular dynamics simulations are run on the femtosecond timescale and relaxation times can be hundreds of seconds, it is simply too computationally expensive to use molecular dynamics as a tool to sample non-equilibrium states or find the most likely equilibrium state in the potential energy landscape of many amorphous organic semiconductors.

The mechanically stable set of states are the minima in this potential energy landscape. These minima of amorphous configurations have a large variation in depth (and there is often a deep minima, sometimes the global minimum, that represents a crystalline

state). Each minimum is surrounded by a basin of states where the surrounding states are all higher in energy. If, in one of the higher energy states of a given basin, all the particles were allowed to come to rest, the structure/configuration found would be the same as that of the minima of the basin. The structure at the potential minima is therefore referred to as the inherent structure - the configurational structure with the additional vibrational/thermal disorder removed (by a minimisation or quenching, for instance) [138].

Inherent structures become useful means of describing the configurations of an amorphous system at constant density by providing a new statistical-mechanical basis (a partition function over only the potential energy) of the system [139]. Therefore one way of representing the structure of the macrostate on a computer simulation would be to capture the distribution (or a representative distribution) of inherent structures in the energy landscape and weight each states' contribution to a structural property according to this new partition function. They also bring practical benefits, such as providing sharp features in structural measurements such as $g(r)$ (the radial distribution function) and $s(q)$ (the structure factor) without needing to average over many states due to the lack of vibrational disorder [139, 140]. In terms of organic semiconductors, using the inherent structure also provides self-consistency as the Marcus theory picture relies on a minima in energy of the product and the reactant states. To determine the charge transport properties using a Marcus transfer model, the inherent structures would then be used as the atomistic morphologies which a quantum chemical transfer integral calculation and a subsequent KMC or ME simulation performed and their contribution to a physical property (e.g. diffusion constant) weighted according to the new partition function in potential energy of the inherent structure.

2.9 Alternative Simulation Methods

The problem remains of finding these different inherent structures given the slowing down of relaxational dynamics in dense amorphous packings, including those of organic semiconductors. Typically an MD simulation is performed with a high-temperature annealling phase, then followed by a cooling phase (often a nanosecond to tens of nanoseconds in length) to reach room temperature and pressure, and a sampling phase where frames of a simulation are captured periodically (often once a nanosecond or tenth of a nanosecond) [141, 57, 116, 118, 24]. While this annealing phase helps avoid some of the out-of-equilibrium basins, the room temperature simulation does not allow large-scale rearrangement or the escaping of deeper minima. For instance, if a semi-

crystalline region is found, it is unlikely it will break and each of the frames included in the average measurement will contain this feature. The sampling is then instead limited to sampling within a single basin or a small set of basins with the same or similar inherent structures.

This restricted sampling occurs because MD and MC are thermodynamic methods [142]. Slow rate (long-timescale) processes such as the escape of energy basins in the potential energy landscape are less likely than high rate (short timescale) vibrations within the basin. If one is interested in sampling many basins or many minima, then a different method which either alters the energy landscape so as to make it thermodynamically more likely a barrier/saddle-point is crossed or a method which biases away from thermodynamically likely states is required. Methods such as the minima-hopping method of Goedecker [142], which discourages the system from returning to previously visited states, have been shown to be effective at finding the structure of the global minimum of potential energy. Other methods such as metadynamics, described by Laio & Parrinello [143], where the currently occupied basin is 'filled' with additional potential are designed to increase the rate at which barriers are crossed. In the simulation of glasses, procedures such as Swap Monte Carlo where two different types or species of particles are swapped with an acceptance rate like that of the metropolis algorithm has also been shown to 'avoid' the glass transition, increasing thermalisation/sampling rate at a given temperature compared to MMC which does not include swap moves [144, 145]. Other techniques such as coarse-graining also alters properties such as the glass transition temperature and increases sampling rates.

For sampling structures of organic semiconductors, a desirable method would therefore be one which is able to overcome the energetic barriers between different potential minima very quickly and in an unbiased way such that, with a large enough sample and after performing a minimisation, the distribution of the potential energy of the simulated inherent structures represents the real density of states of inherent structures. It is not necessarily desirable to find the lowest energy configuration (like in the method of Goedecker) if the fabrication process of the device that one intends to model does not include a high temperature annealling process. It is also clearly undesirable to be thermodynamically biased to low energy states as this still experiences slow relaxation [142]. Using MD in NPT conditions followed by a high temperature anneal in NVT conditions and performing periodic minimisations is still hindered by the slow-down experienced while the system is yet to reach its physical density (this author is yet to reach equilibrium density for simulations in the range of 0 to 300 K and at 1 atm for a system of 1,000 PCBM molecules after 160 CPU hours on a 16-core node on a

supercomputer). A method where structural relaxation is not constrained by the shape of the energy landscape is therefore highly desirable.

Lessons from simulations of proteins

In simulations of proteins, long-dynamical processes have also historically created significant limitations on the processes that could be modelled [146, 147, 148]. Proteins are very large organic molecules consisting of one or many long chains of amino-acids (they are polymeric) which, due to its size, exhibit dynamics on much more disparate timescales than those of organic semiconductors [149, 150]. The high-frequency atomic vibration is similar, but the oscillations of these large structures or parts of these large structures —events such as protein folding— are performed over timescales of milliseconds [151, 148, 146]. However, an atomistic molecular dynamics simulation must be performed on time-steps of femtoseconds and a coarse-grained molecular dynamics simulation still requires roughly a picosecond time-step (although strictly it depends on the coarseness of the groupings) [113, 57]. For those considering atomistic molecular dynamics, MD simulations were often performed on relatively small systems for durations of hundreds of picoseconds to tens of nanoseconds [152, 153, 154]. Simulating events like evolution of the active pathway and protein folding has therefore been a significant computational challenge.

In more recent years, researchers moved to other methods which have allowed these long-timescale processes to be studied. Some early approaches involved forcing the rotation of torsional angles to sample new conformations [155, 156]. Other researchers used multi-canonical molecular dynamics to enhance sampling rate [157, 158]. Another approach was to identify the fully-constrained atoms [159] within the proteins' structure (constrained by bonding and, later, other electrostatic forces) and treating sections of the protein as rigid. In some work, this rigidity classification was performed using the pebble game rigidity analysis algorithm between atoms and across bonds to determine which atoms were fully constrained by their neighbours [160, 161]. One such package, Floppy Inclusion and Rigid Substructure Topography or FIRST, also attempts to identify when non-bonded forces acting between two atoms (such as hydrogen bonds) are sufficiently strong that their separation is unlikely to vary [161, 162]. In these scenarios where atoms are constrained by bonded and non-bonded forces, their relative positions are maintained by forcing subsets of atoms to be within a small distance from their 'ghost'/template positions. These templates are re-positioned and rotated to minimise the squared distance between template and atom before applying any distance

constraint [163, 164, 148, 165]. Subsequent large-scale net displacement then occurs without changing the relative separation between atoms that are members of the same template. These templates, effectively, become the coarse-grained units which relative motion occurs between. The coarse-grained simulation performed in this manner was implemented in Framework Rigidity Optimized Dynamic Algorithm (FRODA). Using FRODA, random atomic displacements were used with templates that spanned large areas of the protein, with large-scale diffusive motion observed across the protein structure [148].

A subsequent expansion to FRODA with relevance to this work is the elastic network model (ENM) [166, 167]. In the FRODA-ENM model, the random motion is replaced by directed motion. Normal mode analysis is performed between C_α chains or other large-scale groups and the large-scale dynamics (the functional motion of the protein) induced by displacing the coarse-groups along the direction of the low-frequency modes. This method of combining both atomistic structural information and coarse-grained dynamical information does not consider or encourage high-frequency motion in the system (such as atomic vibrations and changes in bond-lengths or angles) however these contribute an insignificant amount towards the large-scale, long-time motion, which ENM methods do encourage [147, 166, 168]. Normal-mode analysis and studies of the low-frequency motion have therefore been used extensively in protein simulation [169, 147, 170, 166, 167, 171, 172, 173, 174, 175].

2.10 Low-frequency Modes

Low-frequency modes have also been widely studied in amorphous, glassy and jammed systems [176, 177, 178, 179, 180, 181, 182, 183, 184], in systems of ellipses/ellipsoids [185, 186], in spheres undergoing active oscillation [187], under shear [188, 189] and under compression [178, 190]. Links between the physical structure and the vibrational modes as well as the collective relaxation and the low-frequency modes have been made extensively [182, 183, 180, 181, 176, 178, 179, 190, 191, 188, 192, 193]. Common to both studies of glasses and protein structures is the concept of rigidity, in some cases, also identified through playing the pebble game on a network of constraints [193] or, more generally, counting the number of contacts in a system [176, 194, 179] and comparing to the Maxwell criterion for rigidity [176, 195].

The Character of Vibrational Modes

The low-frequency modes describe the more collective vibrational motion. By 'collective' I mean that the direction of motion and the amplitude is spatially correlated - the neighbouring points will move in a direction and at a speed, reaching an end point, which is more similar than that of further particles. For the most collective motion this will encompass more neighbours than less collective motion. Collectivity is often assessed using measures such as the participation ratio and the mode localisation [196, 197, 187, 180, 181, 182, 183, 192]. The participation ratio is calculated as

$$p_m = \frac{1}{N} \frac{(\sum_i^N \|\boldsymbol{u}_i\|^2)^2}{\sum_i^N \|\boldsymbol{u}_i\|^4} \tag{2.6}$$

where N is the number of particles and $\boldsymbol{u}_i$ is the displacement vector for particle i in a vibrational mode. This describes the similarity in the distribution of the length of the displacement vectors of a given vibrational mode, with the most collective motion (here, the most uniform amplitudes/displacements) having a value of 1 if all particles have the same displacement in a given mode and a value of $1/N$ when only one particle has a non-zero displacement vector. This measurement alone is therefore able to distinguish between different types of low-frequency modes: the collective low-frequency modes where displacement is often system-wide [183] and the soft-modes which are localised or semi-localised and occur when a given particle has too few contacts with neighbouring particles [176, 194]. A localised soft-mode can be created when an isotatic system (a system considered rigid by the Maxwell criterion with the average number of contacts per particle, $Z \geq 6$) has enough contacts removed such that $Z < 6$ or the number of local contacts on an individual particle is less than 4 (the local criterion for rigidity) [176, 195]. This soft-mode would then be identified by a low frequency and a very low participation ratio [176].

These low-frequency modes and soft modes (they will be distinguished in this way throughout this work) have been studied in relation to their role in structural relaxation processes. In a series of papers by Widmer-Cooper et al., the low-frequency modes and, later more specifically, the soft modes were liked to 'irreversible structural reorganisation' in 2D systems of super-cooled liquids [180, 181, 182, 183]. These soft-modes were therefore seen as an indicator of structural relaxation - the system was able to find new configurations following the detection of a soft mode. These have also been described as 'failure points' where a system under shear stress is likely to give way due to their indication of mechanical instability (they can also be induced by the shear

stress) [194, 193, 188]. The soft modes are considered excess modes (beyond that of the Debye model) and exist in the shoulder/plateau in the vibrational density of states (the Boson peak) which can be continuous in frequency down to zero-frequency, when a system is near jamming [178, 176, 192, 132, 193].

The vibrational modes are regularly studied through the creation of a Hessian matrix (although sometimes the modes are inferred from particle dynamics in molecular dynamics simulation) [189, 192, 194, 178, 198, 147, 187]. The Hessian describes the curvature of the potential in space around each particle and can be calculated from the second derivatives of the potential with respect to each of the Cartesian components, which is often approximated as being harmonic. The Hessian matrix where elements are weighted by their masses leads to the dynamical matrix which can be used to find the Newtonian vibrational dynamics (within the harmonic approximation). The diagonalisation of this matrix produces the square of the vibrational frequencies (the eigenvalues) and the displacement vectors of the corresponding normal mode (the eigenvectors) [198, 166, 147]. The instantaneous normal-modes describe the curvature for any given structure and, if particles are not at their local minima in the potential field, imaginary modes will exist that define the directions to minima (the mechanical equilibrium structure) [199]. When particles are at equilibrium separations the normal modes will all point to increases in energy and only positive real modes will exist, as well as the zero energy modes of uniform translation (one per dimension) and the soft modes if there are too few contacts and a low/zero energy displacement exists [199, 176]. The low-frequency modes produced by this matrix are usually high in participation ratio, and therefore more collective, than the modes at the high-frequency end of the spectrum [184, 192, 197, 187]. This is not necessarily always true however, as is the case of the modes in the Boson peak [192].

Collectivity can also be measured in terms of the local average dot product [187] and the phase quotient [197, 184] and other measures such as the stretching character. The average local dot product, here

$$\langle u_i \cdot u_j \rangle = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{N_c^i} \sum_{j=1}^{N_c} \hat{\boldsymbol{u}}_i \cdot \hat{\boldsymbol{u}}_j \tag{2.7}$$

where N_c^i is the number of interacting neighbours surrounding particle i and $\hat{\boldsymbol{u}}_i$ is the unit vector of particle i in a given mode. This therefore represents the local directional collectivity between neighbouring particles and does not consider the relative amplitude between a particle and its neighbours. The average local dot product produces

values from 1 to -1 where a value of 1 represents all of the particles moving in the same direction and -1 all particles moving in opposite directions to their neighbours, with 0 representing a lack of local collectivity as would be the case in a large system and averaged over all particle pairs or in an ideal gas. This measure, along with the stretching character (which represents an amplitude weighted version and measures the amount of changing spring length) [184], allows classification of the behaviour of the vibrational modes. In the case of the linear-chain spring model with a basis, there is a clear division in the type of behaviour of the modes with distinct acoustic and optical branches producing highly collective and in-phase (acoustic) and less collective and out-of-phase (optical) motion [200]. In this case, the optical motion would represent a negative average dot local product while the acoustic modes would produce highly positive values (both of which become less collective with increasing frequency).

The analogy of the linear chain model does not apply itself well to amorphous/disordered solids or liquids, however [201]. There is a distinct lack of periodicity in an amorphous system and the phonon dispersion and group velocity are ill-defined [201, 192]. In an article by Silbert et al., each vibrational mode was found to contain contributions of multiple wavevectors in Fourier-space [192]. In these amorphous systems, the modes are diffusons/locons and not generally described by travelling waves (propogons). Instead of considering amorphous *or* optical modes, the vibrational modes of amorphous systems are more commonly described as being amorphous-like or optical-like in character through measures such as the phase quotient [201, 184, 197]

$$\varphi_m = \frac{\sum_i^N \sum_j^{N_c} \boldsymbol{u_i} \cdot \boldsymbol{u_j}}{\sum_i^N \sum_j^{Nc} \|\boldsymbol{u}_i \cdot \boldsymbol{u}_j\|} \tag{2.8}$$

which is again measured between particles and all of their interacting neighbours. The phase quotient, like the average local dot product, produces values in the range of 1 to -1, however a value of 1 describes a vanishingly small frequency acoustic-like mode which represents a complete coherence of the direction of travel of the neighbours (generally found in simulations as the translational modes) while a value of -1 describes an entirely out-of-phase displacement of an optical-like mode [197, 201]. Distinct from the average local dot product which captures the average angular separation in displacement vectors, this measure describes the coherence of the vibration. In several studies, the low-frequency vibrations are modes with phase quotient that is very high, falling to lower values at larger frequencies [201, 184]. The low phase-quotient modes are also often accompanied by low participation ratio modes, indicating that the modes are

also localised or semi-localised in amplitude as directional collectivity is reduced [201], meanwhile the low-frequency (non-soft) modes are generally high in both participation ratio and phase quotient.

Low-frequency, collective modes and structural relaxation

The highly collective low-frequency modes, considered distinct from the soft modes, have been linked to structural relaxation in amorphous solids and super-cooled liquids [202, 187, 203]. An argument for this takes root in their being long-time period processes: the short-time processes may represent individualistic (uncollective) vibrations of within a configurational basin on the potential energy landscape, while the long time processes represent much larger-scale (collective) rearrangements of the system as a whole. While the high-frequency motion, occurring more often, may induce more local structural relaxation in physical systems, the low-frequency motion still plays a role and regions of high amplitude in low-frequency modes correlate with regions of strong relaxation of regions of the material [183, 203]. This reliable structural relaxation has been seen in comparisons of the participation ratio of the low-frequency modes and the areas where structural rearrangement was strongest in a variety of systems [183, 202, 187]. One such system using active oscillation at selected frequencies found fast relaxation and diffusive motion when enhancing the low-frequency modes in an athermal 2D jammed system [187], demonstrating the propensity low-frequency motion has for achieving structural relaxation.

In comparisons between participation ratios and frequency, localisation and frequency, and phase quotient and frequency, the low-frequency modes were the most collective and the amplitude of these collective low-frequency modes has been found on multiple occasions to correlate with structural relaxation processes [187, 182, 202]. Other studies have attempted to link structural causes to relaxation events or relaxation pathways [190, 194, 179, 191, 203]. Two relatively recent studies, one of elastic polymer rings at high density and one of soft repulsive spheres at variable density have shown how relaxation time can increase upon increasing the density, finding ballistic and diffusive motion respectively [191, 190]. This dynamical slowdown followed by a re-emergence of liquid-like dynamics is known as a re-entrant fluid-glass phase transition. This has been observed in soft systems and would appear to suggest that there exists a regime in which the dynamics are again able to relax some of the densest systems. A method such as that of Tjhung & Kawasaki that actively enhances the low-frequency motion [187] or a softening of the structure as with Gnan et al. [191] or softening of the potential

as performed by Berthier et al. [190] may therefore prove a useful way of sampling the configuration space of the energy landscape diffusively [187].

2.11 Summary

In this section, the requirements placed upon simulations of organic semiconductors, particularly those of amorphous systems, such that the physical reality could be captured were put plainly. Charge transport in organic semiconductors must be performed using an appropriate transport model which can vary from non-adiabatic quantum mechanical treatments [92] to point-to-point hopping/transfer models [2, 78] with a range of parameterisations in between [72, 67]. These simulations should be performed on a morphology that is representative of the system being studied which, in the case of amorphous materials, is often captured on small scales by a molecular dynamics or Monte Carlo simulation.

Organic semiconductor morphologies must also be generated in a way that captures the thermodynamics of the system being studied with considerations given to the fabrication process and the appropriate configurations on the potential energy landscape [138, 130, 35]. However, simulations of organic morphologies must also overcome the thermodynamic constraints of the energy landscape which is often exhibits slow structural relaxation at such high densities and, instead, consideration should be given to methods which overcome this slow relaxation [142, 138]. Methods such as coarse-graining are able to alter the dynamics such that relaxation is somewhat faster and the computational workload reduced. However, the parameterisation of a fully-atomistic system into a system of larger atomic point-clusters can also lead to artefacts in the structure (in the case of a Marcus theory transport model, these morphologies must also be converted back to an atomistic morphology for calculation of the transport rates too). This approach has been used widely as it allows the simulation of much larger systems, required to mitigate finite size effects in the subsequent charge transport simulations [43, 45] and to capture the variability in domain size required for simulations of exciton dissociation in a bulk hetero-junction, for instance.

Methods in the diverse field of soft matter, including protein simulation and glass physics, have also attempted to study the slow relaxation of structures. In proteins, the large-scale motion of the molecule(s) requires a timescale several orders of magnitude larger than the vibrations of individual atoms within the structure. Attempts to model the large scale motion have therefore attempted to 'freeze-out' the small-scale motion

by treating sections of the protein as rigid. Using low-frequency modes to determine the displacement of the larger-clusters, the large-scale motion can be further enhanced, effectively increasing the simulation speed. Similar approaches have been taken in simulations of amorphous, jammed (or very dense) soft spheres where the enhancement of the low-frequency motion (at the expense of the high-frequency) can enable diffusive motion of the particles [187, 191]. Other approaches also suggest that a soft potential enables the rediscovery of fluid-like motion following compression starting in the glass-phase [190].

Considering the importance of the contribution of the low-frequency modes to structural relaxation in non-ergodic dense amorphous glassy systems and their enhancement in simulations of glasses and proteins provides us with a way in which these organic semiconductor structures might be generated and sampled. Combining this with the benefit of a suitable coarse-graining scheme may also reduce the computational burden and improve sampling rate, as well as further enhancing the rate of structural relaxation. A combined scheme such as that of the FRODA-ENM method applied to dense amorphous packings of organic semiconductors under the fluid-like regime past the re-entrant glass-fluid transition, may therefore be a suitable approach worth studying further. This thesis will focus on the development and application of such a method applied to organic semiconductors and contain studies of their low-frequency vibrations and charge transport properties.

Chapter 3

Methodology

In order to study the electronic processes of organic semiconductors, we were forced to face the fact that current simulation methods are a computationally-arduous process and that the scale of any simulations that were to be performed would be limited, for reasons previously described. Combining a suitable coarse-graining and low-frequency mode displacements, inspired and informed by previous work on proteins and soft glasses, a new method for simulating morphologies of organic semiconductors was created.

The method, Simulation of Atomistic Molecular Structures using an Elastic Network, or SAMSEN, takes the wisdom and computational speed-up from the coarse-grained techniques of protein simulations while retaining its atomistic structure. It also gains the benefits of selectively displacing along the low-frequency modes in a soft potential regime where dynamics, and specifically structural relaxation, begins to speed up again. In this chapter, SAMSEN (the model and its implementation) will be described in great detail, outlining its relationship to the FRODA-ENM method and explaining the creation of an initial simulation state, the coarse-graining scheme, the construction of the Hessian and the displacement along the vibrational modes. The simulation parameters that control sampling performance will be described in this chapter and explored in the following chapter, providing the reader with a recipe for achieving an approximate structure of organic morphologies and exploring new configurations with computationally-fast relaxation rates.

The end of this chapter will then describe the details of the charge transport model that will later be used to calculate charge transport properties such as the diffusion constant of electrons and holes and the diffusion length of excitons.

3.1 From Molecules to Rigid Sections

3.1.1 Molecules, Polymers and Functional Groups

To begin, we must first design the coarse-graining scheme that shall be applied to the molecules. In the research on proteins, the large-scale motion that described the protein folding and the functional motion, could be well described as relative displacement between large clusters of atoms spanning large sections of the overall structure. To model motion on this timescale, atoms could be limited to approximately fixed distances from their neighbours within these larger structures and the vibrational modes and low-frequency displacements calculated between the geometric centres of these large functional groups.

The timescale disparity is smaller in smaller structures such as organic molecules, however their large scale motion also ignores individual atom displacements and vibrations and consists of movement between a molecule's functional groups. In molecules and polymers, the functional groups represent bonding motifs in their chemical structure such as alcohol, phenyl and carbonyl groups, and fullerene cages, for instance, often with well defined internal bond lengths and angles which are energetically expensive to change. If these groups are provided with enough dynamical freedom between one another (by not being constrained by double bonds or stiff dihedral potentials which limit rotational motion), then these functional groups with their own characteristic internal structure and relative displacement between them provide a strong analogy to the large-scale motion of the proteins.

3.1.2 FRODA

In the FRODA model, the internal structure of the larger groups of the proteins were maintained by a mismatch-fitting correction procedure which ensured that atoms had approximately static positions within a given functional group. This was performed by creating 'ghosts'/templates (here 'rigid sections') of the relative positions of groups of atoms within sections of the structure where the atomic positions were sufficiently constrained by the chemical bonding (and other strong interactions e.g. hydrogen bonds) of the molecule as determined by the pebble game [161]. Similar rigidity constraints will be considered here.

Throughout the simulation, the atoms will be displaced by either random motion [148], through the chosen low-frequency eigenmodes [166] or by interacting with other atoms.

To test the mismatch condition, the templates are then placed at the new geometric midpoint of the set of atoms in the corresponding rigid section and rotated following the procedure described by Wells et al. (2004) [164] and Wells et al. (2015) [165]. Briefly, this method utilises information about the relative displacements of atoms from the template to determine an optimal direction of rotation of the template. The rotation direction is determined and a rotation performed in a series of steps to iteratively find a least squares mismatch between the template and the real set of atoms. The implementation used here follows the same algorithm, was based upon the FRODA source code and written by Thomas McManus, and was then adapted for the SAMSEN implementation used here.

Once the optimal rotation has been found, a simple distance calculation between the current positions of the atoms and the optimally rotated templates can be performed. In FRODA, a bias is added to the next displacement of the atoms in the direction of their templates [148]. Therefore, at each step FRODA attempts to provide a restoring displacement so that the rigid sections are (internally) in their low-energy configuration.

This fitting bias was balanced against inter-atomic collisions, which act to distort the internal structure of neighbouring rigid sections but protect the effective volume of their own rigid section. In a similar manner, the collisions are determined by measuring the closest approach between two atoms and, if they overlap (if they are separated by a distance less than the sum of their van der Waals radii, specified in Wells et al. 2005 [148]), a bias is added to the next displacement to act to repel overlapping atoms.

In addition to the collision-fitting biases were contributions from the eigenmodes. In the FRODA-ENM model, the eigenmode is selected by the user and the normal modes calculated between the rigid sections (in the larger rigid sections of the proteins, the eigenmodes were calculated between C_α atoms) using the Hessian construction as described by Atilgan et al. [147] which uses relaxed springs with spring constants and masses set to unity [166]. The eigenvectors for that chosen eigenmode were calculated for the initial structure and the biases added to the atom displacements (according to their membership of the rigid sections, taking the average if an atom is a member of multiple sections). This eigenmode bias is unchanged in the FRODA-ENM model and the simulation ends when the protein can no longer move along the chosen eigenmodes without creating overlaps or mismatches larger than the mismatch and collision conditions.

3.1.3 Conversion to Rigid Sections

FRODA created its rigid section templates from pebble game rigidity analysis as implemented in FIRST (which also considered non-bonded constraints, such as hydrogen bonding) [161]. For the bonded-constraints, the pebble game is a method for working out how many degrees of freedom are removed from a given atom by forces exerted by neighbouring atoms [159]. This is in many ways similar to the Maxwell criterion for mechanical stability [195], but considers local interactions which a system-wide average does not [159]. The pebble game is also an explicit degree of freedom counting method, but is played by 'moving' degrees of freedom between neighbouring atoms. If a 'spare' degree of freedom cannot be removed from a section of the molecule (within the rules of the game), then that region is internally rigid and can move with respect to other rigid sections. The pebble game is described in more detail in Thorpe [159] and the FIRST implementation in Jacobs et al. [161].

A similar analysis can be performed for small molecules in order to inform which parts of a given molecule are unlikely to change their relative position (the bonds are approximately rigid rods/bars). However, a common set of rules emerge which, for the most part when considering bonded-forces only, can determine which sections are rigid. These, generally speaking, permit angular changes between functional groups within the chemical structure. Such examples are: fullerene cages and cyclic regions are internally rigid (reflecting the low energy of the delocalised structure), atoms between double bonds are rigid and prevent dihedral angle changes locally, single bonds fix neighbouring bond lengths but do not prevent changes in dihedral angle to an atom four bonds away.

In SAMSEN we consider rigid regions within the molecule to be those where internal deformation is not possible due to bonding constraints. Short lived perturbations such as bond vibrations are ignored and not considered deformation under this definition. We also do not consider other effects such as hydrogen bonding or electrostatics from formal or partial charges. Therefore, only a region where the bonding arrangement of atoms prevents deformation shall be considered internally rigid. This form of rigidity can be calculated using the pebble game upon the molecules network of covalent bonds, however in this work we use a simpler framework to determine which sections are rigid.

This picture of rigidity is similar to that of Thorpe [159] where we do not consider the precise temperature, entropy or free energy (as in Phillips [204]) but instead consider if a section of a molecule can undergo relative motion continuously with respect to a neighbouring region without experiencing any change in energy. When making this

consideration we are only calculating the potential associated with the bond length and bond angle and do not consider any form of intermolecular interaction. In these conditions it is sufficient to determine rigidity by counting the number of contacts (or identifying zero frequency modes) [159] and techniques such as the pebble game, which calculates the number of degrees of freedom in regions of such a network, apply. If the two sections can move freely with respect to each other, then the two sections do not form a single rigid section. If the two sections cannot, then they are considered a single rigid section. In SAMSEN this scheme of rigidity is (with a small amount of permitted bond stretching / bending) how we shall try to retain the molecular structure while permitting some flexibility which will become particularly important in polymeric systems.

In SAMSEN, molecules shall be created and coarse-grained with the following set of rules that attempt to create the chemically intuitive picture just described. Prior to the beginning of the simulation, rigid section templates are constructed from the molecular structure (which shall be considered the geometrically-optimised structure) using these rules and the templates shall remain unchanged. These shall be referred to in the simulation and used to test a mismatch condition between the current atomic positions and the template positions - the template is optimally rotated to minimise the square mismatches (for the precise method used in this implementation, see Wells et al. [164, 165]) and the difference in template to current atomic positions calculated and compared to a maximum distance condition (see section 3.2.3). If the maximum mismatch condition is not broken then the simulation proceeds. If the condition is broken, all atoms in the system are then returned to the positions of their (optimally rotated) rigid section templates. This method therefore allows rotational motion of a rigid section as a whole, but only places constraints on changes in bond length and angular changes of bonds within and between rigid sections, maintaining the small-scale geometry of the molecule. The dihedral angles between bond-pairs in different rigid sections are not constrained by the scheme described below and becomes the dominant form of intra-molecular structural change within SAMSEN (but could, in principle, be restricted for molecules where it is required such as polymers with a large dihedral potential between repeat units). These allow medium- and large-scale changes, such as the breaking of polymer head-tail configurations and the curling-up of polymers respectively, as have been seen in simulations of P3HT, for example [57].

The procedure for determining the rigid sections produces similar results to the methods described in Jacobs et al. [161] (detecting zero frequency modes and the pebble game) if one neglects non-bonded constraints. The five-step procedure published by Smith

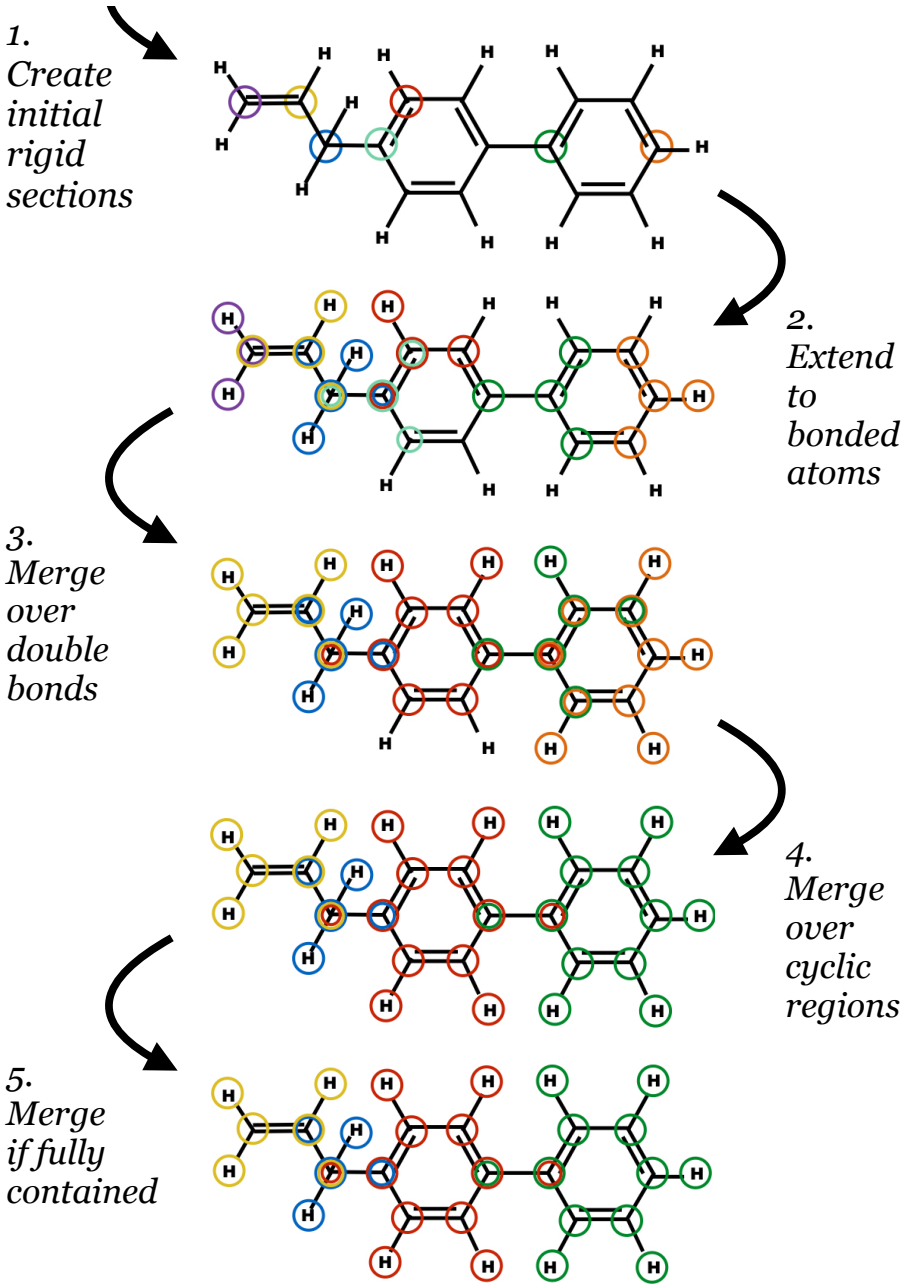

Figure 3-1: Step-by-step demonstration of creating rigid sections for an example molecule. In step 1, rigid sections are created containing each atom (excluding hydrogen) - a subset are shown here. In step 2, rigid sections extend to include bonded atoms. In step 3, rigid sections connected by double bonds are merged (one section now contains all atoms from the sections involved in the merger). In step 4, rigid sections in cyclic regions are merged if their inital atom is part of a ring. In step 5, any rigid sections fully contained by another are merged into the one that contains it - this step is unecessary in many molecules.

et al. [46] are reproduced here with modifications made for additional clarity and the process is shown in figure 3-1 for an example molecule:

1. Create initial rigid sections. Each atom in the system (excluding the hydrogen atoms) will initially form their own rigid section consisting of a single atom with a geometric centre at the midpoint of the atom. In figure 3-1 this is shown for a subset of atoms in an aromatic compound with rigid sections donated by the colour of the enclosing circle(s).

2. Extend to bonded atoms. Any atoms which are chemically bonded to an atom are also granted membership of the rigid section of their neighbouring atom and vice versa. Atoms can be members of multiple rigid sections. In figure 3-1 we see different rigid sections beginning to overlap with their neighbours as rigid sections expand to include bonded neighbours.

3. Merge over double bonds. Atoms that share a double-bond will have their rigid sections merged. One rigid section will now contain all member atoms from both rigid sections, whilst the other rigid section shall contain none and be deleted. In figure 3-1 we see the purple and yellow sections merge as well as some rigid section mergers within the cyclic regions.

4. Merge over cyclic regions. If a region of a molecule is cyclic, then all of the rigid sections within that cyclic region are merged, as described in step three (in the case of benzene, the entire molecule would now form a single rigid section). In figure 3-1 we see entire cyclic regions merging leaving one unhindered rotational degree of freedom.

5. Merge if fully contained. If, at this point, any rigid section in a molecule contains no unique set of member atoms (if all of its atoms are a subset of another rigid section's member atoms), then this rigid section is merged with the one(s) that contain it. In figure 3-1 we see no change to the number of rigid sections and this is common, however in the case of chlorobenzene the rigid section that begun on the chlorine atom would now be merged and a single rigid section would contain all atoms in the molecule. However, in many cases such as in figure 3-1, atoms will still be members of multiple rigid sections.

For molecules that produce multiple rigid sections, it is important in step 2 that atoms are members of the neighbouring rigid sections - this is what will maintain the overall molecular structure. By restricting only bond length, bond angles and not restricting the dihedral angles, the coarse-graining scheme effectively connects rigid sections by concentric hollow rods with a difference in radius equal to the maximum mismatch.

This picture describes the relative rotational motion that is available between two neighbouring rigid sections - there is completely free rotation about the shared axis of the rods (which run along the connecting bond), but off-axis rotation is restricted, limiting the angle that can be formed between the two rods (and therefore between the template positions for atoms that are members of both sections). The translations between sections are also limited by a value enforced by the maximum mismatch condition. This picture also holds for atoms within the rigid section as well, with an additional restriction that the internal dihedral angles are restricted by the number of bond constraints (the condition for rigidity, as would be determined explicitly by the pebble game).

The mismatch constraints, that hold the molecular structure together, will be combined with a collision/overlap constraint which will act to repel neighbouring molecules and create an excluded volume for the rigid sections that represent (or at least attempt to capture) the shape of the molecule in its current state. As we shall see later, these two conditions together will determine the structure of the organic morphologies in SAMSEN. The dynamics of the morphologies will be determined by the relative molecular structure, with the Hessian matrix constructed assuming a relaxed harmonic potential between the rigid section centres.

3.2 The Model

In this section I will describe the SAMSEN model used to describe both the structures and the dynamics. While this separation is not usually possible in conventional methods, these two components can, in a wide-range of circumstances within this model, be considered separate and will be referred to as the structural model and the dynamical model. The first will control the internal and external structure of the molecular morphologies, dictate how particles interact and therefore dominate measurements of the structure. The second will determine how collections of particles move and how quickly new morphologies are sampled, but, as we shall see later, will have little impact on the measured structure.

While this section will focus on the rules of the SAMSEN model, the following section will describe details of its implementation and setup procedure.

3.2.1 Simulation Volume

We begin by constructing a simulation volume with orthogonal sides of length L_x, L_y, L_z with periodic boundary conditions. In this volume we wish to place our molecules.

3.2.2 Atoms, Rigid Sections and Molecules

Each molecule consists of atoms which are grouped into n rigid sections in the manner described in section 3.1.3. The total number of atoms in the system will be denoted A and the total number of rigid sections denoted N.

Every atom a is a sphere with a position at point $\boldsymbol{r}_a$, a radius R_a, and a mass m_a. Atoms must be members of at least one rigid section. Each rigid section, s, is a set of p_s atoms. Atomic positions and masses give each rigid section a geometric centre $\boldsymbol{r}_s = \frac{1}{p_s}\sum_a^{p_s} \boldsymbol{r}_a$ and a mass $m_s = \sum_a^{p_s} m_a$.

The rigid section retains knowledge of the original relative position of its atoms throughout the simulation. The template position of atom a in rigid section s in the frame of reference of s is defined as

$$\boldsymbol{T}_a^{s'} = \boldsymbol{r}_a - \boldsymbol{r}_s \tag{3.1}$$

and is calculated once the molecules are loaded into the simulation volume. $\boldsymbol{T}_a^{s'}$, therefore, is a vector of constant length. We will use $\boldsymbol{T}_a^s$ to represent the template position in the global coordinate frame.

3.2.3 Conditions of a Valid State

All states that meet the following collision and fitting conditions are considered valid. If the conditions are met, the simulation can proceed to the next state. If not, the state must be modified until a valid state is reached. This modification is analogous to a minimisation, terminating at a desired tolerance. In this model, the minimisation is of overlapping volumes (ΔV_{ab}) and template mismatch ($\Delta \boldsymbol{T}_a^s$) and terminates at collision threshold χ_c and fitting threshold χ_f.

Put another way, this is how the structural model attempts to achieve states where the atoms are in positions which we estimate to be energetically likely at a given

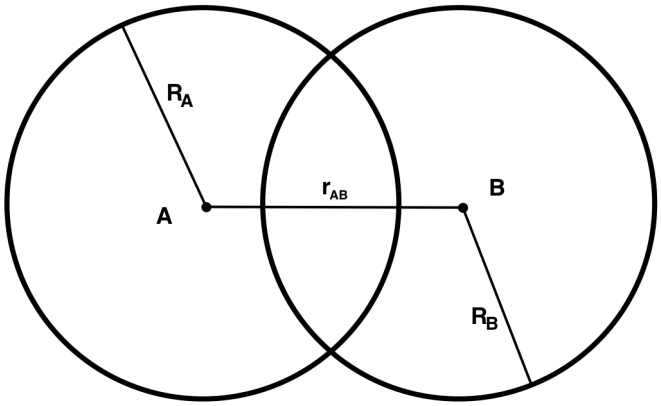

Figure 3-2: For a pair of atoms that are not members of the same rigid section, their separation r_{ij} must be less than the sum of the boundary radii R_a and R_B minus the collision threshold χ_c. χ_c is therefore length at which two spheres can overlap.

temperature —the molecules are not over deformed, the atoms are not too close to one another and the system is not in a state of very high potential. The minimisation of overlaps and mismatches are often competing and are here intended to reflect the competition between the external (Lennard-Jones) and internal (bond length, bond angle) potentials of organic systems in the General Amber Force-Field.

Collision Condition

Any two atoms a and b which are not members of a common rigid section cannot be separated by a distance r_{ab} that is smaller than the sum of their boundary radii R_a and R_b less a collision threshold χ_c. Therefore in each valid state

$$r_{ab} \geq (R_a + R_b) - \chi_c \tag{3.2}$$

or, considering figure 3-2, atoms cannot penetrate/overlap with another atom by a length greater than χ_c.

Fitting Condition

Any atom a that is a member of a rigid section s must be within a distance χ_f of its template position, $\boldsymbol{T}_a^s$

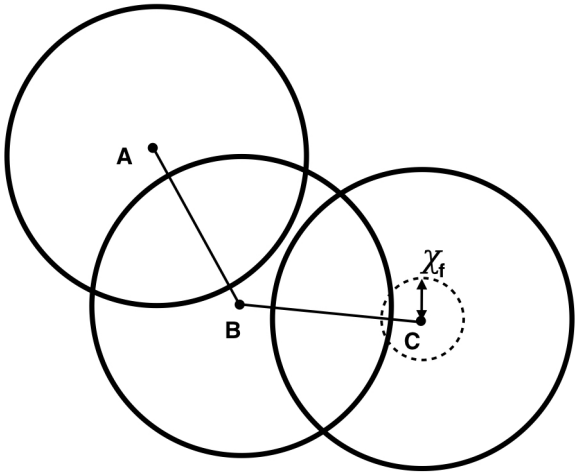

Figure 3-3: An atom must remain within a distance χ_f of its template positions in every rigid section in which it is a member. χ_f therefore controls both the maximum change in bond length and also the change in bond angle.

$$|\boldsymbol{r}_a - \boldsymbol{T}_a^s| \leq \chi_f \tag{3.3}$$

as shown in figure 3-3 for atom C in a rigid section containing atoms A, B and C.

Minimisation of Overlaps and Mismatches

To maintain these rules, at every simulation step, the system is checked for pairs of atoms breaking collision conditions. If a collision rule is found to be broken, then all overlapping atoms in the system (unless they are a both member of a common rigid section) are displaced such that their boundary radii are just touching ($r_{ab} = R_a + R_b$). If an atom has multiple overlaps, then the average correction vector is used.

Rigid sections are also checked to make sure no atoms break fitting conditions. The template $\boldsymbol{T}^s$ is positioned at the new geometric centre and rotated to minimise the squared displacements of atoms. The fitting conditions are then tested. If a fitting condition is found to be broken, all atoms in the system are moved to their template positions ($\boldsymbol{r}_a = \boldsymbol{T}_a^s$). If an atom is a member of more than one rigid section, it is moved to its average template position.

Collisions are checked by measuring the overlap length and atomic positions adjusted (if the condition is broken), then the fitting condition is checked by measuring the

maximum mismatch length and atomic positions adjusted (if the condition is broken). If either condition is broken and a correction applied, both collision and fitting conditions are checked again. This continues *ad infinitum* or until both conditions are satisfied system-wide. The state is then declared valid and the simulation may proceed.

Note on the resulting structures

Enforcement of the collision-fitting conditions produces a system where atoms deviate from their template positions by an amount depending on the local environment, and pairs of atoms overlap depending on how far they have strayed from their template positions. This process attempts to balance the two globally and, if the thresholds χ_c and χ_f were set to zero, would attempt to find configuration with the minimum deviation from template positions and external atomic overlaps. As shown later in figure 4-14 for spheres and figure 5-8 for α-NPD, there are very few atoms on the limit of the collision and fitting conditions, and instead a broad distribution is formed, with the atoms mostly near their template positions and only penetrating others by a small amount. The distribution of overlaps and mismatches have similarities to the Maxwell-Boltzmann distribution and in certain conditions (low χ_c in the case of a sphere) resemble the likely population in an appropriate potential (e.g. Lennard-Jones, Harmonic). We shall see later that these rules alone are enough to create good approximations to the structure of disordered systems.

This is the essence of this geometric simulation - emulating something physical by enforcing rules based upon shape. This is how SAMSEN will approximate configurations of weakly-interacting molecules where short-range interactions dominate and to do so with little computational effort.

3.2.4 Choosing Conditions

Determining the appropriate values for χ_c and χ_f is straightforward - one just needs to make a choice about the force-field and temperature that the simulation should capture. Class 1 force-fields such as the Optimized Potential for Liquid Simulations (OPLS) or the General Amber Force-Field (GAFF) and a suitable temperature, here 300 K, allow both thresholds to be determined. Taking GAFF as an example, where the potential is given by

$$
\begin{aligned}
E_{potential} = \sum_{atoms} \frac{q_a q_b}{\epsilon r_{ab}^2} + [\frac{A_{ab}}{r_{ab}^{12}} - \frac{B_{ab}}{r_{ab}^6}] \\
+ \sum_{bonds} k(r - r_0)^2 \\
+ \sum_{angles} k_\theta(\theta - \theta_0)^2 \\
+ \sum_{dihedrals} \frac{V_n}{2}(1 + \cos n\phi - \psi)
\end{aligned}
\tag{3.4}
$$

where q_a and q_b are the charges on atoms a and b, and r_{ab} represents their separation, A_{ab} and B_{ab} represents the Lennard-Jones constants for their interaction, k represents a spring constant for their chemical bond treated as a harmonic spring with r_0 the bond equailibirum length and r the current bond length, k_θ represents the constant for the harmonc angular potential around their equilibrium bond-angle θ_0 and actual bond angle θ, and V_n the constant for the n-th harmonic of the dihedral with an angular offset of ψ. We can see that the force-field is separated into its component terms: electrostatic, pairwise, bond length, bond angle and dihedral angle. These can be separated into their external set and bonded set of forces. The external set (pair and electrostatic) will be used to inform the atomic overlaps and the bonded set (length, angle, dihedral) will inform the templates.

For the external/atomic set, the pairwise force is a Lennard-Jones 6-12 which acts over short distances between pairs of atoms. This is often called the van der Waals term. The electrostatic potential is a typical Coulombic repulsion which acts over long distances between partial of full charges on pairs of atoms. In this work, we are going to assume that partial charges are very small and we can ignore the Coulombic term. This is not compulsory and electrostatics could, in principle, be incorporated with a suitable scheme. This choice limits the range of systems that can be studied but means χ_c is the only external structural parameter and is only related to how close two atoms can approach one another at a given energy - determined by van der Waals forces alone. This closest approach for a pair of atoms, r_{ab}^{min}, is found when a pair of atoms approaching one another increase their potential by a maximum amount of energy (when the initial kinetic energy at equilibrium has been converted to potential energy). The collision threshold is then determined as the largest atomic overlap across all possible combinations of atom pairs

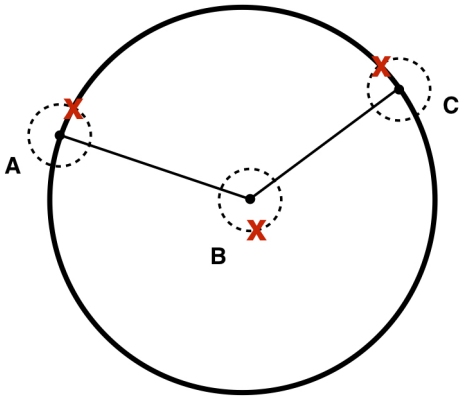

Figure 3-4: The displacements of length χ_f that produces the largest change in angle for a set of three atoms is a chord that joins two points on a circle with a radius of the bond length and a displacement of the central atom down the central axis of the set.

$$\chi_c = \max(R_a + R_b - r_{ab}^{min}) \,. \tag{3.5}$$

For the bonded set, the bond length and bond angles are both harmonic forces within the molecule between groups of two and three atoms respectively. For the bond length, the furthest possible separation is then determined when the chosen kinetic energy of the atoms equals the potential energy as the atoms move away from their equilibrium positions. This change in bond length is χ_f^{ab}. For the bond angles, the equivalent is when the three atoms move from their original positions, changing the angle until the two energies are equal. If three atoms are bound to their original positions by a common limiting distance, the maximal displacement is found when the central atom moves to maximise bond length and the other two atoms move to a position where the circle of radius equal to their limit intercepts with the circle of radius equal to their bond length. This is shown diagrammatically in figure 3-4. If their maximum angle is determined, then the limit of displacement from their equilibrium position can be determined. This maximum change in positions is χ_f^{abc}. For a given force-field, we then determine the fitting threshold as

$$\chi_f = \max(\ \max(\chi_f^{ab}),\ \max(\chi_f^{abc})\ ) \,. \tag{3.6}$$

To determine when the maximum kinetic energy equals the potential energy, we must

first determine the maximum kinetic energy. One way of doing this would be to determine the distribution of kinetic energies and then to truncate the distribution at a high percentile. For the calculations below, we do just this with the Maxwell-Boltzmann distribution at 300 K and truncating at the 95th percentile which we then take to be our maximum kinetic energy (and therefore maximum increase in potential available). This value of energy can then be used to calculate the maximum overlap (collision) length and template mismatch, below which, we aim to create a distribution of overlaps and mismatches.

For the General Amber Force-Field [121], using the boundary radii of Boyd [205] (with the exception of hydrogen, which is set to be equal to χ_c), at 300 K and removing elements such as the heavy metals, it is found that

$$\chi_c \approx 0.6\mathring{A} \tag{3.7}$$

and

$$\chi_f \approx 0.3\mathring{A} \tag{3.8}$$

across all possible combinations of atoms and are the appropriate values of the collision and fitting thresholds for the organic systems of interest. It should be noted once more that χ_c and χ_f are extremal values —we are assuming that all energy is expended in the direction parallel to the steepest increase in potential and choosing a largest likely energy to find the largest possible displacement.

If some atomic species are missing from the simulation, it is possible to reduce the thresholds to match the system being studied. For example, when only H, C, N, O are present, it may be appropriate to use $\chi_f = 0.2\mathring{A}$ to reflect the stiffer harmonic potentials (about the minimum) of the lighter elements.

It should also be restated that without additional constraints, the dihedral angle is not controlled by χ_f. This is perhaps suitable for systems where the dihedral potential has many very shallow minima and less so for systems of deep minima. If deep minima are present, it may be preferable to constrain the dihedral angles using a similar approach. One should create a new set of (larger) templates that cover the sets of atoms that determine the unrestricted dihedral angles, find a new fitting parameter and further constrain the molecules. Such an approach was not taken in this work.

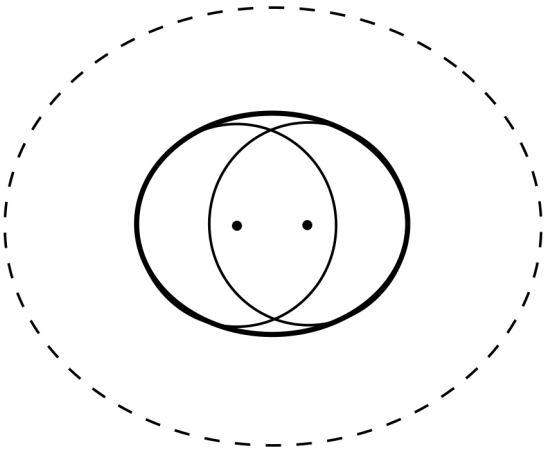

Figure 3-5: A diagram of a hydrogen molecule and its minimum volume enclosing ellipsoid (solid) and inflated ellipsoid (dashed)

3.2.5 Minimum Volume Enclosing Ellipsoids

An ellipsoid of least volume, centred on the rigid section centre $\boldsymbol{r}_s$, shall be constructed around rigid section s, which encloses the volumes of all atoms in s. This minimum volume enclosing ellipsoid (MVEE), then defines a natural volume and orientation for a rigid section. A minimum volume enclosing ellipsoid for a hydrogen molecule is shown in figure 3-5.

Instead of using radial cutoffs we shall use the ellipsoids, which better capture the anisotropy (if present) of the rigid sections, to determine if two neighbours interact. In the case of a radial cutoff, it is possible for a disc-like section to contain more neighbours in the direction of the normal axis, than around its circumference. By using MVEEs, we can reduce the likelihood of this happening.

The size of the ellipsoidal cutoff can be adjusted by multiplying the MVEE axes by an inflation factor α. As we shall see later, the precise value of α will have a significant impact on the simulations (we shall use the inflated MVEEs to select which rigid sections are interacting neighbours). The inflated MVEE of rigid section s will be referred to as E_s. We will use the inflated ellipsoids of each rigid section and test for intersections when building an elastic network (section 3.2.6) and to aid computation we will also use the MVEEs and find intersections when checking the collision conditions.

Like the rigid section templates, the ellipsoid will also hold the same shape throughout the simulation and E_s will always be rotated and translated to maintain a constant

relative orientation and position to s in the global frame.

Ellipsoid Intersections

To determine if two ellipsoids are intersecting we can frame the problem as an eigenproblem. The method used here is that of Alfano & Greer [206].

An ellipsoid E can be described in matrix form by

$$XEX^T = 0 \tag{3.9}$$

where

$$E = \frac{1}{2}\begin{bmatrix} 2A & D & E & G \\ D & 2B & F & H \\ E & F & 2C & J \\ G & H & J & 2K \end{bmatrix} \tag{3.10}$$

and

$$X = \begin{bmatrix} x & y & z & 1 \end{bmatrix} \tag{3.11}$$

This describes the surface of the ellipsoid in any position and orientation in the Cartesian frame. One can then formulate an eigenproblem in the usual manner for a system of two ellipsoid surfaces

$$X(\lambda E_1 - E_2)X^T = 0 \tag{3.12}$$

and as Alfano & Greer demonstrate, a simple inspection of the eigenvalues or substitution of the eigenvectors into the equations of the ellipsoids allow one to determine if two ellipsoids intersect. For instance, if the real-component of the eigenvectors produce points inside or on the surface of both ellipsoids or if there are real eigenvalues, then the ellipsoids must intersect.

If in the simulation, two rigid section ellipsoids intersect, then the rigid sections shall be described as being in-contact.

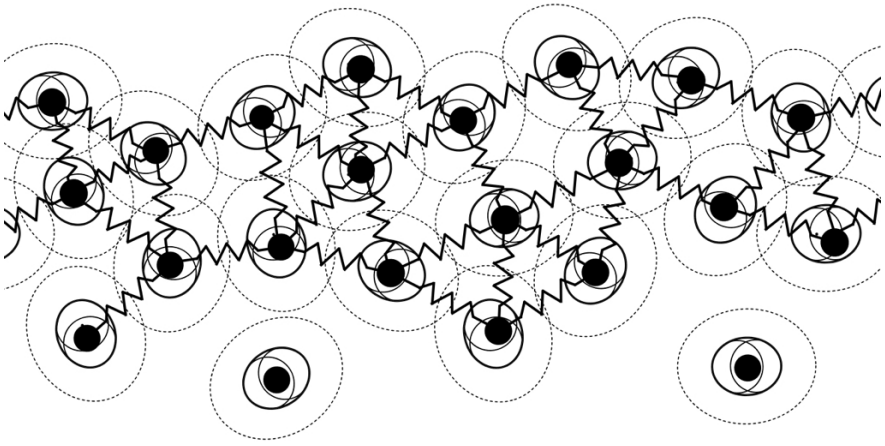

Figure 3-6: The construction of the dynamical system is performed between rigid sections that are in-contact (their ellipsoids intercept) including those over periodic boundaries.

3.2.6 The Elastic Network

To create the dynamics of the system, SAMSEN employs an elastic network model. The elastic network model describes a set of local contacts connected by harmonic springs and forming a point-spring network. The energy landscape of this network can be described by the Hessian matrix – the second-derivatives of the potential energy around each point. The eigenvalues of this matrix describe the vibrational frequencies and the eigenvectors describe the vibrational displacements at the corresponding frequencies. These are the normal modes of vibration. In SAMSEN, we shall use these normal modes to try and overcome the dynamical slowdown experienced by molecular dynamics, by neglecting the higher frequency displacement and instead moving the system along the low-frequency modes.

To construct the elastic network, nodes (points) are placed at the positions, $\boldsymbol{r}_s$, of the rigid sections. The ellipsoid intersections are then determined throughout the system. If any two rigid sections are in-contact, then they are given an interaction in the elastic network: a harmonic spring is created between rigid sections i and j with an equilibrium length equal to their current separation (the spring is relaxed). This is governed by the

potential

$$V(r_{ij}) = \frac{1}{2}k(r_{ij} - r_{ij}^0)^2 \tag{3.13}$$

where k is the spring constant for the pair which shall be assumed to be equal for all interacting pairs $k = 1$ kg s^{-2}. For all pairs in-contact, $r_{ij} = r_{ij}^0$ therefore $V_{ij} = 0$. Construction of such an elastic network is portrayed in figure 3-6. If a rigid section ellipsoid does not intercept with another it is not given an interaction in the Hessian.

To calculate the Hessian, the partial second-derivative with respect to each Cartesian component is determined. This creates a 3×3 array or super-element which fills off-diagonal rows and columns $(3i, 3j)$ to $(3i+2, 3j+2)$ and vice versa. This super-element is of the form

$$H_{ij} = \begin{bmatrix} \frac{\partial^2 V_{ij}}{\partial x_i \partial x_j} & \frac{\partial^2 V_{ij}}{\partial x_i \partial y_j} & \frac{\partial^2 V_{ij}}{\partial x_i \partial z_j} \\ \frac{\partial^2 V_{ij}}{\partial y_i \partial x_j} & \frac{\partial^2 V_{ij}}{\partial y_i \partial y_j} & \frac{\partial^2 V_{ij}}{\partial y_i \partial z_j} \\ \frac{\partial^2 V_{ij}}{\partial z_i \partial x_j} & \frac{\partial^2 V_{ij}}{\partial z_i \partial y_j} & \frac{\partial^2 V_{ij}}{\partial z_i \partial z_j} \end{bmatrix} \tag{3.14}$$

where

$$\frac{\partial^2 V}{\partial x_i \partial y_j} = -k\frac{(x_j - x_i)(y_j - y_i)}{r_{ij}^2}\,. \tag{3.15}$$

The diagonal super-elements for interactions (i, i) and (j, j) are filled with similar components of the form

$$\frac{\partial^2 V}{\partial x_i \partial y_j} = k\sum_j \frac{(x_j - x_i)(y_j - y_i)}{r_{ij}^2}\,. \tag{3.16}$$

3.2.7 Vibrational Modes

Diagonalisation of the Hessian produces a set of $3N$ eigenvalues and $3N \times 3N$ eigenvectors which are the vibrational frequencies and displacements of the normal modes

respectively. As the Hessian is constructed with each node at a minima ($V_{ij} = 0$, $k = 1$ or 0), all eigenmodes are real with positive or zero frequency.

In a system where no springs are formed across the periodic boundaries, there are 6 trivial modes at zero-frequency. These modes describe global/uniform translations and global/uniform rotations - they describe no change of the relative position of the particles, but instead the reference frame. In a periodic volume where springs are formed across the boundary, the rotational modes do not exist and instead there are only 3 trivial modes that describe uniform translation. These modes are to be neglected and we shall instead refer to the lowest frequency non-trivial modes as the lowest frequency modes for the remainder of this work.

Amplitude of Vibration

For a given mode, one can scale its amplitude according to a suitable scheme by stretching the eigenvectors. In SAMSEN, we assume that a wave's amplitude can be calculated through equipartition

$$\frac{3}{2}k_bT \propto \frac{1}{2}\langle A^2\rangle\omega^2 \tag{3.17}$$

where A is the mean amplitude of the wave, ω is the wave's frequency and T is the temperature of the system in which the wave exists. This may not be achieved in a real system: the wave may not reach its maximum amplitude (due to anharmonic effects, change in forces, etc.), not all oscillations are activated at a given temperature, and the vibrational energy may be distributed across the activated modes, for instance. However, we do not intend to recreate the true vibrational dynamics here. We, instead, are interested in neglecting the short timescale dynamics and enhancing the low-frequency dynamics. We shall therefore replace T with ϵ and refer to it as the effective temperature of the dynamic model (to prevent misunderstanding) where each mode is given an r.m.s. amplitude β_m, calculated as

$$\beta_m^2 = \frac{3k_b\epsilon}{\omega^2\langle\delta_m^2\rangle} \tag{3.18}$$

where $\langle\delta_m^2\rangle$ is the average displacement across all rigid sections in mode m

$$\langle \delta_m^2 \rangle = \frac{1}{N} \sum_i^N |\boldsymbol{\delta}_{im}|^2 \,. \tag{3.19}$$

This means that all eigenvectors are re-scaled such that the r.m.s. displacement is equal to β_m, and ϵ provides a way of controlling the kinetic energy of the wave at mechanical equilibrium (by construction) and therefore the displacement length.

Assuming Similar Frequencies

If a subset of the modes are chosen in the range $m_{min} \leq m \leq m_{max}$ (which shall be compactly written $m_{m_{min}}^{m_{max}}$ throughout) which consists of a number of modes over a narrow frequency range, then we can make the assumption that the vibrational frequencies are similar

$$\omega_{m_{max}} - \omega_{m_{min}} \approx 0 \,. \tag{3.20}$$

Under this assumption, the period of the mode oscillations are also similar and all waves reach their peak displacement at the same time. This assumption allows a linear sum of the scaled eigenvectors to represent the displacement, rather than a superposition of waves which vary as a function of time.

Therefore, the total displacement for each rigid section in a SAMSEN cycle will be a linear sum over the displacements in each mode in the selected range

$$\boldsymbol{D}_i = \frac{1}{\sqrt{M_i}} \sum_{m_{min}}^{m_{max}} \beta_m \boldsymbol{\delta_{im}} \tag{3.21}$$

where M_i is the mass of rigid section i, which in this work shall be assumed to be $M = 1$ a.m.u. for all rigid sections - if accurately considered and included in the construction of the Hessian along with pair-specific equilibrium lengths (yielding the dynamical matrix), then the displacements calculated from equation 3.21 would become the Newtonian dynamics (as approximated).

Each atom will then move by the average rigid section total displacement vector that the atom is a member of. This will be performed in a series of steps so that collision-fitting conditions can be maintained as the particles move. As the masses and spring constant are equal and the wave frequencies (and therefore the acceleration) are similar,

instead of displacing by shorter distances as the vibration approaches its peak, the time interval at each step can be re-scaled (physically getting longer and longer) so that at each step an equal distance interval is covered. The displacement at each step is then given by $\boldsymbol{D}_i/S$ where S is the number of steps chosen for the oscillations to reach peak r.m.s. amplitude.

Assuming similar frequencies also means that each displacement period has a time associated with it. This is simply equal to a quarter of the full oscillation period for any of the modes in the range $m_{m_{min}}^{m_{max}}$.

3.2.8 Repeat

At each step in the displacement, the collision and fitting conditions must be checked and maintained before proceeding onto the next step. In this way, the structural and dynamical models are applied in-turn. One could imagine the structural model managing the high frequency restoring forces, and the dynamical model controlling the low-frequency motion and that the disparate timescales allow them to be applied without consideration of the other.

Once the system has reached its full displacement (after performing S steps), a new set of ellipsoid intercepts are found and a new elastic network is constructed. The old set of displacements are then replaced by a new set over the same range of modes. Displacement then continues until another S steps have passed. This is a SAMSEN cycle and becomes the *natural* unit of simulation time (although, if required, it can be related to a physical time through the frequency of the range of modes chosen).

3.3 Simulation

Now the rules that govern a SAMSEN simulation have been established, we can turn to more practical considerations: the creation of an initial state, compression to the required density, polymerisation; to name a few. This section should therefore be considered to be the description of the experimental setup for the series of investigations in this work and specific additional details of the computations performed.

3.3.1 Creating Initial States

Preparation of Molecules

SAMSEN relies heavily on the creation of templates of rigid sections within a molecule which are used the maintain the molecular structure, so it is important that they reflect the likely configuration of the molecules (ideally the ground state configuration). Therefore, before the simulation begins, one must construct or download a suitable molecular structure and optimise its geometry.

This optimisation can be performed at the density functional theory level or, if the molecule is too large, could also be performed using an energy minimisation (e.g. steepest descent, conjugant-gradient) in an appropriate force-field. If, later, one intends to perform charge transport simulations by calculating the electronic overlaps, this optimisation using DFT methods will need to be performed anyway. One would also expect this to present a more accurate representation.

For small molecules, the structures produced from a minimisation in an MD force-field were used. However, for polymers, the optimisation was performed at the DFT level on oligomers and the central repeat unit taken as the optimised structure.

Rigid Sections

Taking this geometry-optimised molecule (or repeat unit - I will, from time to time, use 'molecule' to refer to either), it can then be passed to the algorithm which determines the rigid sections of the molecules, described in section 3.1.3.

For the entirety of this work, this was the method used to produce the rigid sections, although I again note that this could be inferred from calculation of the internal vibrational mode frequencies or other energetic consideration of internal motion and angular potentials.

The membership list for each rigid section is stored for later reference and the minimum volume enclosing ellipsoids are generated.

Polymers

To create polymers in SAMSEN, a series of molecules placed in a row is not sufficient. Like with any simulation, it is important that the repeat units are connected. In

MD/MC that is performed by including additional bonding forces between atoms, here it is achieved by extending the rigid sections between repeat units.

In this work, the rigid section templates of the repeat unit were created with additional atoms present (not removed after the DFT optimisation of the oligomer), left as placeholder atoms. When the simulation is launched, the placeholder atoms will be deleted and the corresponding atoms on the neighbouring repeat units will assume all rigid section memberships of the placeholder atoms. However, the assumed template positions for the atoms in neighbouring repeat units will always be of the initial placeholder atom positions. This subtle point means polymers can be placed/loaded in a twisted configuration without affecting the shape of the templates.

Ellipsoids

To generate the minimum volume enclosing ellipsoids, the coordinates of each atom in the rigid section, as well as the atom types (the element) is given to an optimisation algorithm. This algorithm places each atom in space, centres the set around the origin and places additional markers to represent the size of the atoms (at the boundary radii of Boyd [205]). Using a guess ellipsoid that is much larger than and completely encloses the set of atoms and markers, the program then performs a sequential least squares optimisation on the ellipsoid axes, where the optimal ellipsoid is of smallest volume with the constraint that all atoms and markers must remain within the ellipsoid volume and the volume must be positive. It is similar to a gradient descent, where the gradient is the change in ellipsoid axis length. Other techniques including genetic algorithms also produce similar solutions. The ellipsoid axes are then stored for later use. The implementation of the volume minimisation under constraints used here was written by Ian Thompson of the University of Bath using the optimisation in the SciPy library [207].

Generating an Input

We then place the chosen number of molecules into the simulation volume with random positions and orientations. Each time a molecule is entered, we check to ensure that the ellipsoids between any two rigid sections, that are not a member of the same molecule, do not intercept.

For polymers also, the placement of the first repeat unit is random, the rigid sections

checked for intercepts, but the placement of the neighbouring repeat units is restricted. The rule is that placeholder atoms must be positioned so that the atoms that will assume their membership have the same coordinates (a slight rotation of the new repeat unit is permitted if the rotation angles are small enough to produce a mismatch between the two that is within χ_f). The new repeat units must also contain no ellipsoids which intercept with ellipsoids from other polymers/molecules and also not intercept with ellipsoids of rigid sections in the neighbouring repeat units.

If any of these ellipsoid intercept tests are failed, then that entire molecule or polymer is deleted and a new attempt is made with random positions and orientations.

This method of ensuring no ellipsoids overlap is able to produce initial morphologies at quite high density. In the case of spheres, more than a third of the box volume can be filled using this random placement, but in the case of more awkwardly shaped small molecules and polymers (for example, the ones mentioned in this work), a starting density in the range of 0.05 to 0.15 g/cm^3 can be readily achieved. Of course, high density starting configurations create correlations in position/orientation - particularly in the case of polymers - but we expect to experience this during the compression stage and not be constrained by the initial conditions after the system repeatedly relaxes as it follows the low-frequency modes.

Loading Inputs and Polymerisation

SAMSEN begins by loading in the optimised molecular structures in the given positions and orientations. It then immediately generates the set of template positions for the rigid sections of the molecules. The atoms in each rigid section are then moved to match the positions and orientations specified in the input specification (this allows previous simulation states to be loaded and molecular structures to deviate from the optimised position if desired/required) - taking the average position of their template if an atom is the member of multiple rigid sections.

The polymerisation process then takes place, with the placeholder atoms deleted and the corresponding atoms on the neighbouring repeat units assuming all rigid section memberships. The template positions remain unchanged so that the atoms that have just assumed membership must conform to the fitting condition measured against the template of the new rigid section(s) in the neighbouring repeat unit as well as ones it is already a member of. By this mechanism the entire polymer structure becomes interlocked and can be treated as a molecule in the normal way: considering only the

atoms and rigid sections.

It is perhaps worth mentioning briefly that if one is interested in constraining the dihedral angles in polymers, this could also be performed here by creating a set of dihedral templates. More placeholder atoms would be required, and the second neighbour atoms in neighbouring repeat units could then assume their position. A similar fitting condition, perhaps with a different threshold length, may be required to capture the strength of the dihedral potential.

3.3.2 Compression

Converting a gas/liquid-like density input state into a state at a solid thin-film density is performed by a basic compression algorithm which represents the inverse process of the Lubachevsky-Stillinger compression algorithm [208] at zero temperature - instead of inflating particles, the space between particles is shrunk. A compression is performed in steps and, at each step, a fractional amount is removed from the coordinates of all atoms in the system as well as the box dimensions. The system is then allowed to locally relax - here by enforcing the collision-fitting conditions - before proceeding to the next compression step.

In this work, all coordinates and box dimensions are multiplied by a small factor C_f which, unless stated otherwise, is 0.99995. A $C_f < 1$ represents a compression while $C_f > 1$ is an expansion. During this process, we shall also here set $\epsilon = 0$ so that this compression/expansion is athermal.

This compression (or expansion) is an essential, but non-integral, part of the process we shall follow in this work. SAMSEN strictly only governs the validity of a state, the enforcement of the collision-fitting conditions, and the long-time dynamics of following the low-frequency modes. Starting from a crystal phase or from a molecular dynamics output state would be equally valid approaches to SAMSEN. The only requirement is that the input state satisfies the collision-fitting conditions of the templates generated from the optimised molecular structure (if it is not, the collision-fitting correction cannot guarantee that the input can be turned into a valid state).

3.3.3 The SAMSEN Cycle

Once a configuration has been found at the required density, we shall set ϵ to a non-zero value, set $C_f = 1$ and start to use the dynamical model to sample configurations as

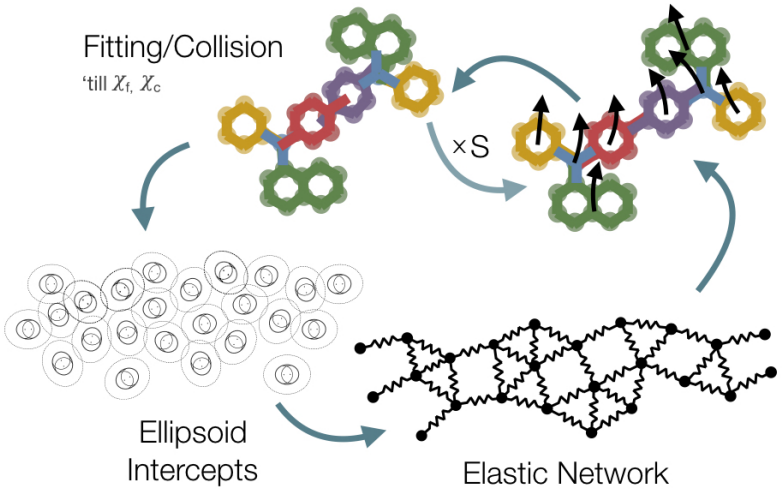

Figure 3-7: A diagram of the SAMSEN cycle. It begins by ensuring the state is valid, checking collision and fitting conditions and performing corrections if necessary. Then an elastic network is constructed and the low-frequency modes calculated. The rigid sections are then displaced along these modes over a period of steps, S, and the conditions for a valid state are maintained at each step. After S steps a new set of intercepts are determined and new modes calculated.

described in sections 3.2.6 and 3.2.7 and use the structural model of section 3.2.3 to maintain a valid state which is designed to reproduce (or at least approximate) the morphologies of other methods.

The stages of a SAMSEN cycle are described in this section and are depicted in figure 3-7. A cycle begins with the determination of neighbouring rigid sections and the calculation of vibrational modes, which are used to displace the rigid sections over a series of steps. At each step the collision and fitting conditions must be satisfied before proceeding onto the next step. The cycle then ends after a predetermined number of steps and a new one begins by calculating a new set of vibrational modes.

Calculate vibrational modes

The elastic network is constructed by placing harmonic springs (with $r_{ij} = r_0$) between any ellipsoids in the system which intercept one another (including across the periodic boundary). This creates a $3N \times 3N$ Hessian matrix.

Consider the memory requirements for a moment: for such a matrix of double precision floating point numbers, the Hessian at $N = 5,000$ fills 1.8 GB, while at $N = 10,000$ a 7.2 GB Hessian is constructed. To reduce the memory-burden, which quickly becomes larger than some desktop PCs are capable of storing, sparse-matrix representations are used.

However, memory is not the only concern. The diagonalisation of such a large matrix is known to be very time-consuming. For instance, in some brief tests, a Hessian for $N = 1,000$ (with an average of 12 contacts each) required tens of minutes to produce the full spectrum of vibrational modes for both eigenvalues and eigenvectors using the GSL library [209] which was limited to a single CPU without parallelism. This, of course, is not a particularly large system and this timescale for calculating the displacements for one cycle is inconvenient, to say the least. This approach also involves calculating information about modes that we do not intend to use - as we are only interested in a low-frequency subset.

To speed up this process, it was found that using an Arnoldi shift-and-invert method from the Arnoldi Package (ARPACK), specifically ARluSymStdEig in ARPACK++ [210], substantially lowered the time required to calculate the required number of modes. For the same $N = 1,000$ Hessian, in sparse-representation, and solved around a central 'guess' frequency that was lower than the lowest frequency mode, the diagonalisation solving for 200 modes could be performed in seconds on a single CPU and seemingly-instantly on 8 CPU cores. This process is faster for fewer modes, however it is important that the central frequency is set so that the lowest frequency eigenmode is found after the first diagonalisation, otherwise further attempts around new guess frequencies will be required. This can be achieved by setting a low central frequency such that the lowest modes are the lowest modes found or by having a wide search range. To ensure that the lowest frequency modes have been found, a check for the presence of the translation modes can be performed (by testing the dot product of the eigenvectors - if they are translational they should all be ≈ 1). This shift-and-invert method is approximate, but it appears using a larger search range also increases the accuracy of the eigenvectors.

For this work, the search range used in SAMSEN was set to at least twice the number of modes in the mode displacement range, $m_{m_{min}}^{m_{max}}$ and, with a low central guess value, usually 0.25, that falls between the lowest frequency eigenmodes and the translational modes. If the translational modes were not in the search, successive diagonalisations were performed with new guesses values until the translational modes were identified. For the eigenvector analysis, the solve was performed with a guess value of 1×10^{-5}

using the SciPy interface to ARPACK [207].

Displace Rigid Sections

With the eigenmodes found, containing modes in the range $m_{m_{min}}^{m_{max}}$, the displacements for each of the rigid sections were calculated using equation 3.21. These are summed over the chosen mode range, creating one displacement vector per rigid section. This is divided by the number of steps per cycle, S, which for the majority of this work was set to $S = 1,000$. This is the set of displacement vectors for the cycle.

Each atom is then given the displacement vector of the rigid section of which it is a member. If it is a member of multiple rigid sections, it gets the average displacement vector. At every step in the cycle the atoms then move along their displacement vectors.

Maintain Collision and Fitting Conditions

At each step, the collision and fitting conditions are checked and enforced, as described in section 3.2.3 and is implemented as shown in the flowchart in figure 3-8.

The collisions checks are performed by identifying neighbouring rigid sections using a cell list. This cell list contains the identifies of rigid sections within each cell and each cell is cubic with sides having a minimum length equal to the maximum MVEE axis length. From this list of local neighbours, ellipsoid intercepts using the MVEEs (not inflated by factor α) are then used to determine if any of the atoms in the rigid sections could intercept with those in a neighbouring rigid section (applicable when $\chi_f < 2\chi_c$). If there are any MVEE intercepts, then all pairs of atoms in the two associated rigid sections are then tested for overlaps. After storing all overlaps in the system, if collision conditions are broken then the correction vectors (moving the spheres so they are just touching, or taking an average if the atom is involved in multiple collisions) are applied.

The fitting checks are performed by finding the optimal rotation of the rigid section template against the current positions of its member atoms. All mismatches are calculated and stored. After performing this for all atoms, if the fitting condition is broken then all atoms are moved to their template positions (or an average correction is applied if an atom is a member of multiple rigid sections).

If either of the collision or fitting conditions are broken then a re-check of both collision and fitting conditions will need to be performed. This loop continues until both

conditions are met in a single loop.

This strategy, of filtering atoms by rigid sections, greatly improves the computational speed without altering the simulation. This makes the collision detection scale quadratically with the number of rigid sections (and number of atoms per section) rather than quadratically with the number of atoms in the system. The fitting and mismatch calculation also scales with the number of rigid sections (and the number of atoms per section). And, as the diagonalisation scales only with the number of rigid sections, then the number of rigid sections becomes the natural way of describing the computational complexity of a SAMSEN simulation.

No further mode displacement is applied until these conditions are met and the state declared valid. When this is achieved the next step begins by reapplying the displacement vectors.

Repeat

This displacement and collision-fitting correction sequence continues until the end of the cycle, defined as after S steps have elapsed, which is when the r.m.s. displacement vector applied for each mode (potentially distinct from the total r.m.s. displacement achieved) is equal to β_m. At this point the next cycle begins with the construction of a new elastic network.

The SAMSEN simulation ends after a specified number of cycles have been completed.

3.3.4 Summary of SAMSEN Parameters

This section is dedicated to highlighting the different parameters introduced in this chapter, in particular those that will have an effect on the simulation result, and is intended as a reference. The effect of changing most of these parameters will be investigated further in chapters 4 and 5.

$\boldsymbol{\alpha}$

The ellipsoid inflation factor, α, was discussed in section 3.2.5. 3.2.5. It is used to inflate the minimum volume enclosing ellispoids that were generated for each rigid section when we come to calculate the vibrational modes. This factor will determine how many ellipsoid intercepts are found and will be examined further in section 4.2.

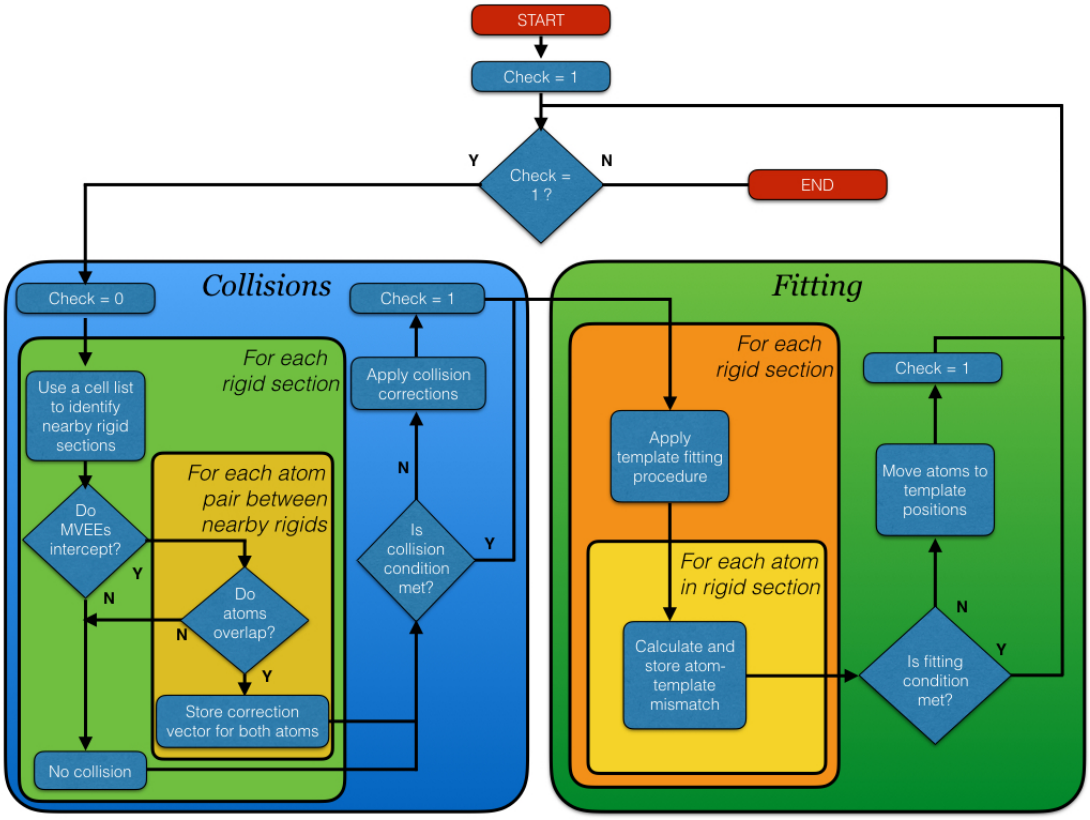

Figure 3-8: Flowchart describing the collision-fitting process implemented in this work, simplified for visual clarity. At the beginning of the sequence, checks for collision and fitting are turned on and we begin the loop of assessing if condition-fitting conditions are met or if we need to enforce and re-check conditions. The loop starts by turning the recheck indicator off and then, for each rigid section in the system, identifying neighbouring rigid sections through use of a cell list and, for each neighbouring rigid section, checking if their MVEEs intercept. If there are intercepts between the MVEEs of two rigid section then, for each pair of atoms in those rigid sections, a distance calculation and collision condition check is performed. Any overlaps are stored and corrections only applied after all overlaps in the system have been calculated and if the SAMSEN collision condition is broken. The loop then continues, looking at the fitting conditions by achieving a best fit of the rigid section templates to the member atoms as they are currently positioned in the system and then by calculating and storing any positional mismatch. If the SAMSEN fitting conditions are broken, atoms are moved to their template positions and another collision-fitting check will need to be performed. If not and if the collision condition was also not broken, the loop exits with the system in a valid SAMSEN state.

ω

The vibrational mode frequency, ω, was first discussed in section 3.2.7. It is a result of diagonalising the Hessian which, in turn, depends on the separation of connected nodes (rigid sections) in the elastic network and constants such as the spring constant, k (which we assume to be unity). ω is primarily used in equation 3.18 to calculate β_m and takes on a range of values which can be picked by m_{min} and m_{max}, however we will assume $\omega_{m_{max}} \approx \omega_{m_{min}}$ (equation 3.20) for reasons which will become clear in section 4.3.

m_{max} and m_{min}

The upper and lower mode numbers respectively, m_{max} and m_{min}, determine the range of modes which displacement vectors will be summed over in equation 3.21. This parameter has quite wide-ranging effects as it also determines ω and as a result β_m. We will also see how the nature (specifically quantities such as the collectivity) of the modes vary in section 4.3 and m_{max} and m_{min} provides scope for the user to vary how collective rigid section displacements are in SAMSEN. Changing these parameters will be discussed (as with ω) in section 4.3 and for the majority of this work $m_{max} = 12$ and $m_{low} = 1$.

ϵ

The effective temperature of the dynamical model, ϵ, is described in section 3.2.7., replacing T in equation 3.17 and used to calculate β_m in equation 3.18. It has units of temperature. This parameter is the primary control of the length over which rigid sections will be displaced in a given SAMSEN cycle as it contributes to β_m. We will investigate the effect of changing this parameter on the structure and sampling rate in section 4.4.

S

The number of steps per cycle, S, was introduced in section 3.2.7 and describes how many displacement steps the calculated rigid section displacements are performed over. In many ways this is a coarse-graining of time within a SAMSEN cycle, with a small value of S leading to large, sometimes unstable displacements and a large value of S

leading to smooth displacements. Computationally, there may exist a most efficient value somewhere between the two extremes for a given accuracy. For this work $S = 1,000$ is commonly used, particularly in the case of molecular systems, however it may be beneficial to increase it for extremly large β_m or reduce it for low β_m (or ϵ).

χ_c

The collision threshold, χ_c, was explained in section 3.2.3 and its value is determined in section 3.2.4. It is used in the collision condition (equation 3.2) and sets the upper limit on which two atoms not in the same rigid section can penetrate one another (or how much their volumes can overlap). It can be described as the maximum overlap length for a valid SAMSEN state and has units of length. The effect of changing χ_c will be examined in section 4.5.

χ_f

The fitting threshold, χ_f, was described in section 3.2.3 and its value is determined in section 3.2.4. χ_f describes the maximum distance an atom can be from any of its rigid section template positions in a valid SAMSEN state. It may also be described as determining how much flexibility is granted both within rigid sections and between rigid sections with $\chi_f = 0$ $\mathring{A}$ representing true rigidity within a rigid section and allowing only co-axial rotation between sections. It sets the upper limit on the fitting condition and has units of length. It will be discussed in section 5.2.1 and 5.2.2 in the context simulations of a small molecule where the method of determining rigid sections is also evaluated.

ϕ

The packing fraction, ϕ, describes the volume of occupied space to the total simulation volume. In this work ϕ is controlled by setting the volume which the original simulation box is compressed to and is measured exactly in the case of hard spheres, and also measured for inflated rigid section ellipsoids (where overlapping is allowed). We will briefly examine how the value of ϕ impacts the vibrational modes in section 4.1. However, for molecular systems, ρ, the density will become the quantity of interest as this makes comparison with other published work much easier.

C_f

The compression factor, C_f, determines the rate at which the simulation is reduced in volume until the desired volume (or density) is reached. It is described in section 3.3.2 and is the multiplicative factor used to modify the lengths of the sides of the simulation volume. Once the required density/packing fraction is reached, C_f is set to 1 and compression ceased for the rest of the simulation. In this work, no displacement of rigid sections along the vibrational modes is performed during the compression phase, however collision and fitting conditions must be maintained.

3.4 Charge Transport Model

To investigate charge transport within the organic semiconductor morphologies, the following model will be used. It solves the master equation, under low-charge density conditions and requires knowledge of the rate of transport between all pairs of sites in the system. These charge transport rates will be later calculated using the Förster equation (section 2.2.1) for excitons in P3HT, using the Miller-Abrahams rate for electrons in PCBM and using Marcus theory rates in the case of hole transport in IDT-BT. With this information it is possible to calculate the steady-state distribution of charges (which amounts to the probability of occupation in a single particle system) and, using the same transfer rates, calculate properties such as the diffusion constant, D and the exciton diffusion length, L_D. Other methods such as kinetic Monte Carlo can also be used to determine such values but will not be used in this work.

3.4.1 Master Equation

The master equation describes the movement of charges within a system through time, in a treatment that works with the charge densities (or probability of occupation) and rates (or probabilities) of transfer between discrete sites [11, 38, 211]. In the case of charge transport, it is based upon the Pauli master equation, which also incorporates the reduced probability of moving to a site that is already occupied, and is given by

$$\sum_{i \neq j} \nu_{ij} p_i (1 - p_j) - \nu_{ji} p_j (1 - p_i) = 0 \tag{3.22}$$

where ν_{ij} is the charge transfer rate between two sites, i and j, and p_i is the probability

of a given site being occupied. Note here that the amount of transfer from site i to site j is reduced depending on the probability that site j is already occupied through the $1 - p_{ij}$ term. At all times, the amount of charge in the system is conserved (although sources and drains can be added), and at steady-state the density moving forwards (i to j) is equal to the density moving backwards (j to i).

In the low-charge limit, when charges are considered isolated, or when particles are entirely non-interacting, $1 - p_{ij} \approx 1$ and the Pauli exclusion term can be neglected and equation 3.22 becomes

$$\sum_{i \neq j} \nu_{ij} p_i - \nu_{ji} p_j = 0 \,. \tag{3.23}$$

We shall use this form of the master equation (without the Pauli exclusion term) in this work. However, it should be reiterated that this restricts the applicability of results to the low-charge limit and will alter the calculated charge transport properties (potentially increasing the calculated diffusion constant in single species systems for instance). To attain self-consistency, the electric potential induced by neighbouring charges will also not be included in the calculation of hopping rates between sites.

3.4.2 Transfer Matrix

The transfer of charges between sites under the master equation can be described by a transition rate matrix [212]. As we shall see later, this matrix contains all the information required to calculate a range of charge transport properties. The transition rate matrix is an $N \times N$ matrix where N is the number of sites and contains elements that describe the rate at which particles will move to all other sites. In population analysis, these are sometimes called transient states and absorbing states respectively [212]. The off-diagonal elements are filled with the rate, ν_{ij} which corresponds to a transfer from state (site) i to state (site) j [212]. This matrix is not symmetric and a transfer ν_{ji} in the matrix element for a transfer from j to i can be different to a transfer ν_{ij} (with the energy differences between sites, it is very likely they are different under a Gaussian disorder charge hopping model, for instance). The diagonal elements are then filled with the negative sum of the off-diagonal elements (in a single forward or backward direction) such that (i, i) contains the negative sum of all values in row i (the rows sum to zero).

The transition rate matrix can also be converted to a transfer matrix (sometimes simply

the transition matrix but for clarity, we shall call it the transfer matrix), which describes the amount of transfer occurring in a given time. In this context, it describes the probability that a charge will move from one site to another in a time, t, which can be calculated from knowledge of the transfer rate [213]. In this way, the transition rate matrix and transfer matrix are related. If the off-diagonal elements in the transition rate matrix can be described by

$$\nu(ij) = -\frac{1}{t}\ln\left(1 - T(ij)\right) \tag{3.24}$$

where $T(ij)$ is the probability of transfer from site i to j (previously p_{ij}) [212], then the off-diagonal elements in the transfer matrix can be written as

$$T(ij) = 1 - \exp(-\nu_{ij}t) \tag{3.25}$$

with ν_{ij} the rate of transfer from i to j as written in the off-diagonal of the transition rate matrix [213]. The diagonal elements of the transfer matrix are then filled such that the sum across a given row of the matrix is equal to 1, which is the condition for the steady-state [212]. It is also possible to modify this matrix to include charges leaving and entering the system (light absorption, exciton recombination/decay) and to use it to consider density-dependent systems [212], however this will not be considered in this work as this can create non-linearity which requires the more direct iterative approach (or KMC) [43, 211, 212]. Readers should refer to books such as Matrix Population Models by Caswell (2001) for further information on the properties and application of the transfer matrix approach [212].

Both of these matrices can be used to find the steady-state distribution of charges. The transfer matrix, which is perhaps conceptually simpler to follow, describes the probability of charges moving. Given an initial distribution of charges or probability of occupation of a set of sites, application of the transfer matrix through matrix multiplication will provide the distribution of sites at time $t_0 + t$ where t_0 is an initial time at which an initial distribution of charges exists and t is the time interval between states. If there is a population of charges (or a probability of occupancy), described by a vector P of N sites, then the population at time $t_0 + t$ can be calculated through the formula

$$P|_{t_0+t} = TP|_{t_0} \tag{3.26}$$

which is accurate for small t and, when applied iteratively, converges to the steady-state distribution, as does the average occupation in a kinetic Monte Carlo simulation, for instance. As the charges are non-interacting in the low-charge limit, the probabilities in the transfer matrix or the rates in the transition rate matrix do not need to be updated.

Unlike the stochastic methods, calculating the steady-state charge distribution and charge transport properties is much faster with the transfer matrix approach. This is because the simulation does not need to be run to convergence as in kinetic Monte Carlo which can take many millions of steps [38]. Instead of iteratively multiplying the transfer matrix, the steady-state can also be calculated by diagonalisation of the transfer matrix (or the transition rate matrix) [212] to find the dominant-eigenvalue which, using a partial diagonalisation such as that provided by the shift-and-invert method in packages such as ARPACK [210], can be calculated very quickly. And as the steady-state eigenvalue is known *a priori*, the calculation can be performed for very few eigenvalues (and eigenvectors) and therefore even faster. For instance, in the transport calculations on systems of 1,000 PCBM molecules in section 5.3.2, the diagonalisation and calculation of the diffusion constant were be performed in less than a second on an 8 core machine.

When diagonalising the transfer matrix, the steady-state is represented by an eigenvalue of 1 (which indicates no further change to the distribution of charges upon multiplication) [212] and for the transfer rate matrix an eigenvalue of 0 corresponds to the steady-state distribution. The corresponding eigenvector, then contains the distribution of charges which is equal to the probability of site occupation times the number of charges in the system.

3.4.3 Diffusion and the Exciton Diffusion Length

With knowledge of the probabilities of a site being occupied and the probabilities of transferring between sites in a time period of length t (held in the transfer matrix), the average displacement of populations/particles over time can be calculated. If vector P is transformed into an $N \times N$ matrix, with the probability of occupation stored in the diagonal and the remaining elements set to zero, the matrix multiplication now described by equation 3.26 can be performed to determine the probability of particles moving from a given starting site to a neighbouring site, weighted by the probability of the particle starting on that initial site. This yields an $N \times N$ matrix with the probability/density of charge that started on the diagonal in any given row reduced

and redistributed to the remaining columns in that row. Repeating the multiplication and summing over columns then shows how the charges move throughout the system, weighted by the initial distribution of the site occupancies. This property can be used to determine the likely 'traffic' of given routes or establish which routes in the system dominate or how easy it is for a charge to travel from A to B. After enough iterations all rows will then converge again to steady-state, albeit with different total populations.

The mean-squared displacement of a single exciton in the system can then be described by the products of distances between sites and occupancy-weighted probability of transfer (a similar approach for calculating velocity and charge mobility was taken by Yu et al. and Koster [42, 214]). Put more succinctly, the mean-squared displacement $\langle S^2 \rangle$ over time t is simply given by

$$\langle S^2 \rangle = \sum_{i}^{N} P_i \sum_{j}^{N} r_{ij}^2 T_{ij} \tag{3.27}$$

Due to the density-based approach and its convergence to steady-state, this calculation is most accurate with a fewer iterations of the transfer matrix and at small time intervals. Therefore, in this approach we will consider a single iteration, which equates to the product of the squared distance between sites i and j and their associated probabilities of transfer, weighted by the probability of a charge being on that site, in a known time period, t, as in equation 3.27.

This then leads us the ensemble average diffusion constant, D, for a single particle the system (through Fick's law of diffusion, which takes the same form as the heat equation)

$$\frac{\partial n}{\partial t} - D\nabla^2 n = 0 \tag{3.28}$$

where the 3D diffusion constant, D, can be determined as the occupancy-weighted transfer probability-weighted mean-squared displacement per unit time $6D = \langle s^2 \rangle / t$ as $t \to 0$ and shall be calculated as such using small t for electrons and holes.

For excitons, the quantity of interest is not the diffusion constant but the diffusion length, which defines the distance excitons are likely to travel before decaying and the electron returns to the HOMO level (in an exciton lifetime, τ) [79, 215, 80]. The diffusion length, L_D, is then given by

$$L_D = \sqrt{\tau \frac{\langle s^2 \rangle}{t}} = \sqrt{2ZD\tau} \tag{3.29}$$

where Z is the dimensionality of the system, in this case $Z = 3$. It is, however, common for the factor of two to be omitted from calculations in the scientific literature [79]. It is also common in the literature to refer to the 1D diffusion length [215, 80]. In this thesis, the one-dimensional diffusion length shall always be given but the diffusion constant shall be given assuming

$$6D = \frac{\langle s^2 \rangle}{t} \tag{3.30}$$

and can be determined from the gradient of the mean-squared displacement as a function of time.

Chapter 4

Behaviour of the Structural and Dynamical Models

In this chapter, the SAMSEN method will be applied to systems of spheres in order to probe the effect of changing parameters upon a simulation and also to demonstrate the validity of the structural and dynamical model in comparison to those of conventional methods.

As established in chapter 3, SAMSEN simulations contain six parameters which can be readily changed, ϕ (the packing fraction), α (the ellipsoid inflation factor), ω (the frequency of the mode, controlled indirectly through $m_{m_{min}}^{m_{max}}$, ϵ (the 'effective' temperature of the mode displacement), and the collision and fitting thresholds χ_c and χ_f, and in principle, they describe the state point of a SAMSEN system. These parameters relate to, or are analogous with, the usual physical/computational parameters such as density, temperature, cut-off, although making the link to physical quantities can be challenging at times. This chapter is therefore dedicated to making this link and showing how changing these parameters alters the simulation and, therefore, what kind of states have been generated.

By using spheres, we can do away with the complexity of the shape of molecules and readily make comparisons with the literature on hard and soft spheres. For instance, we shall look at the changing timescales of the vibrational modes in the context of glasses and the structure (through the radial distribution function) in comparison to dense amorphous packings of hard spheres. This chapter will provide key insight which we shall draw upon throughout the rest of this thesis. We will return to molecules in the following chapter.

In the simulations that follow, systems of 8192 spheres of 1 $\mathring{A}$ radius were randomly inserted into a cubic box volume of sides 51.6 $\mathring{A}$. Unless specified otherwise, they were compressed following the method of section 3.3.2 at $C_f = 0.9999$ every step to a box volume of sides 38.316 $\mathring{A}$ ($\phi = 0.61$), calculating the modes only after compression ended, and the remaining parameters set to $\epsilon = 6$, $\alpha = 1.15$, $m_{m_{min}}^{m_{max}} = m_1^{12}$ (see section 3.2.7), $\chi_c = 0.001$ and $\chi_f = 0.00001$ (although this last parameter is irrelevant as each atom constitutes a single rigid section so in all circumstances, $\boldsymbol{r}_a - \boldsymbol{T}_a^s = 0$). For each of the variables we will change that one parameter in turn and study its effect. We will assume the spring constant, k and masses m are 1 and unitless throughout.

4.1 Changing ϕ

The first parameter, ϕ, is the packing fraction. This is simply the volume of objects in the system divided by the volume of the simulation box. This is a direct representation of the packing fraction used in simulations of spheres [132, 194, 179, 190, 193] and is similar to density, ρ. For simulations of molecules we may switch between ϕ and ρ, however for situations where the mass is less important, it is natural to choose ϕ to measure the sparseness of the simulation volume. When calculating the ellipsoid packing fraction, ϕ is the only choice and can take on values greater than 1.

The packing fraction is an important state parameter in simulations of glasses as the structural relaxation time has a strong dependence on ϕ [190]. This is the dynamical slow-down which, when increasing the packing/density, turns a liquid into a substance so slow moving that it appears to be solid at accessible timescales i.e. a glass [130, 132] (it should also be noted that temperature, pressure and the preparation pathway also alters the structure and properties during manufacturing [216, 130, 217]). Therefore, during a compression, we should expect to see the dynamics slowing down until, as written by Parisi and Zamponi [218],

> "the relaxation time is expected to diverge so that the system freezes in a metastable state, on the experimental timescale, [at the glass transition temperature]..."

This effect can be seen widely across molecular dynamics and Monte Carlo simulations of spheres [190, 219, 145]. The argument goes that particles get more and more confined until their neighbours trap them in position and motion is split into fast motion inside these new neighbour cages, and larger, slower, collective motions between cages [218]. This is the argument that motivates the use of SAMSEN as a method for sam-

pling a large number of states and overcoming this timescale barrier to accessing new states/morphologies.

If this slowing down of the dynamics until the time timescale for relaxation diverges has structural causes, it should be visible in the spectrum of vibrational modes. The Hessian only knows about the relative distribution of separations and the increase in potential for making particles move in certain directions. At high ϕ, the nearest neighbour peak in the radial distribution of hard spheres tends towards a delta function [132, 217], which means approximations about spring equilibrium lengths ($r_o = r_{ij}$) become accurate as they all approach a single fixed length. The radial distribution function is given by

$$g(r) = \frac{V}{N^2}\langle \sum_{i}^{N} \sum_{i \neq j}^{N} \delta(r - r_{ij}) \rangle \tag{4.1}$$

where δ is the Dirac delta-function and V is the volume of the system [220]. Figure 4-1 shows that we achieve such structures with the SAMSEN structural model using the compression procedure described in section 3.3.2. This means our calculated vibrational mode spectrum approaches an accurate representation that is greatly informed by the structure (in a vast energy landscape of many states).

To test if this dynamical slowdown could be observed, a randomly placed set of 8192 unit-radius spheres were compressed until collisions could not be resolved (the spheres were too too structurally arrested to be able to locally relax and maintain a valid state - we will consider this distinct from the jamming point, where there is one state with vanishing entropy, insofar as we will not establish that we are truly at this point and stop attempting to resolve collisions after 10,000 attempts). The radial distribution function for this state is shown in figure 4-1 and it has all the features of hard-sphere glasses approaching jamming [132, 217]. A value of $\alpha = 1.2$ was used for this compression and ϵ was set to zero. At each step during the compression, a Hessian was constructed using the ellipsoid intercept method (although in this case α simply represents a radial cut-off distance) and the eigenvalues calculated for the lowest 200 modes. The result is displayed in figure 4-2 as a function of ϕ (on the lower axis, we calculate ϕ using spheres of unit radius, on the top axis we calculate the packing as spheres of radius α).

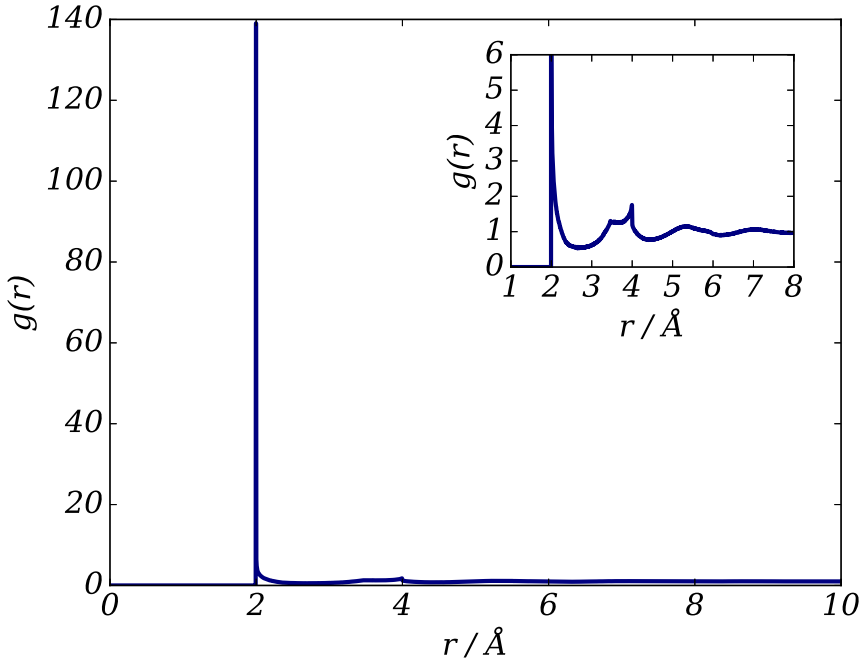

Figure 4-1: Radial distribution function for a system of 8192 spheres compressed until $\phi = 0.6162$. The nearest neighbour peak becomes very large as the system approaches jamming (like O'Hern et al. [132]) . A zoomed-in version is provided in the inset. These peaks broaden as the system gets sparser or when spheres get softer (see figure 4-19).

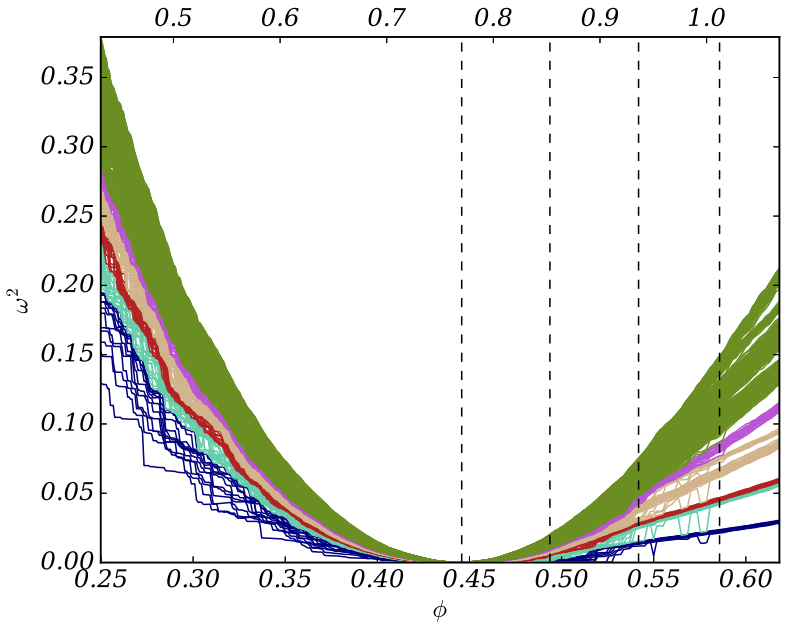

Figure 4-2: Frequencies of the vibrational modes for the hard sphere system during compression as a function of packing fraction (bottom) and inflated MVEE packing fraction (top) for $\alpha = 1.20$. The dotted lines indicates when the average number of contacts reaches 6 and larger integer values as ϕ is increased.

As can be seen clearly, the squared frequencies of the vibrational modes, ω^2, change dramatically throughout the compression. At low packing fraction, they begin in a range between 0.12 and 0.40). As the system is compressed, the frequencies begin to converge and decrease in value until $\phi \approx 0.442$ ($\phi \cdot \alpha^3 \approx 0.765$). At this point, the frequencies are not visible on this scale and all take values below 0.001. Beyond this point, the frequencies then increase in value until, at $\phi \approx 0.525$ where all ω^2 take values larger than 0.01. At $\phi \approx 0.54$, the frequencies cease to form a continuum and begin to split off into distinct bands. These bands increase in ω^2 linearly with ϕ while the remaining trunk continues to appear parabolic with increasing ϕ. The width of the bands also appears to depend on system size, getting narrower for larger systems, although I do not further examine that effect. On figure 4-2, I have also drawn vertical lines to mark where the average number of contacts takes on integer values of 6 and greater.

There are three key moments during this compression which we should take the time to examine further. The first is the approach to $\phi = 0.442$ where the frequencies reach their minimum. The approach to this point from smaller ϕ appears to match

the Parisi-Zamponi description almost perfectly. The timescales for the oscillations of the normal modes get longer and longer (the frequencies get smaller) until eventually they reach an impractically long time. This $\phi \approx 0.442$ point, is a packing fraction well below the expected glass transition density for hard spheres ($\phi \approx 0.58$), while the ellipsoid packing fraction $\phi \cdot \alpha^3$, takes on a value much larger than would be expected. However the observed increase in dynamical timescales up to this point would appear a good representation of the description of the glass transition and at this point the elastic network satisfies the Maxwell criterion for rigidity (where the average number of contacts is 6) [195] which would indicate mechanical stability. I therefore posit that this is the rigidity transition for the ellipsoid (network) system and that in the hard-sphere limit ($\chi_c \to 0$, $\alpha \to 1$) and with infinite attempts to resolve collisions, this inflection in $\omega^2(\phi)$ would be the jamming point. We shall return to this later.

The second is the passing of this rigidity transition, which leads to the reduction in vibrational period (increased frequency). It would appear that the return of mechanical motion occurs beyond the rigidity transition. Whether or not this is similar to vibrations in a crystal is unclear (it is an amorphous-looking solid and I am unable to find an obvious wavevector - Silbert et al. found many wavectors contributing to a given mode [192]). It could be a re-entrant melting phenomena, observed previously by Berthier et al. and Gnan et al. [190, 191]. In the case of Gnan et al. [191], their simulation was of a chain of particles forming a soft colloid and the deformation of the chain allowed this re-entrant melting to occur and the observation of ballistic-like dynamics beyond the glass transition, while Berthier et al. used a binary-mixture of soft-spheres.

The final moment which I shall remark upon is the appearance of vibrational bands at $\phi = 0.55$ which is also where the average number of contacts is 8, in this case. The bands formed take on varying number of modes, with the lowest containing 12 modes (coloured blue) and the next containing 24 (half aquamarine and half red), the next 14 and the next 8 (both coloured tan) and the next 12 (purple). There are more bands containing further modes, but an uncounted, number of modes. Vibrational bands have recently been observed in amorphous solids [221]. In a linear chain model with a basis where an acoustic and an optical vibrational band emerge, the acoustic modes are collective and the optical modes out of phase [201, 200]. While the analogy breaks down for amorphous solids and instead the modes are described as being acoustic-like or optical-like without a clear single associated wavevector, the question arises: are there analogous characteristic differences in the vibrational bands here also? We shall return to this later too.

4.2 Changing α

As was observed in the previous section, there appears to be some rigidity transition in the elastic network where the average number of contacts is 6 and the vibrational frequencies become vanishingly small. At higher densities we see that the dynamics begin to speed up again before observing the emergence of vibrational bands at yet higher densities. For an elipsoid inflation factor of $\alpha = 1.2$ (each axis is stretched by a factor of 1.2) we identified this rigidity transition at a value of ϕ outside the range of expected values (similarly for the value of $\phi \cdot \alpha^3$ which most closely represents the packing fraction of the elastic network, which we form from intercepts of the inflated ellipsoids).

For hard-spheres we expect to observe a glass transition at a packing fraction in the region of $\phi_g \approx 0.58$ above which exists a hard-sphere glass. Returning to Parisi-Zamponi, they state that a glass transition should occur at $\phi_g \approx 0.56$ and an ideal glass transition (the Kauzmann transition), ϕ_K, should occur around 0.58 to 0.62 [218]. At higher densities, rearrangements of the system become no longer possible and the system reaches a jammed state in the region of $0.62 < \phi_J < 0.67$ [218, 132].

The initial objective of this section is therefore to identify what the values of ϕ and $\phi \cdot \alpha^3$, at which the rigidity transition takes place, signify —have we indeed found the rigidity transition and does this compare at all with the expected Kauzmann density or jamming density? We shall attempt to do this by running the previous compression across many values of α until we find the packing fraction at which the observed rigidity transition takes place. For now, this shall be denoted ϕ_{RT}. In the $\alpha \to 1$ limit, we should expect to obtain the rigidity transition for a hard sphere system and therefore obtain an estimate for ϕ_J. For $\alpha = 1.0$, we also know that $\phi_{RT} = \phi = \phi \cdot \alpha^3$, but we have seen this is not true for all values of α. In which case other questions emerge: how does α effect ϕ_{RT}? And why isn't $\phi_{RT} \cdot \alpha^3$ always in the range of 0.58 to 0.62 if it really is the rigidity transition?

Let us begin by observing the effect of altering α on the spectrum of vibrational modes during compression. In figure 4-3, I have provided the 200 lowest-frequency modes during the compression from three different random configurations of 8192 unit spheres using the parameters and procedure as in the previous section, only altering α which is set to 1.005 (top) and 1.100 (middle). I have also reproduced the $\alpha = 1.20$ (bottom) simulation for convenience. As before, when the average number of contacts reach an integer number of 6 or greater, I have marked it with a vertical dashed line.

First, observing the different frequency scales in each of the simulations it almost appears, at first glance, that increasing α has the effect of stretching or translating the ϕ-axis. Perhaps this should be expected as α controls the packing fraction of the inflated system (α^3 is the relevant scaling factor between the soft and hard systems), although some volume is lost to overlap. However, it is not easy to find a point that is invariant to this scaling. It would be convenient if the stretch occurred around the minimum in the frequencies (the rigidity transition) on the inflated ϕ axes, however this point moves between all three graphs and differently across both the hard and ellipsoid packing fraction axes. This means that ϕ_{RT} according to the vibrational modes is changing. However, this point appears to be moving towards the expected value of ϕ_J (on both hard and ellipsoid scales) as $\alpha \to 1$ as expected.

Second, the emergence of the bands is not trivially related to α. At low α, the bands do not have time to appear at all, while for $\alpha = 1.10$ only some of the bands emerge. This limited emergence of some bands also occurs across a range of α. In figure 4-4 I have produced the frequency spectrum for a range of α at $\phi = 0.6162$. For small α the band structure is distorted by the presence of a large number of low-frequency modes. These *soft modes* have been observed in systems of soft spheres just below the jamming threshold when neighbour contacts are removed [176]. They occur when there are too few contacts to pin a particle in place (4 contacts per particle), in which case a localised low-energy mode exists. An example of a soft mode can be seen here for $\alpha = 1.04$ near $\omega^2 \approx 0$ in figure 4-4. At higher α (1.05 or greater) there are enough contacts in the elastic network to prevent these soft modes from forming and the band structure begins to emerge. However, the size and position of the bands varies until $\alpha \approx 1.15$, seen most notably in the tan-coloured band which begins adjoined to the purple band across most of the the range until it separates at around $\alpha = 1.13$. For $\alpha \geq 1.15$, the band structure appears relatively stable with only some slight variation in the higher frequency modes. In all cases the dynamics speed up after ϕ_{RT}.

Performing compressions for a large range of α, each with three repeats and with each system starting from its own random initial configuration, I have found the first instance where $\bar{N_c} >= 6$ (which corresponds to the inflection in the vibrational spectrum) and taken the value of $\phi = \phi_{RT}$ at that point for both the hard (unit radius spheres) and ellipsoid (spheres of radius α) volumes and, in figure 4-5, plotted them as a function of α. As can be discerned, there is a linear increase in ϕ_{RT} for the hard packing fraction at high α which then increases sharply as $\alpha \to 1$. Similarly the soft packing decreases linearly, turns sharply and meets the hard packing at $\alpha = 1.005$. This is slightly above the $\alpha = 1$ value expected and may be because of the value of χ_c which effectively

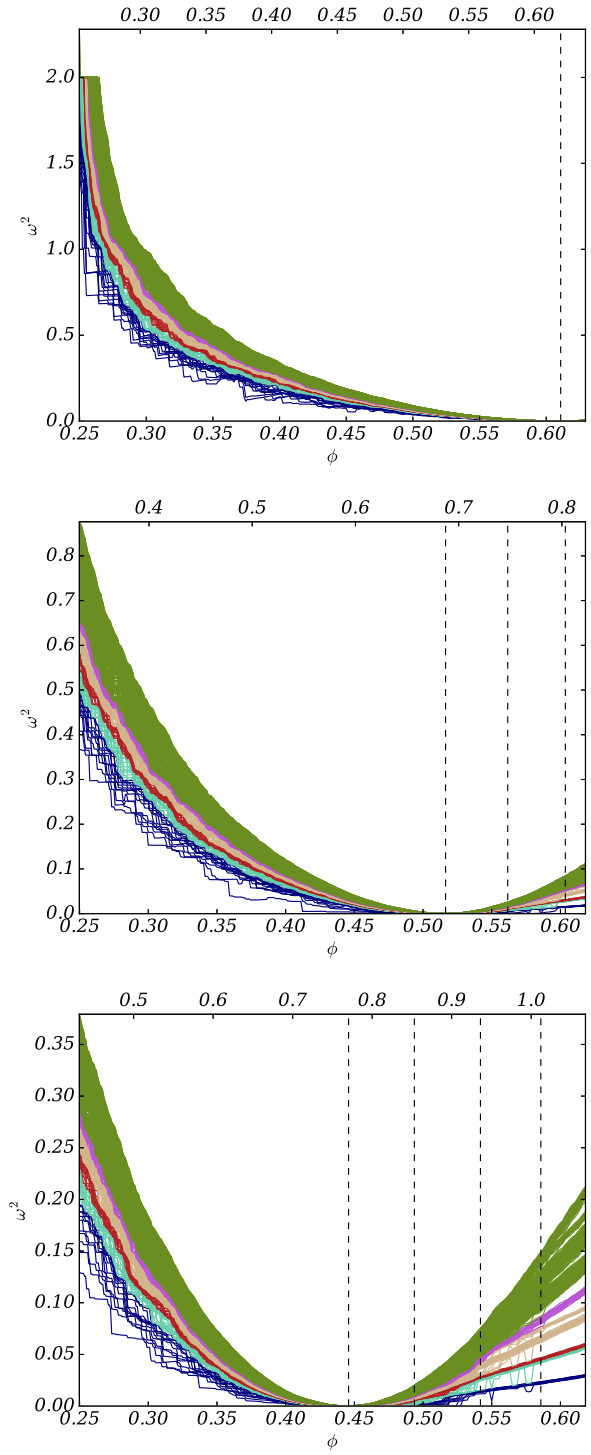

Figure 4-3: Calculated vibrational modes for the hard sphere system as a function of packing fraction (and soft/ellipsoid packing fraction) for $\alpha = 1.005$ (top), 1.100 (middle), 1.200 (bottom). The dotted line indicates when the average number of contacts reaches an integer number of 6 or greater. An inflection always occurs at $\bar{N}_c = 6$.

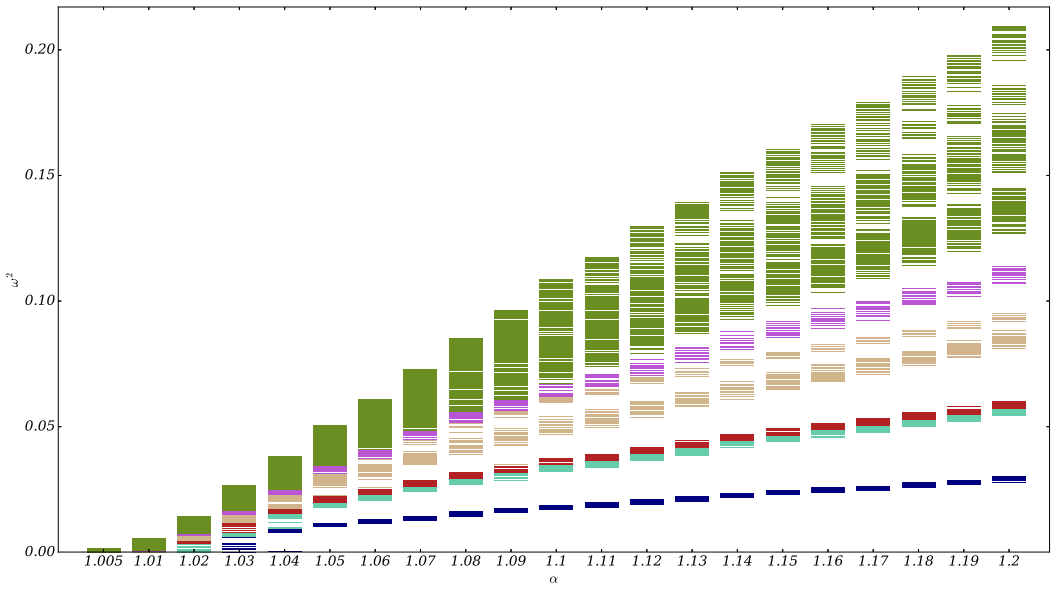

Figure 4-4: The spectrum of vibrational modes for α between 1.005 and 1.2 at $\phi = 0.6162$ for the 8192 sphere systems starting from a random configuration. Each mode is drawn with a width of 0.002, however some still overlap. Bands emerge containing 12 (blue), 24 (aquamarine and red, 12 each), 16 (tan, lower), 6 (tan, higher), 12 (purple), 48 (green, lowest), 48 (green, second lowest), 12 (green, third lowest) and 22 (green, highest) modes.

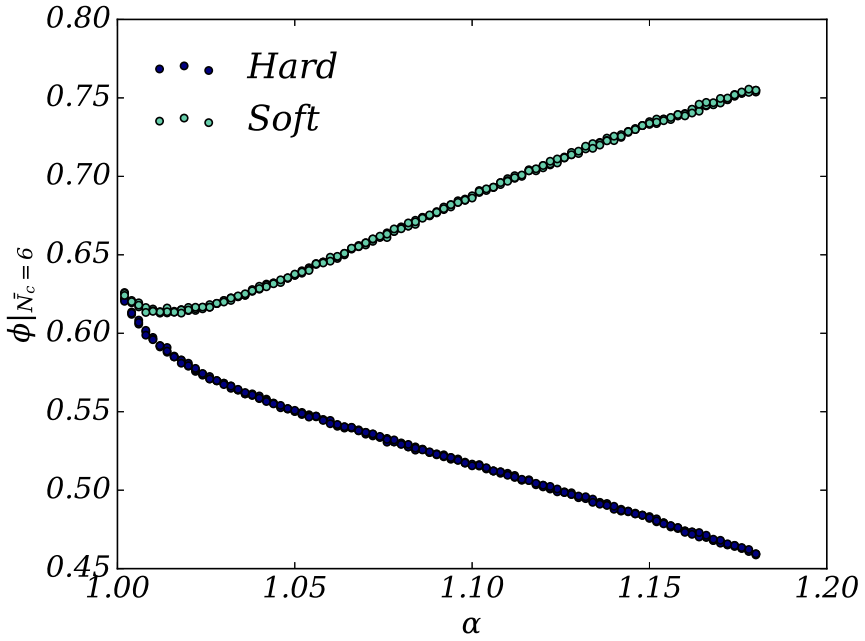

Figure 4-5: The packing fraction (hard and soft) at the inflection point, where $\bar{N}_c = 6$ (the rigidity transition), for a range of α values between 1.005 and 1.18. At $\alpha = 1.0$, the rigidity transition for hard spheres in SAMSEN (with a tolerance of $\chi_c = 0.001$) occurs. This corresponds to a packing fraction, $\phi = 0.63$.

lowers the value of the hard packing and increases the inflationary factor between that packing and the calculated soft packing. The difference in volumes is about 0.6% which accounts for the early intercept between the soft and hard packing.

Extrapolating the curves to $\alpha = 1$, the value of ϕ_{RT} is found to be 0.630 ± 0.001, which is on the high end of the expected range of ϕ_K but slightly lower than the jamming transition of hard spheres at 0.639 reported by O'Hern et al. [132]. This would suggest that, what we have been referring to as the rigidity transition, is indeed the rigidity transition of the ellipsoid system (elastic network) and for $\alpha = 1$ our system of approximately hard spheres would mechanically jammed at this point (although we never achieve this state computationally). We can now refer to this point as ϕ_J and we predict a value of $\phi_j = 0.63 \pm 0.001$ in a system of hard spheres.

It is in this light that the deviation in $\phi \cdot \alpha^3$ and ϕ_{RT} from the expected ϕ_J should be addressed. In the limit of $\alpha \to 1$, the soft system matches the hard system as the amount of overlap between the hard spheres (collisions) and the soft spheres (network cut-off) is approximately the same. This means that the structure of the soft system is perfectly represented by the hard system and the calculated vibrational modes are

those of a dense amorphous solid under a harmonic potential. At higher α, the amount of overlap between the two is different - the soft system overlaps with more neighbours than the hard system can geometrically come into contact with. As α is increased, ϕ_{RT} then takes on smaller values as the rigidity transition, in terms of the average number of neighbours in the elastic network, is reached at lower hard packing fractions, and the soft system deviates away from a hard sphere glass.

This deviation in behaviour should not be considered inherently problematic. Being able to control which dynamical regime we are in during simulation has some obvious benefits if one is only concerned with sampling a larger number of independent amorphous states. By increasing α one is able to ensure that there are enough contacts so that localised soft modes cannot exist and therefore do not contribute to the particle displacements. In the case of amorphous solids, a large number of low-frequency modes appear as systems approach the jamming/isostatic point - sometimes referred to as the Boson peak [176, 178, 179, 140]. In figure 4-6, it can be seen that SAMSEN is also able to detect the excess modes and, if the system is at a packing fraction either side of the inflection point in the vibrational mode spectrum, the anomalous low-frequency modes dominate the calculated spectra. These modes are highly localised [178, 222] (as we shall see later) and under equation 3.18 would produce very large displacements while the rest of the system remained roughly still. This behaviour is perhaps sub-optimal for producing system-wide relaxation. Instead, through α, one is able to choose a regime where an entire band exists and a vibrational spectrum that is is uninterrupted by soft or anomalous modes, with modes approximately constant in frequency, and therefore grants approximate control over the mean squared displacement of particles. In a moment, we shall see that by doing just this β_m^2, as defined in equation 3.18, can be controlled (through ω and ϵ) and used to determine the rate at which states are generated. Of course, if one is studying the dynamics of real-systems where the harmonic potential is a good representation then α will need to be selected appropriately (as we have attempted to do here for the comparison to hard-sphere systems).

Looking again at figure 4-4, it is highly reminiscent of figure 4-2 for $\phi > \phi_{RT}$. The squared frequencies increase increase almost linearly with α here, while in figure 4-2 they increased linearly with ϕ. It would again appear that increasing the packing fraction of the soft system by increasing α has a similar effect on the modes produced by the elastic network as increasing ϕ does. So while the regime where the modes split into bands, and the dynamics associated with those modes, is inaccessible to the hard sphere system, it is possible to recreate those dynamics by increasing the cut-off radius when constructing the elastic network. We will therefore fix ϕ for the sphere

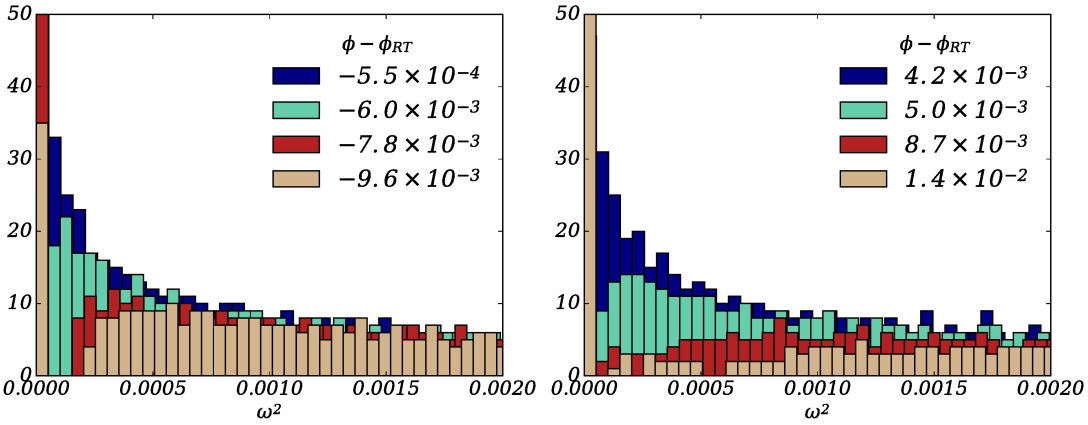

Figure 4-6: The population of vibrational modes for the $\alpha = 1.005$ system at several packing fractions around $\phi_{RT} \approx 0.6096$. The left pane depicts the approach to the isostatic point, $\phi = \phi_{RT}$, and the right pane shows the distribution of modes after compressing beyond this density.

systems and fix the density, ρ, for molecular systems to a desired density (ideally the physical density, if known) and control our position on the spectrum of vibrational modes through α.

4.3 Changing ω

Having now established that the frequency (ω) spectrum of vibrational modes have a band-like structure after ϕ_{RT} and that we can access them by increasing α (rather than constructing the Hessian based upon the physical contacts between touching spheres), it will be useful to probe the dynamics of such modes and attempt to establish how materially different these modes are from each other. Reviewing the literature [201, 197, 184, 221], one may expect that lower-frequency bands have more collective dynamics than higher frequency ones in a similar (but perhaps less discrete) way to the comparison between the acoustic and the optical branches in the linear-chain model with a basis. If collectivity can be found in these dense amorphous systems, then perhaps this provides us with a suitable pathway for making large displacements and achieving structural relaxation - even in materials approaching structural arrest. We expect this is the case and if so one would expect this to be a highly efficient pathway to generate new configurations and escape the dynamical slowdown experienced in conventional simulation methods (and in the physical world too).

I therefore begin by taking a new set of random configurations of 8192 unit radius

spheres and compress them down to a packing fraction of $\phi = 0.6162$ (a direct comparison to figure 4-4) and afterwards, for $\alpha = 1.20, 1.15, 1.10, 1.05, 1.005$, calculate the eigenvectors of the lowest 200 vibrational modes. I then calculate quantities which will allow us to assess the collectiveness of a mode: the average local dot product, the phase quotient and the participation ratio of the modes.

Briefly, the average local dot product, $\langle u_i \cdot u_j \rangle$, describes the how similar the average travelling direction of a given particle's neighbours are to that particle, the phase quotient, φ_m, is a system-wide measurement of vibrational coherence often used in the literature, and the participation ratio p_m considers the distribution of vibrational amplitude in the system. In all of these measures a value near 1 is the most collective motion. In figure 4-7, I present the average local dot product (top), the phase quotient (middle) and the participation ratio (bottom) against ω^2 of the given mode for the five values of α at $\phi = 0.6162$.

Starting with the average local dot product, you immediately get a sense of the highly collective nature of the lowest-frequency modes. For the lowest 12 non-trivial modes at all $\alpha \geq 1.1$, the eigenvector of a given sphere is pointing within 20 degrees of the eigenvector of a neighbouring sphere, on average. This means they are moving in a very similar direction to their neighbours. For $\alpha \geq 1.05$, you can also clearly see a drop in this local dot product as you go to higher frequency bands - although there are a couple of exceptions for $\alpha = 1.20$ and 1.15 in the upper tan band (m_{53}^{58}) and the third green band (m_{167}^{178}), referring to the colours used in figures 4-4 and 4-3. For the lower α values this does not occur and those bands do not exist. Here, the average local dot product falls off very quickly with frequency, significantly between the lowest bands, with the 200th mode of $\alpha = 1.05$ possessing an average local dot product of ≈ 0.55. Meanwhile $\alpha = 1.10$ yields $\langle u_i \cdot u_j \rangle \approx 0.65$, $\alpha = 1.15$ produces $\langle u_i \cdot u_j \rangle \approx 0.7$ and $\alpha = 1.2$ finds $\langle u_i \cdot u_j \rangle \approx 0.73$.

Similarly for the phase quotient, the smaller values of α produce much quicker reductions in collectivity throughout the system as ω increases. The distinction between the bands is less clear than in the average local dot product, although there are some points, such as the transition from the purple (m_{59}^{70}) to the higher frequency green bands, where there is a noticeable change - particularly in the gradient of $\varphi(\omega)$ (if such a quantity is meaningful). It would appear to me, looking at the plateau (particularly in the case of $\alpha = 1.10$), that a new type of mode appears at ω_{71} where the variability becomes becomes less directional - perhaps a new style of movement appears which is still mostly collective and in-phase but distinct from the modes at lower frequencies.

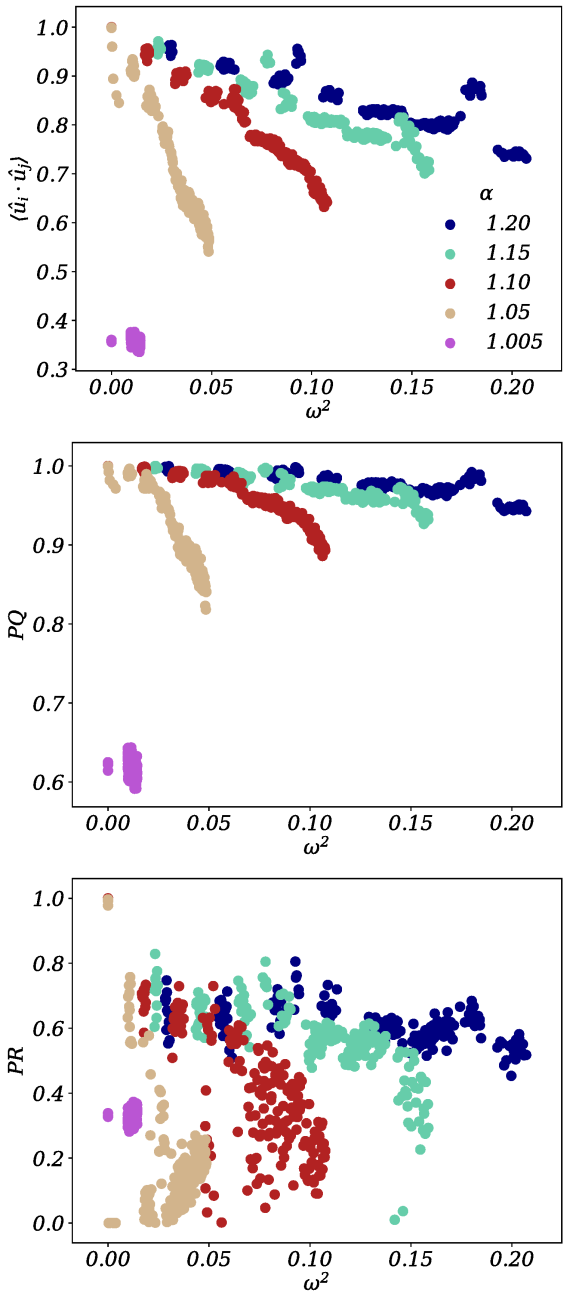

Figure 4-7: The local average dot product (top), phase quotient (middle) and participation ratio (bottom) as a function of ω^2 for the 200 lowest-frequency modes of 8192 spheres of unit-radius at $\phi = 0.6162$ at five values of α (1.005, 1.05, 1.10, 1.15, 1.20) immediately after compression. This behaviour is typical in SAMSEN simulations, including at many cycles after compression.

Looking at the mode participation ratios, it would appear that there is now a large variability in the distribution of amplitudes after ω_{71} for the $\alpha = 1.10$ case. This variability is a reduction, not an increase, in the uniformity of the distribution of amplitudes in comparison to the shape of $p_m(\omega)$ for higher values of α. This would suggest either that some particles are taking a role similar to nodes in an acoustic wave analogy and not moving at all, or very slowly, some particles have a very high amount of energy, or a combination of the two. For the remaining $\alpha \geq 1.15$ values the participation ratio is similar for low modes at all α, with a small drop in participation occurring in the green bands (m_{71}^{200}) and a much larger drop off for m_{179}^{200}. For $\alpha = 1.05$ this is not the case and, while the participation of lowest 12 modes (excluding translational and soft modes), appear similar to those of the higher α modes, they become much more localised for the remainder of the spectrum.

$\alpha = 1.005$ is the one simulation which consistently does not fit any of these patterns and approximately has a constant average local dot product (≈ 0.36), constant phase quotient (≈ 0.62) and constant participation ratio (≈ 0.38), for modes in the range m_1^{200}, including the soft modes where $\omega \approx 0$ and distinguish themselves from the translational modes by not having a value of 1 in any of the measures (it is possible the translational modes were obscured by soft modes). What is remarkable, is that there is now no obvious extended collective character, indicating these may be the localised/semi-localised modes in the Boson peak. Spheres are moving more coherently than random displacements but not at all as collectively as at $\alpha \geq 1.10$ for instance.

This perhaps gives credence to the benefit of displacing particles along the low-frequency normal modes at a chosen α and not just running molecular dynamics, for instance, for a long time. The benefit here is that we can calculate the normal modes for given structure, for which ever portion of the spectra we like. In the case of hard spheres, instead of being trapped with a lot of long-timescale processes which are not that collective and appear highly localised, we can use a larger cut-off (increase α) and then calculate the vibrational modes for the faster dynamics (such as those in the re-entrant fluid), which appears to have highly collective behaviour and should lead us to relatively-efficient structural relaxation. We also now know that we can find a regime where narrow vibrational bands emerge and chose which bands we want to use to displace the particles. Looking at figure 4-7, I would suggest using m_1^{12} at any $\alpha \geq 1.05$ as they appear to have the most collective character and I would expect would produce more computationally-efficient structural relaxation than higher frequency modes at this ϕ.

We can test this by performing a SAMSEN simulation, compressing the spheres to a solid density, calculating the normal modes, and displacing our spheres by setting $\epsilon > 0$.

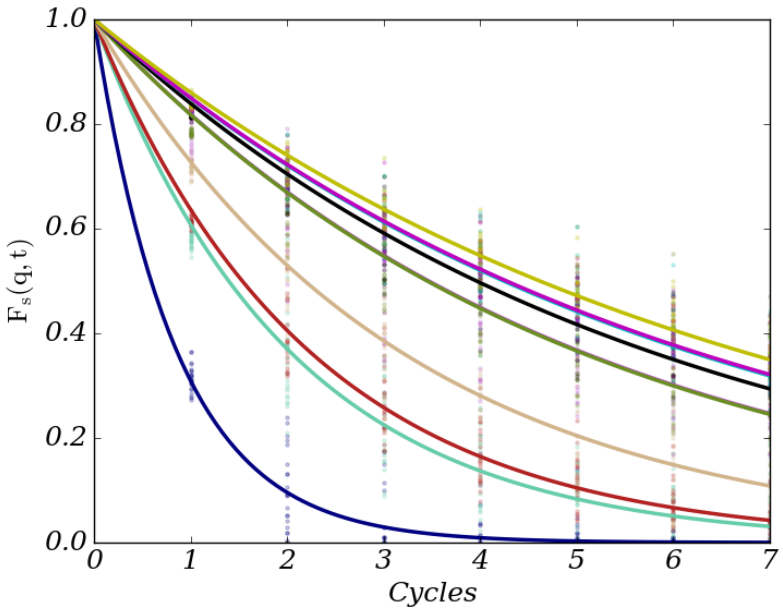

Figure 4-8: The intermediate scattering function as a function of the number of SAMSEN cycles for hard sphere simulations following different modes captured over 7 cycles. An exponential decay has been fitted to the raw data. The colours indicate the range of modules used (and the decay constant follows the same order) - from highest to lowest they are: blue (m_1^{12}), aquamarine (m_{13}^{24}), red(m_{25}^{36}), tan(m_{37}^{48}), magenta (m_{49}^{60}), olive (m_{61}^{72}), black (m_{73}^{84}), sky blue (m_{85}^{96}), pink (m_{97}^{108}) and yellow (m_{109}^{120}).

We can then follow the procedure outlined in chapter 3, recalculating the modes when we reach peak amplitude and continuing the displacement. By using a different set of modes, therefore altering ω and the collectivity of the displacements, we can observe how the structural relaxation time changes with mode frequency.

Another 10 randomly placed sets of 8192 unit radius spheres were randomly generated as before and compressed to $\phi = 0.61$. We set $\chi_c = 0.001$, as before, and we set α to 1.15. For the displacement, $\epsilon = 6.0$ and $S = 1000$. Each simulation used a different range of modes, but all summed over 12 modes. One simulation encompassed the entire first band (blue, m_1^{12}), another the lower half of the second (aquamarine, m_{13}^{24}), the third took the upper half of the second band (red, m_{25}^{36}), the fourth took the lowest twelve in the third band (lower tan, m_{37}^{48}), and the remaining took m_{49}^{60}, m_{61}^{72}, m_{73}^{84}, m_{85}^{96}, m_{97}^{108} m_{109}^{120}. The displacements were completed after S steps and a new set of modes recalculated. Approximately 200 of these mode displacement cycles were carried out before the simulation terminated.

To measure structural relaxation, the self-intermediate scattering function, $F_s(q,t)$, is commonly used. This measures the correlation of the particle positions between a chosen initial time and a later time and is expressed in terms of an inverse length q. The self-intermediate scattering function is given by

$$F_s(q,t) = \langle \frac{1}{N} \sum_{i}^{N} exp(-iq[r_i(t+t_0) - r_i(t_0)]) \rangle \tag{4.2}$$

where N is, again, the number of particles, $q = 2\pi/2R$ ($2R$ being the particle's diameter, in this case 2) and t_0 is some initial time to compare positions against and $t+t_0$ is some later time. $F_s(q,t) = 1$ at $t = 0$ and 0 when the structure has completely decorrelated. The structural relaxation time is usually given as when $F_s(q,t)$ falls to $1/e$ and that convention shall also be used here [183, 187, 223, 220].

In figure 4-8, I have measured $F_s(q,t)$ for each of the simulations at the end of every cycle up to a period of 8 cycles, resetting positions and then restarting the measurement for another 8 cycles. The raw data is provided as a scatter plot and a decaying exponential fit (least-squares) performed for each range of modes used. Looking at the self-intermediate scattering functions, it seems quite clear that the SAMSEN simulation following the lowest frequency modes m_1^{12} produces the fastest decay of $F_s(q,t)$ (in terms of the number of SAMSEN cycles) and therefore produces the most structurally decorrelated/independent states for a given number of cycles. This would make following the lowest frequency modes the most efficient way to generate many independent amorphous structures. For higher frequencies, the rate of decay is markedly reduced. However all of the simulations show a single step relaxation and do not appear to show any signs of structural arrest (such as a plateau in $F_s(q,t)$ at short times) usually found in simulations of glasses. It would imply that the dynamics are at least diffusive, something which could be verified by looking at the mean-squared displacement (which I do find indeed to be linear with time and we will study more closely for C_{60}).

The similarity in decay rates between m_{13}^{24} (aquamarine) and m_{25}^{32} (red), as well as the visual similarity to the band-structure in figure 4-4 raises an obvious question. Is the rate of relaxation simply proportional to $1/\omega^2$? If we refer back to the equation for the amplitude of the waves (equation 3.18), you will see that β_m^2 is proportional to $1/\omega^2$, so one would indeed expect relaxation time to depend on the frequency of the chosen mode(s). However, it is unclear how changing the collective behaviour changes the rate of decay of $F_s(q,t)$ - one would expect more collisions in higher frequency modes but it is unclear if that effect is larger or smaller than changing β_m^2.

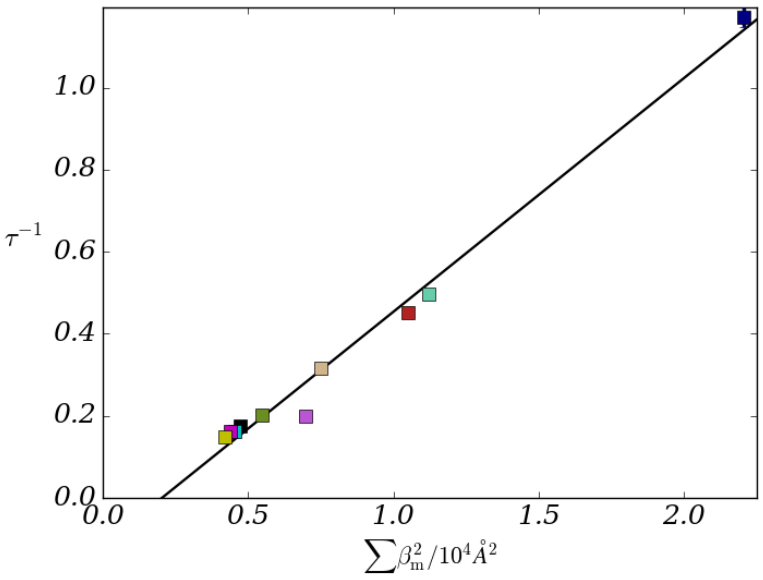

Figure 4-9: The reciprocal of the relaxation time, determined by an exponential fit to $F_s(q,t)$, against $\sum \beta_m^2$ (the applied mean squared displacement for each simulation following different modes).

To separate the effect of the collectivity of the modes and differing amplitudes upon the rate of structural relaxation, I have displayed the decay constant of the self-intermediate scattering function (the reciprocal of the structural relaxation time), τ^{-1}, against the sum of β_m^2 across the chosen range of modes in figure 4-9. This is therefore a direct comparison between applied displacement and relaxation time. While it is again made clearer that the lowest frequency modes (largest β_m^2) have the shortest relaxation time ($\tau < 1$ cycle), it should be immediately obvious that the relaxation is not directly proportional to the applied displacement (and therefore ω^2) as a linear fit shows a negative intercept on the τ^{-1} axis. This would imply that there is a frequency at which structural relaxation becomes impossible. However, a negative time is unphysical and I would instead expect an intercept at $\tau^{-1} = 0$ and therefore, at even higher frequency modes, would find that a linear fit is no longer sufficient. Either way, it appears that the highest frequency modes must suffer from the a similar kind of structural arrest that is experienced in conventional simulation methods (and physical processes) and one may even expect a plateau region in $F_s(q,t)$ at very high frequency modes. If the relaxation time were related to the amplitude alone (and not to other factors such as the collectivity of a mode) then one would expect direct proportionality, but we see something else here. This suggests that there is indeed some added benefit in

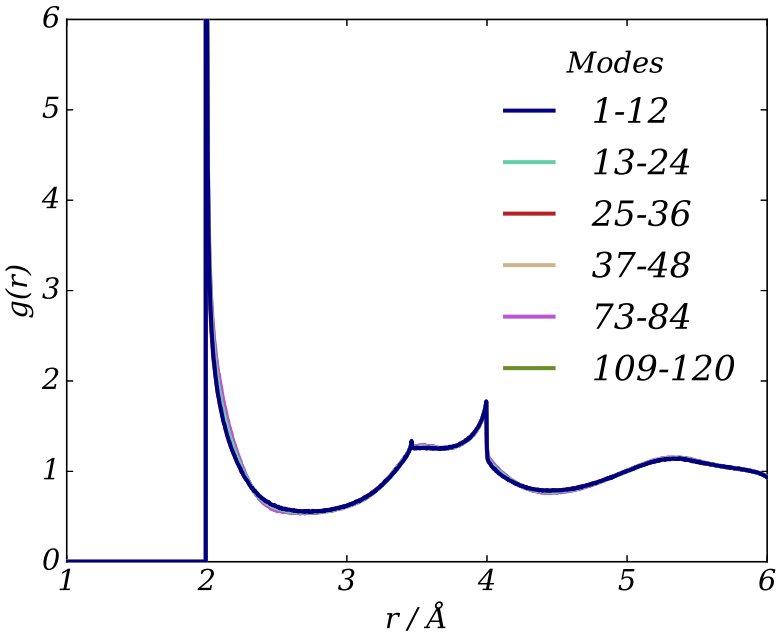

Figure 4-10: The radial distribution function, $g(r)$, averaged at the end of each cycle after compression for a selection of the mode-ranges studied, including the very lowest and highest-studied modes. The features of $g(r)$ are retained regardless of the mode-range used.

terms of structural relaxation caused by the character of the low frequency modes which cannot be explained by increasing the amplitude. It also cannot be explained in terms of differing structure, as across the range of modes studied, there is almost no change in measures such as the radial distribution function (shown in figure 4-10). In order to verify and quantify that it is the character of the modes themselves, we must subtract the effect of solely increasing the amplitude (which we shall do in a moment by varying ϵ over a single mode-band). However, at this point we can still say that the lowest frequency modes are the most efficient at generating structures per applied displacement, and therefore one may also expect them to be computationally efficient too.

Before I move on, I should be clear that the lowest frequency modes producing the fastest structural relaxation in SAMSEN, does not necessarily mean that these modes contribute the most to structural relaxation in molecular dynamics or in the real-world. The frequencies of these oscillations (and therefore the time period) are very different, such that the real contribution to structural relaxation may be less pronounced than shown here.

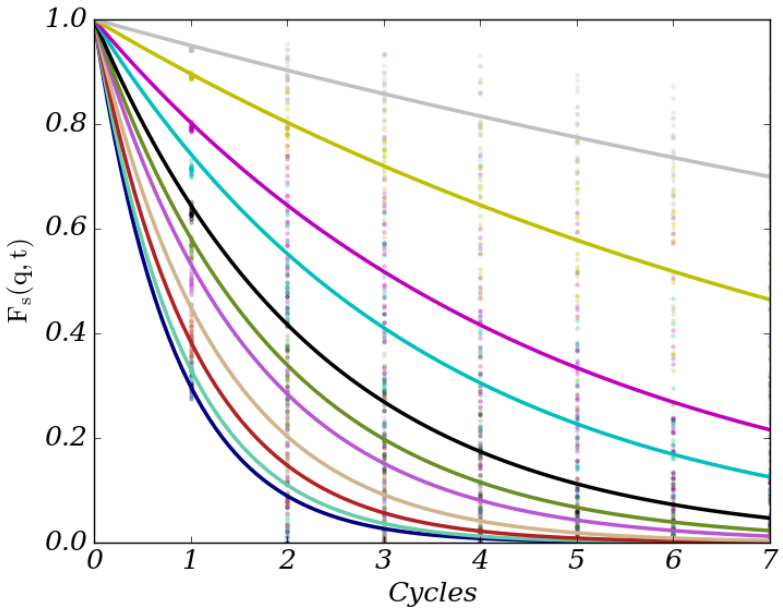

Figure 4-11: The intermediate scattering function, $F_s(q,t)$, as a function of the number of SAMSEN cycles for hard sphere simulations with different values of ϵ. As before, an exponential decay has been fitted to the raw data. This time the colours (with the addition of silver) go from highest to lowest ϵ and this corresponds to fastest to slowest decay of $F_s(q,t)$.

4.4 Changing ϵ

To finish examining the effect of the choice of vibrational modes on the structural relaxation time, it is important to quantify the effect of solely altering the assigned amplitude of the mode, controlled readily by ϵ through equation 3.18. This will then allow us to separate the effect of changing amplitude (via the frequency) and changing mode-character. We therefore follow the same procedure as before, compressing the spheres to $\phi = 0.61$ and with $\alpha = 1.15$, however we select the 12 modes with the lowest-frequencies (m_1^{12}, the entire blue band) and, for $\epsilon = 0.3, 0.6, 1.2, 1.8, 2.4, 3.0, 3.6, 4.2, 4.8, 5.4, 6.0$, displace the spheres over a period of $S = 1000$ and repeat this for the same number of cycles and produce the intermediate scattering function in the same form as before.

In figure 4-11 the fastest decays in $F_s(q,t)$ correspond to the highest values of ϵ. Therefore the fastest relaxation corresponds to the highest values of β_m^2 as before. However, looking at the relaxation times in figure 4-12 we now find something close to direct proportionality between amplitude and inverse relaxation time. If there is zero displacement it follows that relaxation never takes place, so an intercept at (0,0) should

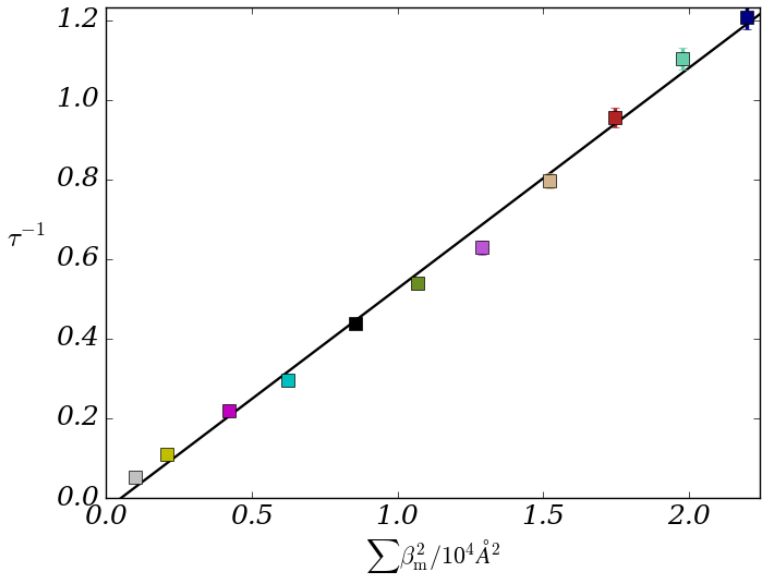

Figure 4-12: The relaxation time, determined by an exponential fit to $F_s(q,t)$, against β_m^2 for simulations of varying ϵ.

be found for smaller ϵ. The fact it does not might suggest that the fit is dominated by the variation in the middle of the range.

However, what this shows is that the gradient, and therefore the relationship between total displacement and inverse relaxation time, is stronger in figure 4-9 than in figure 4-12. This means that there is additional benefit in switching from a simulation using high frequency modes to one following low frequency modes, in terms of how many decorrelated states you are able to generate per SAMSEN cycle. In the case of these spheres at this ϕ we find it is only marginally more effective than increasing the amplitude through ϵ - around a 20% increase in the gradient - however, in the case of small molecules (most notably here, C_{60}) we find this effect to be much stronger.

That the relationship is linear for ϵ means that, in principle, almost full control of the relaxation time per cycle is provided by ϵ - something that α is unable to achieve by varying the spectrum (the soft, localised modes at $\alpha = 1.005$ for instance) and something which controlling ω^2 through m_{low}^{high} cannot achieve due to the discreteness of the bands. So with ϕ fixed by the material, α and ω limited by the desire for collectiveness, ϵ becomes the main lever for controlling dynamics. This is, however, entirely natural given its status as effective temperature for the dynamical model. We can once more look at the $g(r)$ for some values of ϵ across the range (figure 4-13)

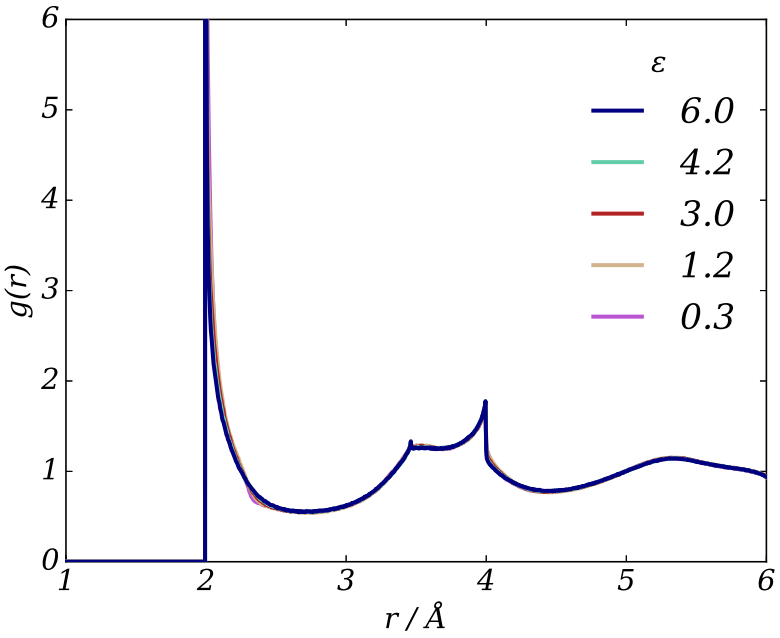

Figure 4-13: $g(r)$ for a subset of the ϵ values studied. The highest and lowest studied have been included as well as some in the middle of the range, the rest have been omitted for clarity only. Again, the features are fully retained across the range, although there is some alteration at $\epsilon = 0.3$ (the lowest epsilon) with an almost indistinguishable peak shift or broadening to higher separations and a slight dip at 2.3 angstroms.

and see that the structure is, again, almost undisturbed by changing the amplitude of displacements (perhaps there is minor broadening at low ϵ). This suggests that changing both ϵ and ω (the dynamical temperature) has little effect on the structure of the morphology (the structural temperature) which, as we shall see in a moment, is controlled by the χ's.

4.5 Changing χ_c

The last controls for the user are the χ's, controlling the amount of overlap permissible between unbound spheres and the range of separations between particles bound together by the rigid sections. In the case of hard spheres, the fitting threshold, χ_f, is meaningless as each particle constitutes an entire rigid section, leaving the collision threshold, χ_c, as the only parameter available to the user for controlling the temperature of the structural model.

As stated in section 3.2.3, χ_c, is the parameter which controls the amount of overlap

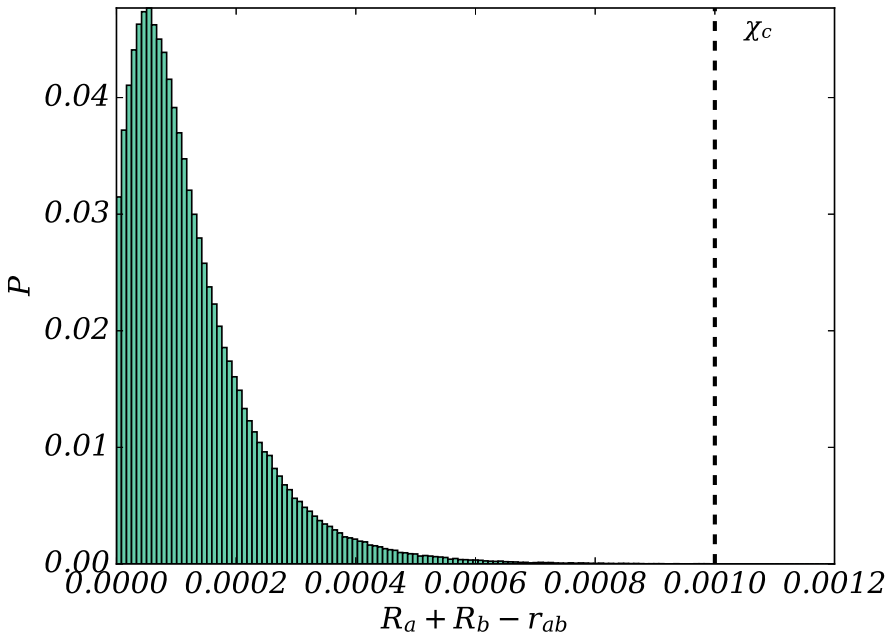

Figure 4-14: Distribution of collisions (overlap lengths) for the 8192 sphere system at $\phi = 0.61$ with $\chi_c = 0.001$, sampled over 30 cycles following m_1^{12} with $\epsilon = 6$.

between two atoms. This gives the user direct control over the closest approach between two atoms and, indirectly, control over the maximum potential increase that can be achieved by any atom approaching another. Once this maximum has been passed, all atoms in the system are displaced away from any atoms which share their volume - on an individual basis, by a distance so that they are just touching, but by the average of these if that atom is overlapping with many neighbours. This process is repeated until the collision condition is met. This creates a distribution of overlaps, with the largest at or below χ_c and, often, a modal and mean overlap length greater than 0. Such a distribution of overlap lengths is shown in figure 4-14 and has a mean overlap of 0.0001258 and a modal overlap of 0.00005452 with no overlaps greater than $\chi_c = 0.001$.

As a SAMSEN simulation has no simple description of potential or kinetic energy or temperature in the same way one might describe a molecular dynamics simulation having for a given force-field, we seek a different way of expressing (or at least representing) these quantities. In SAMSEN, the obvious proxies for quantities such as temperature are the χs, which as we have just seen for the case of the collision threshold, χ_c, in figure 4-14 represents the upper bound on the distribution of atomic overlaps. As we shall see in a moment, χ_c grants control over other aspects of this distribution, which (were a force-field applied) is related to the systems potential energy. The fitting threshold, χ_f,

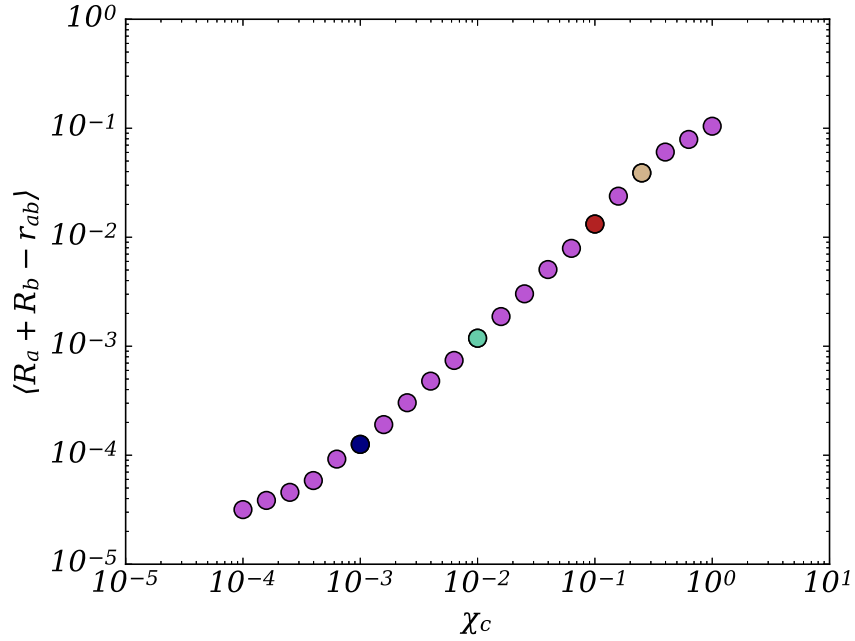

Figure 4-15: The mean collision length for values of χ_c in 8192 unit sphere systems at $\phi = 0.61$, sampled over 30 cycles following m_1^{12} with $\epsilon = 3$. The blue, aquamarine, red and tan points correspond to χ_c values of 0.001, 0.01, 0.1 and 0.25 respectively and produce radial distributions shown in figure 4-19.

which controls the rigidity of rigid sections also has an impact on the structures (which are considered in chapter 5). The χ's will therefore be referred to as the 'structural temperature' and should be considered an analogy rather than a direct comparison. We will investigate the effect of changing χ_c in this section on that distribution and consider its effect (if any) on the dynamics.

One can recreate the above distribution for a range of χ_c's and observe how the distribution in overlaps change. This is SAMSEN's analogy to observing how the potential energy of a system changes with temperature. In figure 4-15 I have plotted the mean overlap, averaged over all collisions across 30 cycles for 8192 spheres at $\phi = 0.61$ using $\alpha = 1.15$, m_1^{12} and $\epsilon = 3$ for 21 values of χ_c in the range 0.0001 to 1. This produces a smooth function mapping χ_c to mean overlap length. A similar and proportional mapping can be also produced from a comparison of the implied Maxwell-Boltzmann temperature and χ_c by fitting a Maxwell-Boltzmann distribution to the overlap distribution. There are three key regions of figure 4-15 which must be discussed.

The first region is the low-χ_c region where the relationship to mean overlap appears approximately linear. Assuming a linear relationship in this region, a fit to the first

4 points finds proportionality constant 0.112 and an intercept at 1.75×10^{-5}. That there is an intercept implies that, at a value of $\chi_c = 0$, there would not be any valid states (as the mean overlap would be greater than the maximum permissible overlap), even this far below the jamming density. This intercept represents the limit of the collision correction procedure's ability to minimise overlap. There is something inherent about the system being at this packing fraction which forces the spheres to have a non-zero overlap. This is perhaps the structural model's expression of pressure. By releasing the pressure (expanding the box) one would expect this intercept to reduce, at least up until the point where the low-χ_c region is indistinguishable from the mid-χ_c region and the high pressure no longer dominates the structure of the morphology. In mid-χ_c region, spanning several orders of magnitude, the mean overlap is directly proportional to χ_c. This region, where the permissible overlap is roughly 0.1% to 40% of the radius of the particle. As we shall see later, this middle region is where we shall perform simulations of molecular systems. Beyond this range, in the high-χ_c region with overlaps lengths of between 40% and 100% of the radius permitted, there is another inflection and a break in the linearity. This high-χ_c region is where the concept of structural temperature begins to fail and the structure becomes increasingly altered by the dynamical temperature - that of ω, ϵ and the nature of the modes. At $\chi_c = 2$, full overlap is permitted, pressure should be zero, and the dynamics - those of the low-frequency eigenmodes - entirely determine the relative positions of the particles.

The beginning of the dominance of dynamical temperature on the structural model can be seen as low as $\chi_c = 0.4$ in a few measures. The number of collisions with an overlap length of χ_c, for instance, does not approach 0 for these low χ_c values, the number of collision-corrections attempts required to achieve a valid state is less than ten (it is zero in the $\chi_c = 1$ case) in the latter half of the simulation. Most interesting, is the average number of collisions per step (averaged over collision-fitting minimisations) during the simulation, which is shown for 4 values of χ_c in figure 4-16. For $\chi_c = 0.001$ there is a brief period immediately after compression where the average number of collisions is high and the remaining period where it remains lower. More importantly, there is a pattern of sharp drops in the number of collisions followed by a sub-linear increase which has smooth form with time. This pattern repeats every SAMSEN cycle. However, for such low χ_c values, the change in the number of collisions at the end of a cycle and the beginning of the next cycle is largest. Meanwhile for mid-χ_c values, the correction is much smaller and, for the higher-χ_c values after the initial post-compression period, there is no sudden change in the average number of collisions and is a smooth function across both steps and cycles as collision conditions are always met throughout this part of the simulation (after the first 3 cycles). The distribution in separations is therefore,

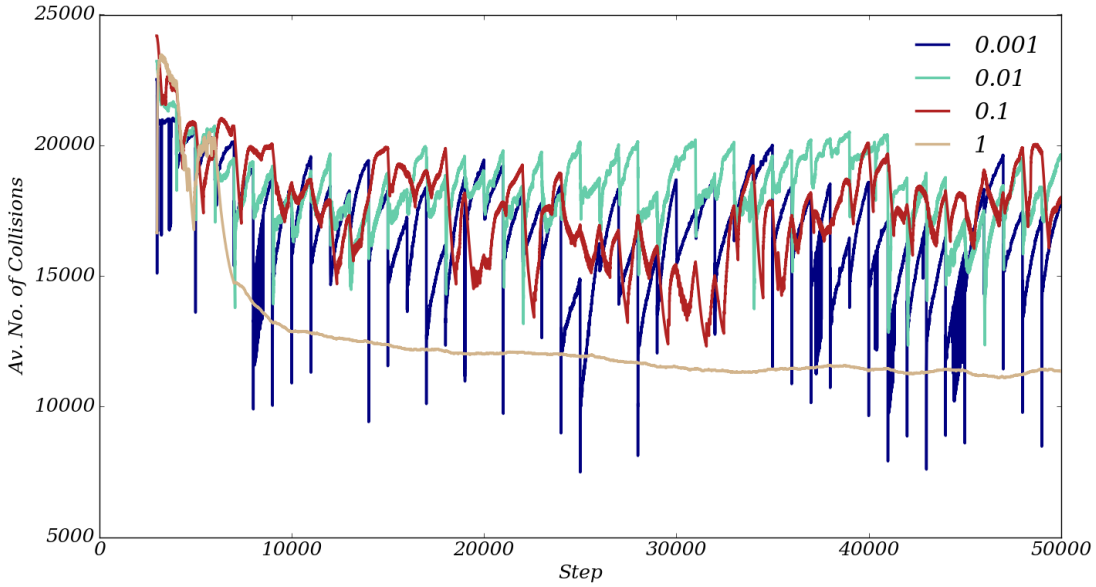

Figure 4-16: The average number of collisions per step (averaged over the fitting-collision attempts within each step) for $\chi_c = 0.001, 0.01, 0.01, 1$. The cycle period is $S = 1000$ and this may be seen as collisions build-up until the new mode calculation is performed.

for the most part if not entirely, governed by the low frequency modes for $\chi_c \geq R_a$ (although one must acknowledge the contribution of the collision condition, $r_{ij} \geq 1$ for $\chi_c = 1$, immediately after compression).

It should also be noted that this period of high collisions occurs for up to the 1st structural relaxation time immediately following compression. We therefore acknowledge that, in future results, it is important to stay clear of this post-compression region before beginning any averaging or sampling, just as one would wait for the system to equilibrate in molecular dynamics, for instance.

We can also study the reverse effect: what impact does the structural temperature have on the dynamical temperature? ϵ and ω represent the dynamical parameters of a SAMSEN simulation. ϵ (under the right conditions) might be considered a direct comparison to temperature. ω is closely linked to time. However, a further quantity, τ (the relaxation time) is also of interest as represents the time it takes for an object (in this chapter, spheres) to travel a characteristic distance. Therefore τ might be considered a good proxy for representing the kinetics. As τ is expressed in cycles which, in principle, have a link to time through ω and is dependent upon ϵ (which the user has direct control over) and also the collective behaviour of the modes, we

will refer to these quantities collectively and consider them as part of a 'dynamical temperature'. However, as with χ, this is to build an analogy rather than to provide a direct comparison with other methods. While the structural temperature relates to the severity of atomic overlaps, the dynamical temperature will relate to the speed at which structural units move throughout the system.

To examine the link between the structural parameters and the dynamical properties we can look at the effect of χ_c on the vibrational modes and on the relaxation time. One may expect that, with an increase in χ_c, comes an increase in the amount of overlap and an increased variation in the nearest neighbour distance (we shall verify this in a moment), so one may expect a change in the vibrational modes and therefore, at the same ϵ, a change in relaxation time. This is not what is seen. In figure 4-17 is a plot of the average frequency for the 50 lowest-frequency modes at each value of χ_c studied. All have similar values, similar variability and the bands have similar widths across the range of χ_c, with perhaps the exception of $\chi_c = 1$ which, as we have just discussed, is dominated by the dynamical temperature. So despite an increased disorder in the nearest neighbour separations, the modes do not change in frequency in response to a change in structural temperature in the linear region. For the same ϵ and the same ω, any observed change in the relaxation time could perhaps therefore be attributed to resistance caused by the collision condition. However, looking at figure 4-18 where, τ, the relaxation time (the decay constant of the intermediate scattering function) for each value of χ_c is shown, it would appear there is little correlation between χ_c and τ. There is perhaps weakly negative correlation, so that smaller χ_c values are more resistive to relaxation, but this is a small effect over many orders of magnitude of χ_c. This, perhaps, has the effect of lengthening the relaxation time by one cycle over the entire range of χ_c —the variance is itself much larger than this, however. If one assumes none of the data points are outliers, then the gradient of a linear fit is found to be very small across the range and there is very little dependence of τ upon χ_c. In both cases, structural temperature (particularly in the linear region of figure 4-15) has little-to-no effect upon the dynamical temperature.

Having examined the effect of changing χ_c within SAMSEN, we can finish by examining its effect upon the output structures (which has already been alluded to and has consequences for the remainder of this work). We shall mostly restrict ourselves to the the mid-χ_c region where structural temperature can be controlled, in great part, by χ_c. In figure 4-19, I present the radial distribution function for $\chi_c = 0.001, 0.01, 0.1, 0.25, 1$, zoomed so as not to capture the change in peak height but to focus on the change in shape of the various neighbour peaks. With lower values of χ_c the radial distribution

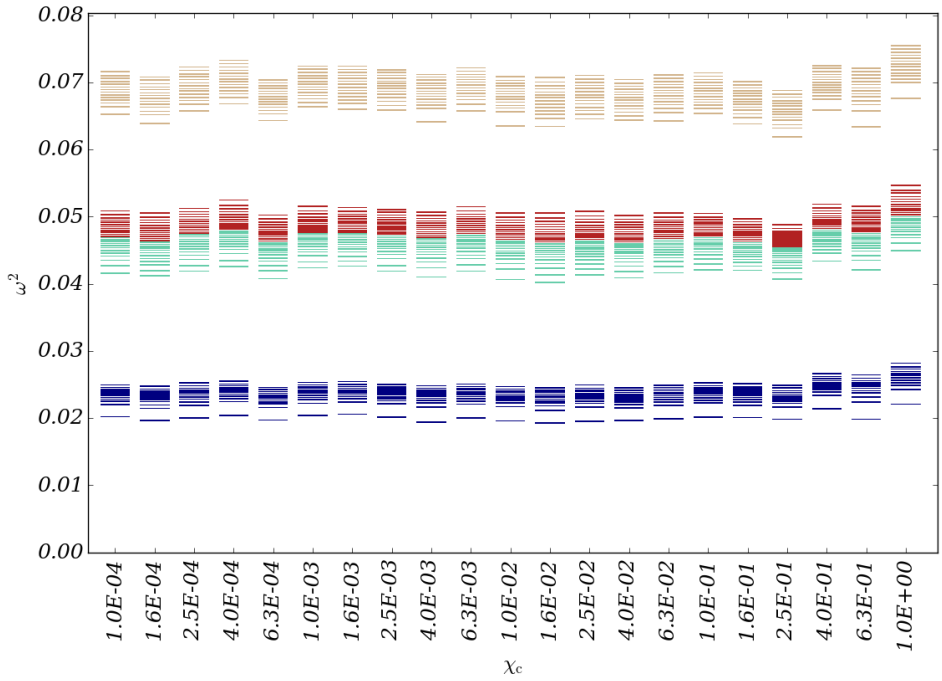

Figure 4-17: The spectrum of the lowest 50 vibrational modes averaged at the end of each cycle (22 cycles after compression) for each of simulations running different values of χ. With the exception of $\chi = 1$, which experiences a minor increase in frequency, the spectrum remains nearly constant for all values of χ.

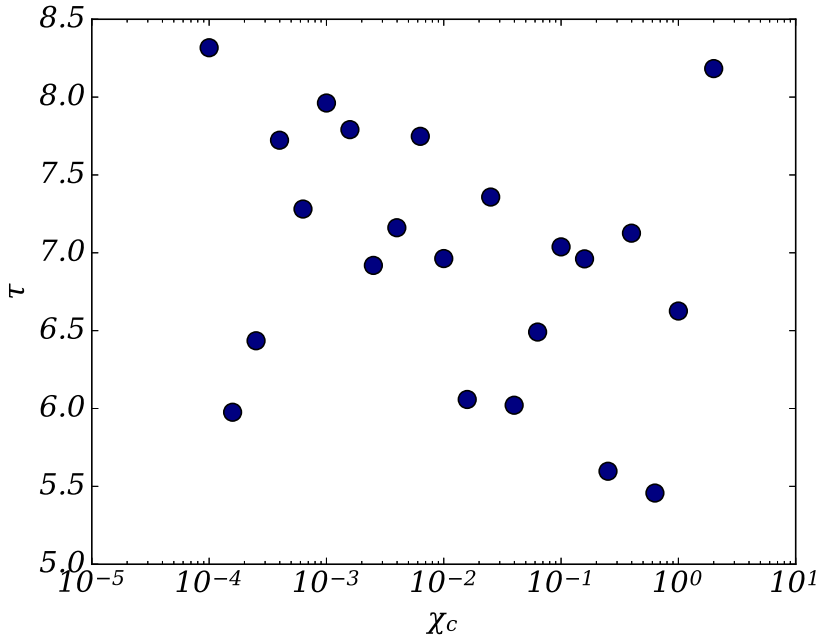

Figure 4-18: The relaxation time, τ, as a function of χ_c used in each of the simulations, determined from an exponential fit to the decay of $F_s(q,t)$.

looks like the O'Hern et al. result for an amorphous solid (I have reproduced this simulation from the beginning of the chapter) and, as χ_c (the structural temperature) is increased, the peaks all broaden and the $g(r)$ begins to resemble a liquid - a feat near-completed for $\chi_c = 1$, albeit with a skew to higher separations. We, of course, understand that the peak height diminishes when a peak broadens and this pattern across χ_c is true of the nearest neighbour peak which is largest at low-χ_c and smallest at high-χ_c (for $\chi_c = 1$ the amplitude of the nearest neighbour peak is below 5, while we recall it is approximately 140 for $\chi = 0.001$).

There are, of course, some differences in the shape of the nearest neighbour peak when compared to simulations of a Lennard-Jones spheres for all χ_c. Instead of forming a smooth crest, it is sharp at $r_{ij} = R_a + R_b$, and broadens unevenly on each side, with a larger density at larger distances. There is also a kink at 2.2 angstroms, followed by a divot at 2.3 in the high-χ_c distributions. This shows the limit of the ability of the collision condition/correction to capture the likely distribution of separations. This perhaps presents the limit of SAMSEN's applicability - it most accurately captures the structure of systems at high density and with small overlap (the same argument was also made about constructing the Hessian for the dynamical model at the beginning of the chapter). Were these details, smoothed out, only then would they appear to

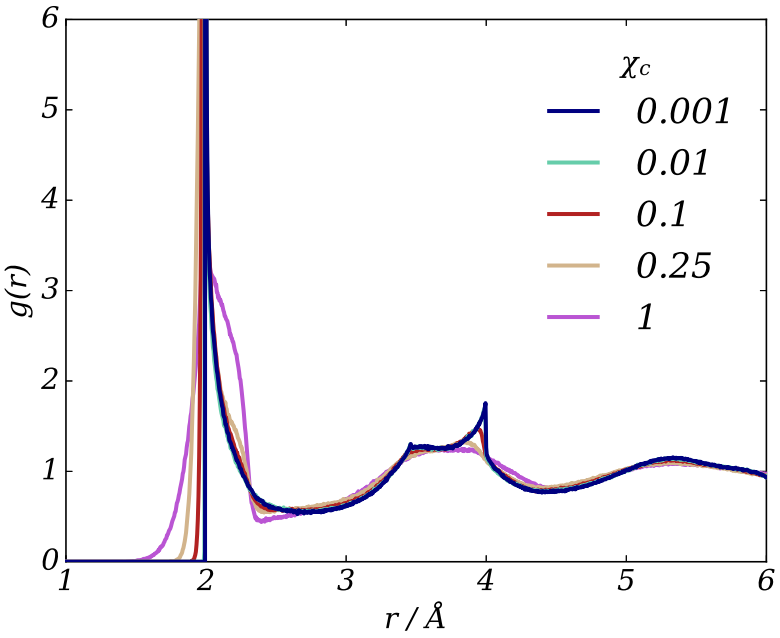

Figure 4-19: The radial distribution function, $g(r)$, for $\chi_c = 0.001, 0.01, 0.1, 0.25, 1$. Compared to $\chi_c = 0.001$, we now see large changes in the features, including diminished peak heights (from 140 at $\chi_c = 0.001$ to 5 at $\chi_c = 1$), broadening (with a bias to higher separations, a transition from a hard-sphere system near jamming to something resembling a liquid, and a divot forming at $r_{ij} = 2.3$ as χ_c is increased. For $\chi_c = 0.001$ and 0.01 there is little noticeable change (except in the nearest neighbour peak height) and for $\chi_c = 0.1$ there is a similar story, except in the peak at $\approx 3.95\mathring{A}$ and the broadening at $r_{ij} < 2\mathring{A}$

look like the expected radial distribution. We will later show, for small molecules, how this can be achieved and how these states relate in structure to those on the energy landscape in a classical force-field.

4.6 Summary

Through studying SAMSEN simulations of spheres, we have gained some insight into how the structural and dynamical models behave across a wide parameter-space.

We understand that both components of SAMSEN require high-ϕ (at least locally) for the approximations to produce results that compare well to known results. We argue that this is because a narrower distribution of nearest-neighbour separations means the network is better-represented by the harmonic approximation and the equivalent

equilibrium length approximation and that the lack of long-range interactions means that the generated structures are only informed by collisions at small distances. This is not to worry us, as many of the molecular systems we intend to study are at high density and have weak electrostatic interactions.

We also see that increasing α has a similar effect upon the vibrational modes to increasing the density (it is increasing the packing fraction of the ellipsoid system and therefore determines the vibrational spectrum produced). By using this parameter, we are able to ensure that our network is above its rigidity transition, without localised soft modes, and produces bands of vibrational modes - each with a different collective nature. Using the most collective band (the lowest ω band), we find that structural relaxation (in terms of decorrelation in a particle's position) is achieved fastest and that going from high-frequency to low-frequency modes, increases the rate of relaxation faster than simply increasing ϵ (and therefore increasing the amplitude of the modes). However, that does not mean we cannot still further increase the rate of relaxation by increasing ϵ.

We then considered the effect of structural temperature through χ_c and how it is, for the most part, able to determine the output structures, distribution of overlaps and the number of collisions without altering the vibrational frequencies or significantly altering the relaxation time. This means the method allows amorphous structures to be generated, with some influence on desired properties (such as width of nearest neighbour peak), without altering the rate at which these states are generated. We also find the opposite is true - that the dynamical temperature can be changed without significantly altering the structure.

In considering small molecules in a moment, we will attempt to draw similar comparisons to the results found here. I hope that, for anyone using this method, that these pages have so far provided an understanding of the role of each of the parameters, appreciation that control over parameters such as χ and ϕ is limited by the systems you intend to study and that a prescription for α and ω has been provided, and that ϵ is, in principle, the only variable one needs to consider before running a simulation and that this has, for the most part, no effect on the structures produced (only the rate at which they are produced).

Therefore, if one is only considering studying structures and not dynamics, one does not need to put any effort into finding the correct parameters. With χ_c and χ_f extracted from an appropriate force-field such as the General Amber Force-Field [121] and density determined for you by the system-at-hand, one can leave the dynamical parameters α,

ω and ϵ entirely unchanged across a range of molecules.

Chapter 5

Generating Structures of Small Molecules

This chapter intends to demonstrate the wider applicability of SAMSEN by applying the method to three small molecule organic semiconductors still used in current OSC architectures. This chapter is expands upon work previously published in the Journal of Chemical Physics, 150(16) 2019 [46]. Here we shall study C_{60}, α-NPD and PCBM, starting from random low-density configurations, compressed to their respective thin-film densities and displaced using the SAMSEN dynamical model to produce a series of states which can be used to compare SAMSEN structures to those of other simulation methods.

The SAMSEN parameters described in the previous section were set to $\alpha = 1.2$, m_1^{12}, $\epsilon = 200$ K, $\chi_c = 0.6$ $\mathring{A}$, $\chi_f = 0.2$ $\mathring{A}$, $M_i = 1$ a.m.u., $k = 1$ kg ps^{-2}, $S = 1,000$, $C_f = 0.99995$ for all systems studied, except where noted for the C_{60} simulations.

The optimised structures for these molecules were created by downloading structure files from the Royal Society of Chemistry's ChemSpider service [224] and minimising them in the Universal Force Field using the Avogadro package [225], with the exception of PCBM which was minimised in the Large-scale Atoic/Molecular Massively Parallel Simulator (LAMMPS) [226] using an OPLS-AA (Optimised Potentials for Liquid Simulation-all atom) force-field [227] previously published by Cheung & Troisi [24] who's structures we will compare SAMSEN structures to.

We will then go one step further and try to recreate the Cheung & Troisi PCBM structures by running a short molecular dynamics simulation using the SAMSEN states as

inputs. We will then consider the distribution of energy minima sampled by SAMSEN and how that may relate to the thermodynamic distribution of states. In later sections we will consider how the position on the energy landscape affects the charge transport properties in the amorphous phase and consider SAMSEN's treatment of polymers and polymer blends.

5.1 C_{60}

One thousand buckminsterfullerene (C_{60}) molecules were placed randomly as described in section 3.3.1. C_{60} forms a single rigid section (due to it being completely cyclic) so $N = 1,000$ and the MVEE for the rigid section represents that for the entire molecule and has axis lengths 5.23 $\mathring{A}$, 5.21 $\mathring{A}$ and 5.17 $\mathring{A}$. The packing fraction of the initial configuration was slightly below 0.25. This initial system was then compressed as described in section 3.3.2 until the collision-fitting conditions could no longer be maintained. This occurred at a packing fraction of $\phi_H = 0.627$. Here, for C_{60} we define the hard packing fraction, ϕ_H, as the packing fraction of hard spheres with radius equal to the half the nearest neighbour separation (inferred from the peak position of the radial distribution function) at this maximum density - this is measured to be ≈ 4.25 $\mathring{A}$. We will also make use of the ellipsoid packing fraction using the inflated MVEE axes to calculate the total ellipsoid volume and try to use both values to make comparisons to the unit sphere systems studied in the previous section.

5.1.1 Vibrational Modes During Compression

Performing the compression of C_{60} with $\epsilon = 0$, the elastic network was constructed at each compression step and the 200 lowest-frequency eigenvalues were calculated. In figure 5-1, I have plotted the values of these modes (using the same colour scheme as with the unit spheres) as a function of the hard-packing fraction (bottom), ϕ_H, and the soft (ellipsoid) packing fraction (top), $\phi_E = \alpha^3 \cdot NV_{E_s}/(L_x L_y L_z)$ where V_{E_s} is the volume of the C_{60} MVEE. In this case, because the hard-packing fraction is calculated as the effective radius, $\phi_E \neq \phi_H \cdot \alpha^3$.

From the frequencies found during compression, we find that there are some features common to both the unit spheres and C_{60}. Beginning with a decrease in frequency, there follows an increase in the vibrational frequencies after $\phi_H = 0.342, \phi_E \approx 0.75$ which we previously identified as indicating the elastic network's rigidity transition.

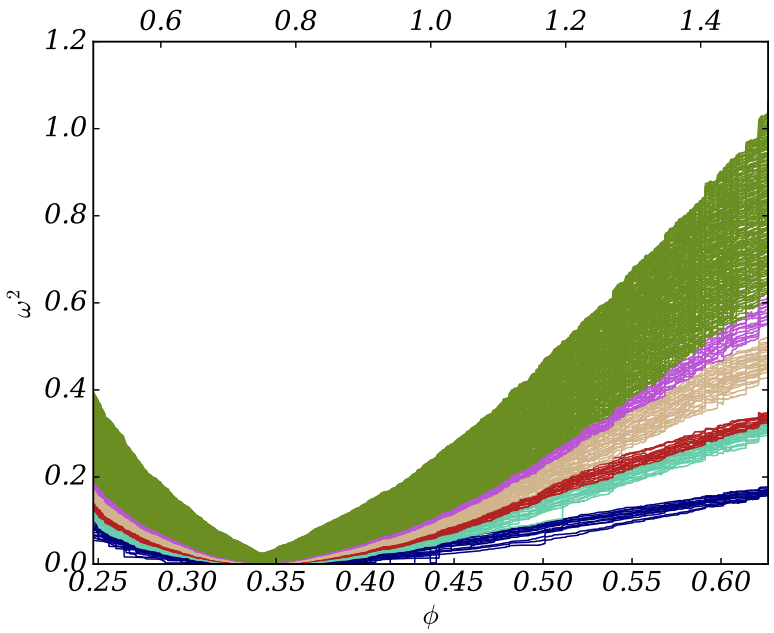

Figure 5-1: The frequencies of the 200 lowest-frequency vibrational modes for the C_{60} system during compression at $\alpha = 1.2$.[1]

As the modes begin the increase in frequency, we find zero-frequency (and non-trivial) modes - the soft modes - which are present up to $\phi = 0.44$ where the fewest number of network contacts reaches four - the minimum number required to pin a particle in place (the condition of local rigidity) [176]. There is also a quadratic increase in ω^2 around this point and is a linear increase for each of the vibrational bands as they split off from the main trunk as ϕ is increased. The band structure takes a very similar form, with the exception that the tan band is not split into a lower and upper half (as in the case of the spheres) - something which could possibly be attributed to system size. The magenta band is also slow at breaking away from the main tranche of modes and only does so at the very last moment, where the system approaches jamming at $\phi_H = 0.627$.

The appearance of these bands in molecular systems (they are also found in α-NPD, PCBM and P3HT, for instance) justifies the approach suggested in the previous section: that there is a natural choice of low-frequency modes to use which we expect will achieve the fastest structural relaxation and that we can access these by creating an elastic network of inflated MVEE contacts.

[1]Reproduced from Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. Journal of Chemical Physics, 150(16):164115, 2019, with the permission of AIP Publishing

5.1.2 Structural and Dynamical Model at High Density

Using the configuration at $\phi_H = 0.62$, five simulations were performed using different modes ranges at $\epsilon = 200$ K, each for 200 cycles. The modes used were m_1^{12}, m_{13}^{24}, m_{25}^{36}, m_{37}^{45} and m_{59}^{70} which represent the lowest 12 within each band and the upper half of the second band (as coloured but excluding the top 10 most in the tan band and the remaining $3N - 73$ modes). This density was chosen as, after a cycle, there is a minor increase in frequencies and the magenta band can be distinctly identified. Beyond a cycle the modes remained approximately constant in frequency.

In figure 5-2, the mean-squared displacement, the self-intermediate scattering function (using $q = 2\pi/9.5$ $\mathring{A}^{-1}$ as the length-scale to probe correlations) and the radial distribution function are shown for each of the simulations following their given mode ranges. These plots are analogous to those in section 4.3 and reproduced from Smith et al. [46]. Figure 5-2a shows the mean-squared displacement over a 20 cycle period, averaged over 10 lots of 20 cycle periods. For all modes ranges used, this appears completely linear, even at this high density, which would imply that our system is diffusive and unhindered by nearest neighbour cages. This is confirmed by the intermediate scattering function in figure 5-2b which follows an exponential decay and shows the system is able to avoid a two-step relaxation. The lowest frequency modes (blue) produce the fastest structural relaxation (fastest decay rate constant, τ^{-1}) which takes approximately three cycles and can be compared between modes ranges in the inset. Here, the applied displacement $\sum \beta_m^2$ is shown against τ^{-1} and, again, we find a linear relation and that there is a negative intercept on the τ^{-1} axis - implying the highest frequency modes produce are unable to relax the structure or, as is more likely, there is a different regime here where the high frequency modes are just highly inefficient per displacement applied. As it stands this picture looks exactly like the one presented for the unit spheres and the conclusions drawn about the dynamical model appear to apply to C_{60} too.

This direct analogy also appears to extend to the structural model. Looking at the radial distribution function for C_{60} (which was averaged ever 5 cycles) in figure 5-2c and comparing back to figure 4-10, we find a very similar structure with C_{60} also resembling an amorphous configuration of spheres at high-density. The sharpness of the peak is much lower than the unit sphere system owing to this difference in density (the difference in the measure of packing too) and C_{60} possessing many flat faces rather than a smooth surface. However, we also find that C_{60} demonstrates a similar resistance to changes in the features of $g(r)$ as ω is varied - through the majority of

the structure there is perhaps little or no change with the only noticeable differences being in the peak height and an almost imperceptible divot between 10 and 10.5 $\mathring{A}$. For clarity, an inset has been provided which has zoomed in on the peak at nearest neighbour separation. From this it can be gleaned that the lowest frequency modes (with the largest displacements, but the most highly collective motion) produce the sharpest structure which is marginally diminished as the frequencies are increased.

We can also once again try to distinguish the benefit of exploiting the collective character of the low-frequency modes from the increased applied amplitude under equation 3.18. We can, as before, vary ϵ for a fixed range of modes and compare the relationship between applied displacement and relaxation time. We can also use this to finally attempt to observe if the dynamical temperature produces noticeable changes to the structural temperature. For this set of simulations, modes were used in the range m_1^{24} and ϵ set to 50, 100, 200, 250 and 300 K.

In figure 5-3 we present both the radial distribution function and intermediate scattering function ($q = 2\pi/9.5\ \mathring{A}^{-1}$) for each of the ϵ simulations. Again in the $g(r)$ shown in figure 5-3(a), there is a very similar structure found between the simulations, regardless of the value of ϵ used, with only a minor change in peak height for the nearest neighbour peak (shown again in the inset) - perhaps less obvious than in the ω-varying simulations - and with the highest value of ϵ (the simulation with the highest amplitudes) producing the sharpest structures. This counter-intuitive behaviour was seen before in section 4.4.

In the intermediate scattering function (figure 5-3(b)), we find, naturally, that the largest amplitude simulations produce the fastest structural relaxation and that there is a linear relationship between applied amplitude and relaxation time (inset). We also find that this is, almost, direct proportionality (there is a slightly positive intercept, which is impossible because at zero amplitude, there is zero motion and therefore $\tau = \infty, \tau^{-1} = 0$, so direct proportionality, $\tau^{-1} = C \cdot \sum \beta_m^2$, may be the expected behaviour). Comparing the gradients of both the ϵ-varying and ω-varying simulations, we find that the gradients are 1.63×10^{-7} and $3.87 \times 10^{-7} \mathring{A}^{-2}\text{cycle}^{-1}$ respectively, making the effect of increasing the amplitude at the same time as increasing the collectivity more than twice as influential on the relaxation time as increasing the amplitude alone.

The recipe for generating states of small molecules is therefore the same as spheres: use a value of α that puts you in a regime where modes split into distinct bands, follow

[2]Reproduced from Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. Journal of Chemical Physics, 150(16):164115, 2019, with the permission of AIP Publishing

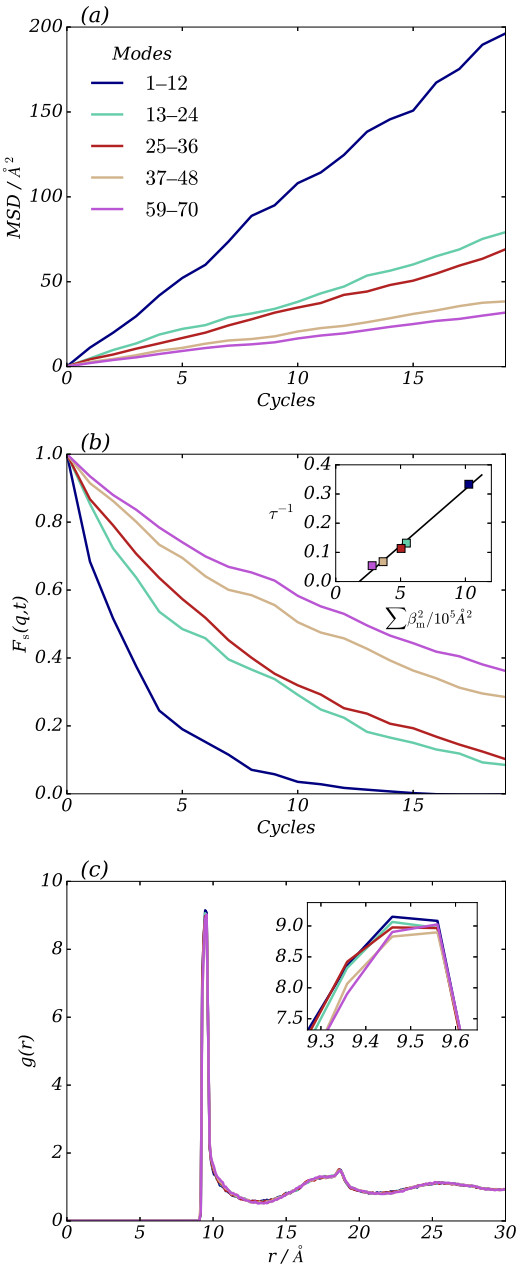

Figure 5-2: (a) the mean-squared displacement, (b) the self-intermediate scattering function and (c) the radial distribution function for the C_{60} system following different ranges of modes, each containing 12 modes. The inset of (b) shows the reciprocal of the relaxation time against total applied displacement and the inset of (c) shows an enlarged view of the nearest neighbour peak in the $g(r)$.[2]

the lowest-frequency band and apply a large ϵ. This will not (significantly) alter the output structures generated, but let you generate many and do so very quickly. The short-range structures are dominated by the value of χ_c picked, while the long-range structure resembles a dense liquid or amorphous solid.

Distribution of collisions and mismatches

One must recall that in the unit sphere simulation, we varied χ_c in order to investigate the effect it had on the structure produced. Looking at $g(r)$ when both varying ω and ϵ, we see that they produced near identical radial distribution functions. This again is true of the C_{60} simulations and it would again appear that the 'temperature' of the dynamical model is very much separate from that of the structural model. Of course, there are differences in the structure (between a sphere and a fullerene, as we have already noted) but, if one were to try and make the comparison, it would be hard not to suggest that the C_{60} $g(r)$ using $\chi_c = 0.6$ $\mathring{A}$ most closely resembles the unit sphere $g(r)$ at $\chi_c = 0.25$ and would probably be closest to one in the range $0.1 < \chi_c < 0.25$. This can again be seen in the distribution of particle overlaps for the m_1^{12} simulation, shown in figure 5-4, with very few overlap lengths near χ_c but a modal overlap of approximately $\chi_c/22$ and a mean of 0.0655 $\mathring{A}$, placing this distribution in the linear region of mid-range χ_c values. This is perhaps surprising, as χ_c/R is closer to 0.5 for the carbon atoms in C_{60} so one may have expected a broader, smoother $g(r)$ as was the case with the spheres. However, this is perhaps due to the poor quality of the comparison and this simply demonstrates the enhancement of the local structure by having many atoms bound together in rigid sections. We shall come back to this thought when considering α-NPD.

α-NPD will also present an opportunity to examine the distribution of mismatches. In the case of C_{60}, there is a single rigid section and, as $\chi_c > \chi_f$, it is impossible to produce a valid state that has a non-zero mismatch - in the case of no collision correction the template will make the same displacement as the atoms in the rigid section do following their modes and, if there is a collision correction, it will always trigger a mismatch correction which will always put the atoms on their ideal template position. C_{60} is therefore a special case and α-NPD will be the first chance to examine the distribution of mismatches and the effect of adding more degrees of freedom as we split a geometry-optimised molecule into groups of locally optimised sections.

[3]Reproduced from Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. Journal of Chemical Physics, 150(16):164115, 2019, with the permission of AIP Publishing

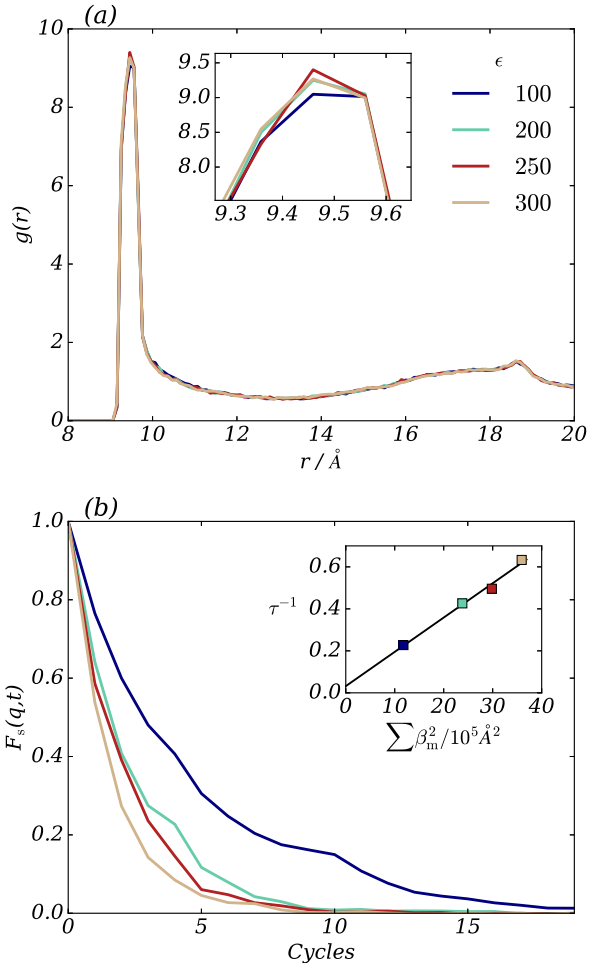

Figure 5-3: (a) the radial distribution function with an enlarged view of nearest neighbour peak in the inset and (b) the self-intermediate scattering function with the inverse relaxation time against applied displacement in the inset for the C_{60} system under varying applied displacement for the lowest 24 modes (changing ϵ).[3]

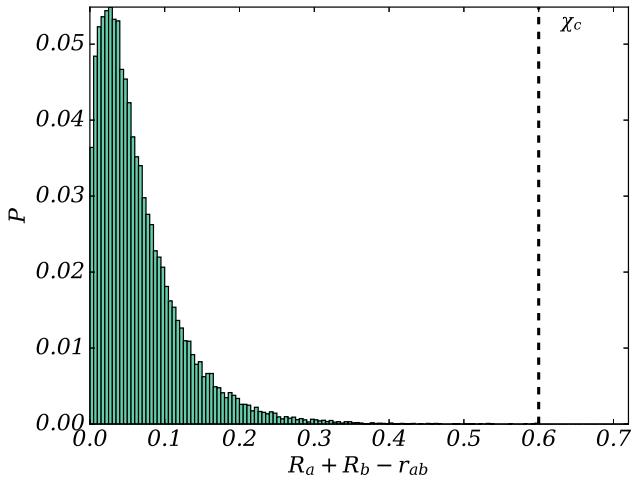

Figure 5-4: The distribution of sphere overlaps of atoms in different rigid sections for the C_{60} system following the lowest 12 frequency modes (overlaps for atoms between molecules). No collisions are larger than the threshold χ_c (also marked).

5.2 α-NPD

N,N'-Di(1-naphthyl)-N,N'-diphenyl-(1,1'-biphenyl)-4,4'-diamine, or α-NPD, is an organic small molecule that has found common use as a hole-transporting material in OLED devices [228, 119, 229]. α-NPD is found to form amorphous layers, with only a slight preference in orientation in thin-films [119], which means it falls into a class of materials which we expect SAMSEN will be able to simulate. This molecule also has a chemical structure which we will find convenient for examining how altering the rigidity of the molecule alters the simulated structure and the rate of relaxation.

Molecular Structure and Rigid Sections

α-NPD is a molecule built from several cyclic regions, each separated by at least one single bond as shown in figure 5-5. The two phenyl groups at the centre are connected by a single bond (this forms the biphenyl group) and on the opposing side of the phenyl rings they are connected to nitrogen atoms. These two nitrogen atoms are then connected by single bonds to another phenyl group and to a napthyl group (which is the double-ringed section). These phenyl and napthyl groups, in the α configuration, are positioned at opposite ends and on opposite sides of the molecule. This large number

[4]Reproduced from Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. Journal of Chemical Physics, 150(16):164115, 2019, with the permission of AIP Publishing

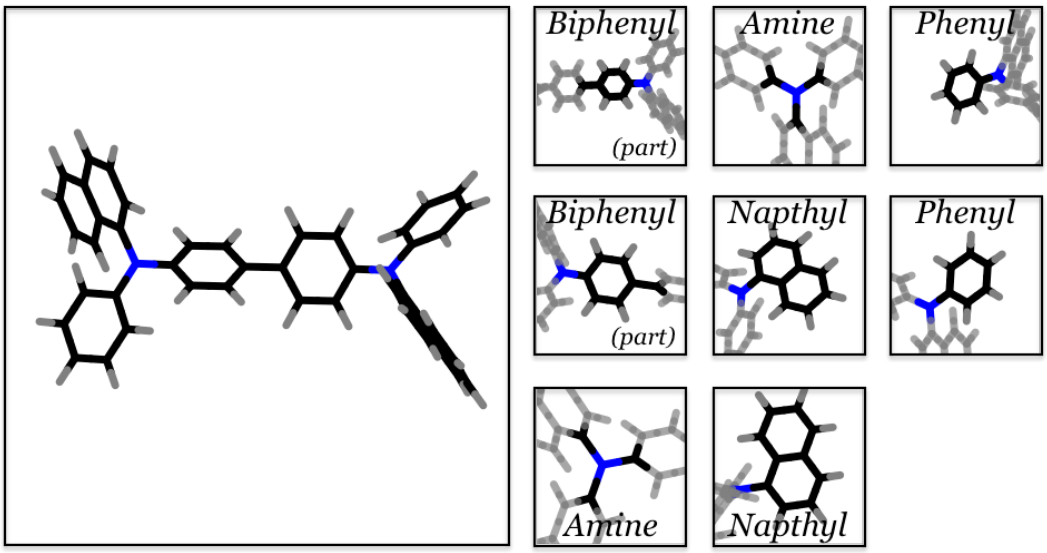

Figure 5-5: α-NPD and the 8 rigid sections generated following the method used in section 3.1.3. The molecule contains two amine-centres, two phenyl and two napthyl groups each with their own rigid section as well as a biphenyl group in the centre which forms two separate rigid sections. The measurements used in this section relate to the geometric centre of the molecule as a whole or the geometric centre for each of the rigid sections as depicted, with the exception of the biphenyl group where measurements are taken from the middle of the bond connecting the two parts (rings) of the biphenyl group.

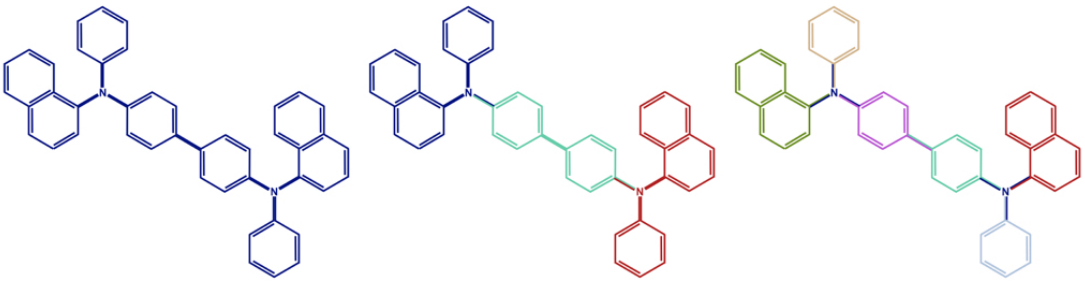

Figure 5-6: The three rigid section groupings of α-NPD used in this section to study the impact of the rigid section approximations. The colours indicate the different rigid section groupings. The left image represents a single rigid section. The middle represents the three rigid section grouping comprising the biphenyl centre and the two amine-phenyl-napthylene ends. The right groupings shows the eight rigid sections formed by the method in 3.1.3 which consists of each ringed group forming a rigid section as well as the amine-centres. Recall that the atoms bonded to each of the described sections has membership of that section e.g. nitrogen has membership of the phenyl sections.[4]

of completely rigid groups (cyclic) connected by single bonds gives us great flexibility in observing how the structural and dynamical models respond to altering the rigid sections, in particular, by merging different sections.

In figure 5-6 is α-NPD with three different rigid section groupings which we shall use to investigate the method for determining rigid sections in section 3.1.3. The first is α-NPD as a single rigid section. This mimics a scenario where the dihedral potentials are so large that there is great restriction of the ability of the phenyl and napthly groups to rotate or flex with respect to their neighbours. This is a full-merger of the molecule's rigid sections that would be produced by the algorithm of section 3.1.3. The middle rigid section grouping then represents a situation where the these phenyl and napthyl groups can rotate as a group with respect to the biphenyl group at the centre of the molecule. This is a merging of the phenyl, napthyl and amine-centre sections on either side of the molecule, as well as merging the two phenyl ring sections in the biphenyl group at the centre of the molecule. Finally, full rotational freedom for each of the ring sections is granted in the final rigid section breakdown, including within the biphenyl group in the centre. This final grouping is the one produced using the method of generating rigid sections used in section 3.1.3. We will run simulations on each of these three rigid section groupings and examine how they alter the local structure of this organic molecule and assess the applicability of the calculation of rigid sections used in section 3.1.3.

Simulation

A random input state of one thousand α-NPD molecules was generated for each of the rigid section groupings (each with their own set of MVEEs). Each simulation began at a density of 0.12 g/cm^3 in a cubic volume with periodic boundary conditions and was compressed at $C_f = 0.99995$ to a density of 1.15 g/cm^3 which matches that of of Symalla et al. [228]. This density is in the range of reported thin-film densities which varies from 0.9 to 1.45g/cm^3 [230, 228, 231]. This took 22,818 compression steps and at each step the usual collision-fitting conditions were enforced as appropriate with $\chi_c = 0.6$ $\mathring{A}$, $\chi_f = 0.2$ $\mathring{A}$.

Once the three α-NPD simulations reached their target density, the compression was halted and the rigid sections displaced by the lowest 12 eigenmodes (m_1^{12}) at $\epsilon = 200$ K for 200 cycles with a period $S = 1,000$ steps.

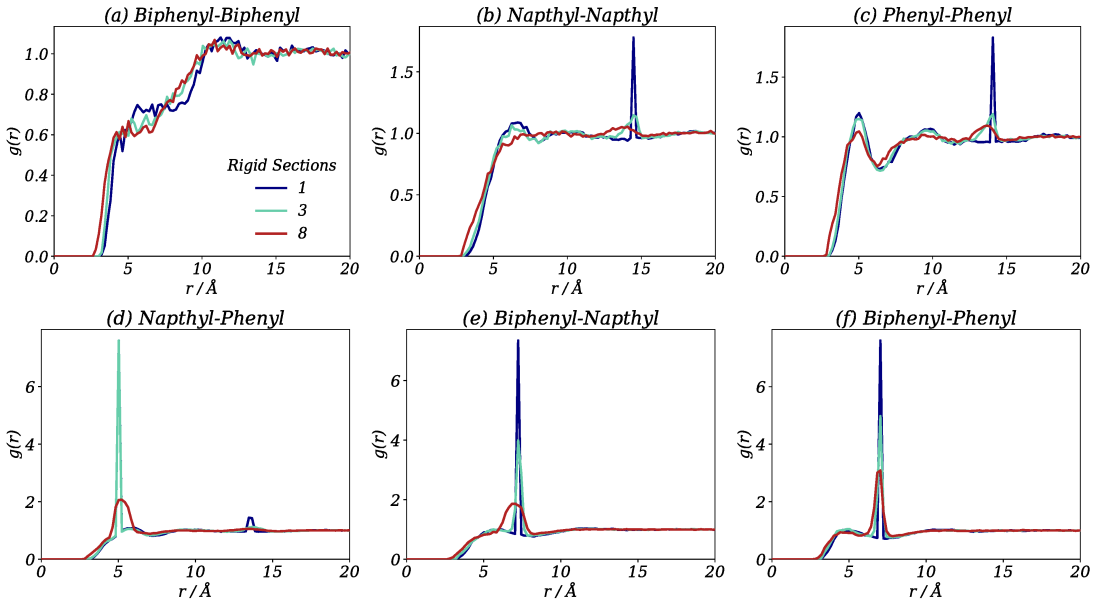

Figure 5-7: The radial distribution function ($g(r)$) for the three rigid section groupings of α-NPD as measured between the different functional groups within the molecule. The sharp peaks for the fewer rigid section simulations correspond to internal separations (neighbouring points within the same molecule). These broaden as extra rotational freedom is allowed within the molecule as the number of rigid sections increases. The $g(r)$ is measured between the geometric centre of (a) the biphenyl groups of each molecule, (b) between all napthyl groups in the system (two per molecule), (c) between all phenyl groups (not including those in the biphenyl groups) in the system and (d,e,f) between combinations of aforementioned groups.[5]

5.2.1 Structures of α-NPD

For each of the simulations we aim to compare their structures through the radial distribution function as measured between the different functional groups in the molecule. To do this we compute the $g(r)$ in the normal way, but instead take a point of one type and only count contributions to the $g(r)$ at a given r for points of a second type —in this case other functional groups in the molecule. The normalisation factor is altered to account for the number of points in the modfiied sample. The radial distribution function can be seen in figure 5-7, averaged over 20 states separated by 10 cycles for 6 different pairings built from the sections described in figure 5-5: (a) between the centre of the biphenyl groups, (b) between the centre of the napthyl groups, (c) between the centre of the phenyl groups (not including biphenyl), and (d-f) the three mixed combinations, comparing the positions of one type to positions of other types.

From each of the radial distribution functions, it is clear that SAMSEN produces a dense fluid-like $g(r)$ again without any obvious long-range order and only with sharp features when there is a neighbour to sample within the same molecule (these particularly sharp peaks correspond to the separation in the optimised molecule structure). What can be seen, however, is that when additional rotational freedom and angular freedom is granted, the structures diminish and tend towards a featureless $g(r)$.

The simulations each produce similar short-range structure. In the case of biphenyl-to-biphenyl (representing the α-NPD centre-to-centre $g(r)$), each of the rigid section groups produce a minor peak between 5 and 6 $\mathring{A}$ and another larger peak at 11 $\mathring{A}$. In coarse-grained MD these peaks should reach an amplitude of 0.8 and 1.2 respectively [232]. These slightly muted peak amplitudes, albeit at the correct radial positions, followed by poorly pronounced minima appears to be a common feature of SAMSEN simulations - the distribution of separations is broader than would be expected. In coarse-grained MD, one would expect a reduced structure when the representation is coarsened however, for SAMSEN, the opposite is true and the coarser simulation (the fully rigid molecule) produces the best representation - perhaps indicating that the bending and rotating of the sections is significantly restricted in α-NPD.

For the internal-internal peaks, this is especially true. On the whole, these are only prominent for the 1 rigid section and 3 rigid section simulations and become significantly broader for 8 rigid sections. In the case of napthyl-to-phenyl there is little change in the leading edge, no change in the sharpest peak at 5 $\mathring{A}$ and the loss of the peak at $\approx$ 13.5 $\mathring{A}$ (the internal separation on the far side of the molecule is lost) between the 1 and 3 rigid section simulations. The 13.5 $\mathring{A}$ peak does not appear very strongly in MD simulations [232], which means the 3 rigid section simulation produces the most accurate representation of the $g(r)$, as measured from that particular pair. For the biphenyl-to-napthyl $g(r)$, a similar pattern emerges where additional sharpness (here too large a peak amplitude at $\approx 7\mathring{A}$) occurs for the 1 rigid section, but the 3 rigid section simulation best captures the $g(r)$ from the all-atom MD [232]. The same is true again for biphenyl-to-phenyl and napthyl-to-napthyl. Something similar is true for phenyl-to-phenyl, however, the far internal peak at 14 $\mathring{A}$ is too tall and narrow in the 3 rigid section case. Across all of the simulations the 3 rigid section appears to match the all-atom MD closest, while the 8 rigid section perhaps most closely matches a CGMD simulation [232] - perhaps with the exception of figure 5-7(a) as discussed.

[5]Reproduced from Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. Journal of Chemical Physics, 150(16):164115, 2019, with the permission of AIP Publishing

I lay this out in such detail because the alteration in peak heights can be explained by the additional rotational freedom granted as the number of rigid sections are increased and the overall rigidity of the molecule lowered. This informs us, in an indirect way, about which parts of the physical molecule are relatively rigid (or can be considered as such) and allows us to evaluate the scheme for constructing rigid sections used here. From the additional napthyl-to-phenyl 13.5 $\mathring{A}$ peak and the enhanced sharpness of the first internal peak in all but panel (a), it is clear that complete rigidity (within a distance χ_f, at least) is a step too far and the molecules are too structurally cold and low in energy (internally). That the 3 rigid section simulation best captures the structure as measured between the biphenyl-to-napthyl groups means that the napthyl group, which has gained partial rotational freedom from the biphenyl rigid section as part of the amine-napthyl-phenyl rigid section, has a similar distribution of bond angles (napthyl-nitrogen-biphenyl) to this in other simulations (with freedom on the atomic level, granted by an all-atom MD). This would suggest that additional bending created by not including the napthyl group in the amine-centred rigid section (in the case of the 8 rigid section simulation) does not capture the real internal potentials. Meanwhile, for the phenyl-to-phenyl 14 $\mathring{A}$ internal peak, additional freedom is required to diminish the peak amplitude and the reverse —including it in the amine-centred rigid section—doesn't accurately represent the internal potentials. This might suggest that taking the 8 rigid section groupings, merging the amine-centre with the napthyl group on both sides and treating α-NPD as 6 rigid sections might be appropriate the recreate the 3 section $g(r)$ in figure 5-7(a,b,d,e), but broaden the far internal peaks in (c,f) and produce the best representation of potentials and therefore internal and short-range structure. Altering the rigidity of the molecule will, however, have consequences for the sampling of amorphous states by altering the number of nodes in the elastic network as we shall soon see.

Distribution of overlaps and mismatches

As well as altering the separation between the different functional groups in α-NPD, changing the number of rigid sections also alters the distribution of overlap and mismatch lengths. As discussed, this represents the first opportunity to examine this effect. Calculating the overlap and mismatch lengths between atoms not in the same rigid section and against all assigned template positions respectively, figure 5-8 shows (a) the distribution of collisions/overlaps and (b) the distribution of mismatches for each of the three α-NPD groupings, sampled over states produced at the end of the final three cycles using equal bin widths and normalised so that the sum of all points

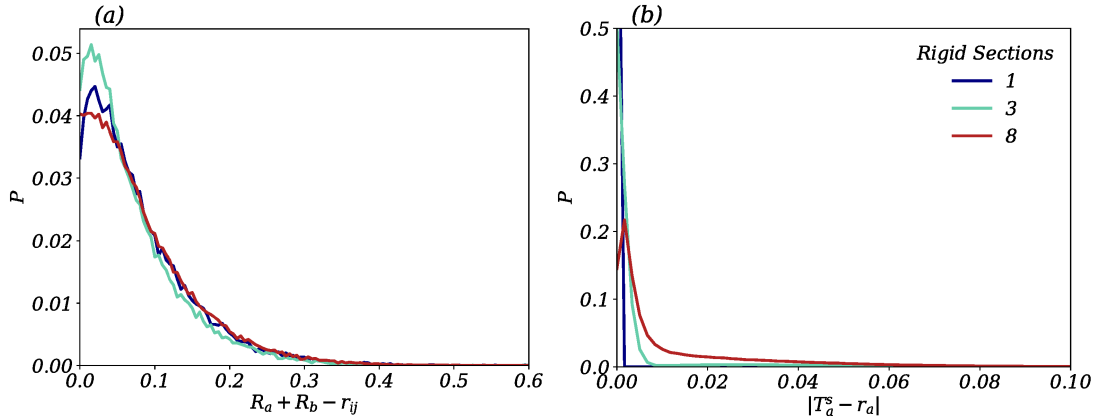

Figure 5-8: The distribution of collisions (a) and the distribution of mismatches (b) for the three α-NPD rigid section groupings. The average number of collisions was 16,175.666 for the 1 rigid section, 13,163.333 for the three rigid section and 20,852.0 for the 8 rigid section simulation.

is equal to 1.

In the distribution of collisions we find a similar shape to that in the unit-sphere and C_{60} systems - a broad distribution with a low mean and modal values with very few overlaps near χ_c. It is clear that the chance of finding an atom at high overlap increases for the 8 rigid section compared to the 3 rigid sections grouping. However, we find a similar distribution at higher overlap for both the 1 and 8 rigid section simulations, with some differences as $R_a + R_b - r_{ij} \to 0$ (the 1 rigid section reduces in probability here, the 8 rigid section reaches its peak). It would appear that the 1 rigid section simulation moves to higher overlaps (mean length of 0.082 $\mathring{A}$ with 16,175.666 collisions per state) with respect to the 3 rigid section simulation (mean length of 0.073 $\mathring{A}$ with 13,163.333 per state) while the 8 rigid section simulation (mean length of 0.085 $\mathring{A}$ with 20,852.0 collisions per state) broadens the distribution to higher and lower overlaps. Looking at the peak heights, they appear to shift to lower overlaps as the number of rigid sections is increased.

The mismatch distribution shows simpler behaviour. For the 1 rigid section simulation all mismatches are zero just as they were for C_{60} due to the order in which the collision-fitting corrections are applied and $\chi_c > \chi_f$. For the 3 rigid section simulation, the distribution broadens away from an mismatch of $|r_a - T_a^s| \approx 0$ (which now only accounts for 50% of atoms) and we find almost all atoms have a mismatch of less than 0.01 $\mathring{A}$ with an almost flat region up to χ_f. For 8 rigid sections we see significant change. $|r_a - T_a^s| \approx 0$ is no longer the most likely position to find atoms in and the distribution

significant broadens with a significant number of atoms having mismatches up to 0.10 $\mathring{A}$ and a near flat, but more populous, region beyond that up to χ_f.

It is perhaps unexpected that there should be fewest collisions for the 3 rigid section grouping, just as it was perhaps unexpected that this produced the most accurate structure. It should also be noted that, of the three rigid section groupings, the 3 rigid section simulation was also the fastest to compute.

5.2.2 Relaxation of α-NPD

The number of rigid sections in the coarse-graining of α-NPD alters the elastic network. It does this due to the differing structures (the relative position of the rigid section centres —the nodes in the network) but also because it increases the number of nodes and harmonic springs in the network as well as changing the size and shape of the MVEEs used to find the neighbour contacts. These effects will therefore be expected to alter the eigenvectors and the eigenvalues produced and therefore also alter the structural relaxation time.

We can evaluate this by measuring the intermediate scattering function of α-NPD, measured at the molecular centres across each of the simulations, and extracting their relaxation time from a fit of an exponential decay, as performed previously. In figure 5-9, the intermediate scattering function has been measured across a period of 30 cycles and averaged over the length of the simulation for each of the rigid section groupings. The 1 rigid section grouping (blue) shows the slowest relaxation, with a relaxation time (inferred from an exponential fit) to be 30 cycles. The next slowest is the 3 rigid section simulation which has a relaxation time of approximately 10 cycles and almost decorrelating by the end of the 30 cycle period. The 8 rigid section has the fastest relaxation time for approximately 8 cycles and, in this measurement, finds that positions entirely decorrelated over the same time period as the 3 rigid section simulation - it is as if there is some slowdown, or a second relaxation step, after the characteristic relaxation time.

It is challenging to find an explanation for why the intermediate scattering function for the 3 and 8 rigid section simulations align so closely after their characteristic relaxation time. One argument could be that the increased number of collisions for the 8 rigid section simulation show that there is clearly some action against the mode displacements which could be hindering particle motion. This may be what is occurring for the 1 rigid section simulation, where there is deviation from an exponential fit of the first

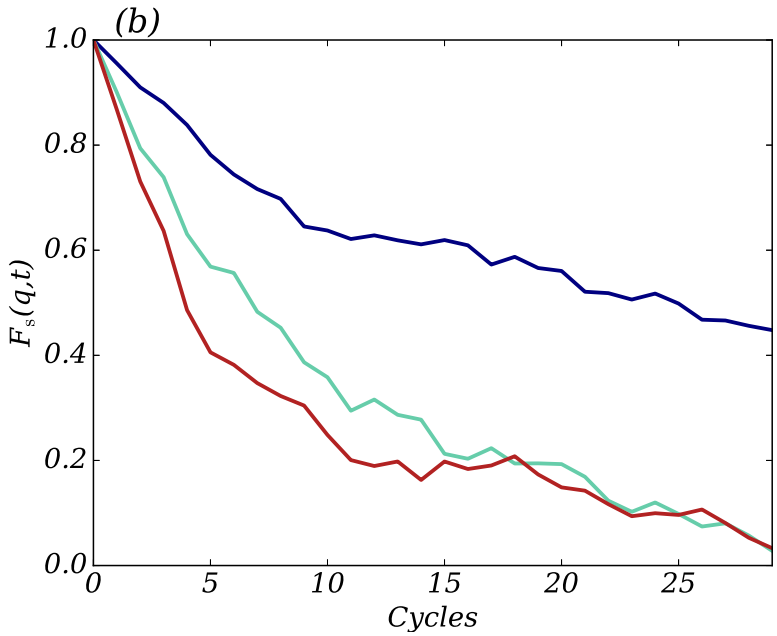

Figure 5-9: The intermediate scattering function for the α-NPD simulations for the 1 (navy), 3 (aquamarine) and 8 (red) rigid section groupings.

10 cycles at high $F_s(q,t)$ which corresponds to slowdown at shorter distances. This hindrance in the 8 rigid section simulation could also be due to enforcement of the mismatch condition if the displacements applied to the individual rigid sections are not highly collective - the opposing displacements would diminish the displacement of the molecule as a whole. However, this argument would only suggest that the relaxation is consistently slowed (beyond a distance of χ_f) system-wide, not that it is slowed after some number of cycles or intermittently or at distances similar to the molecule's diameter. Perhaps this is instead due to the behaviour of the dynamical model, rather than the structural.

Looking at the average frequencies of the vibrational modes in figure 5-10 for the different number of α-NPD rigid sections, we can begin to see some large differences in the vibrations of the elastic network. Looking at 1 rigid section, we see a picture very similar to C_{60} with distinct bands, including the splitting of the upper tan band, however the frequencies are significantly higher and over a broader range and the gaps between the bands altered. In the 3 rigid section simulation, the first 200 modes are over a much narrower range of frequencies and, correspondingly, the bands become narrower. For the 8 rigid section α-NPD, the bands move yet again to a lower and narrower range of frequencies and the band splitting becomes somewhat obscured by

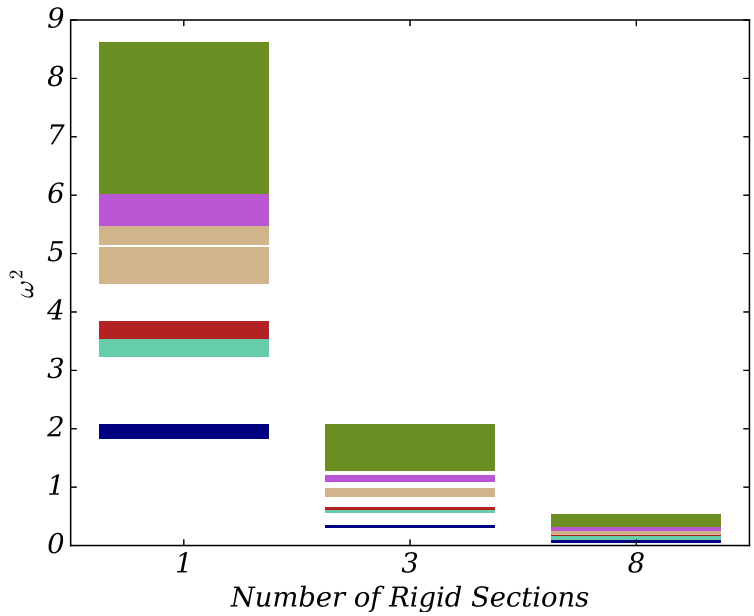

Figure 5-10: The value of the average eigenvalue of the 200 lowest frequency modes for each of the α-NPD rigid section groupings, plotted as horizontal lines with a width equal to the standard deviation.

the variation in frequencies.

The change in characteristic relaxation time can, here, be mostly explained by the change in ω and the corresponding change in β_m^2 and the therefore the applied displacement (which we have established for spheres and C_{60} has a linear relationship to relaxation time when following the collective modes). However, the blurring of the bands due to the large variation in frequencies in the 8 rigid section simulation may help us begin explain why the decay of $F_s(q,t)$ begins to slow at large number of cycles. If one studies the vibrational modes of the 8 rigid section simulation, it becomes clear that this variation is caused by the appearance of a number of soft modes which periodically disrupt the distribution of frequencies and, most notably, alters the frequency distribution of the modes selected during the SAMSEN simulation to include some very low-frequency modes. This means that, occasionally, the character of the modes used is altered significantly. Recall the reduced collective character of the unit spheres in the presence of soft modes in figure 4-7. If similar behaviour is occurring here and, periodically, the collectivity is compromised, then one may expect a relaxation at variable rates and an increase in collisions too. This may be what is occurring here and creating the late slowdown. This could perhaps be overcome by modifying the ellipsoid

inflation factor, α, to remove the presence of soft modes and performing larger averages to produce a more accurate $F_s(q,t)$.

A note on rigid sections

In this section we have explored how approximating parts of molecules as rigid alters the behaviour of the structural and dynamical models. We have seen that the method described in section 3.1.3 produces structures that are less accurate than other approximations about which parts of the molecules should be considered rigid with respect to one another and that finding a scheme, which best captures reality (inferred from measurements such as $g(r)$, is preferable to the one discussed here. For those considering following a similar approach, I would encourage them to determine the applicability of the scheme for generating rigid sections in section 3.1.3 and also consider other approaches, also discussed in that section. Again, it should be reiterated that if dihedral angles require more control than provided by the scheme for determining rigid sections of section 3.1.3, then the method may require extension to include dihedral templates, as also discussed previously. In the case of biphenyls (such as those in α-NPD), there are known minima in the dihedral angular potential representing highly likely configurations which SAMSEN will not capture and one may instead expect to see a uniform likeihood at all angles or one dominated by collisions with other molecules. Similarly, we do not include electrostatic forces and effects such as π-π will also be weaker in SAMSEN than in reality.

This, of course, is only true if the exclusion of the high energy internal configurations of the molecule is of importance. If one is trying to produce an accurate structure from which one immediately attempts a charge transport simulation, then this approach of further constraining the molecule so that the dihedral angles between rigid sections are maintained is an appropriate route. If, on the other hand, one would prefer to obtain a large sample of states and intends to use these states as starting points for molecular dynamics simulations, then restriction of the dihedral angle is of less importance and can be performed with full accuracy in a fully-atomic MD force field. The only limitation is if one is only performing a minimisation in a force-field, not an MD simulation, and the dihedral potential has multiple minima. In this scenario it may be better to constrain as much of the molecule as possible around its optimised geometry and limit disorder to disorder external to the molecule. We will take such an approach in the next section.

For the remaining work presented in this document, the method of creating rigid sec-

tions in section 3.1.3 will be used in systems where the dihedral angles should be allowed to significantly change e.g. polymers such as P3HT. However, in systems where they are more strongly bound or the molecules are sufficiently small, such as PCBM in the following section, we will assume the entire molecule is rigid and therefore constrained by the collision and fitting conditions near its optimised geometry. In situations where there is a mixture of molecules, some with highly constrained angles and others where the potential is not strong enough to prevent dihedral angles from changing, the method of determining rigid sections described in section 3.1.3 will be used for all components of the system.

5.3 PCBM

Now that the applicability of SAMSEN's structural and dynamical models to molecular systems has been examined, we can now employ SAMSEN to generate a series of states for a commonly-studied small molecule and make comparisons between the structures produced by SAMSEN and those of other methods such as molecular dynamics. Most notably, we shall also attempt to convert the series of SAMSEN output state into a molecular dynamics input state and find the local minima in a chosen force-field. This will allow us to begin to examine SAMSEN's ability to sample different parts of an energy landscape. By then using these minimised states and performing molecular dynamics under constant NPT conditions at room temperature and pressure and also performing a similar procedure under constant NVT conditions, we can then access SAMSEN's suitability as a platform for improving sampling rate of the distribution of configurational states and also consider the diffusion of charges through the material at different points on the potential energy landscape.

Phenyl-C_{61}-butyric acid methyl ester ($PC_{60}BM$, here PCBM) is perhaps the most widely-studied semiconducting small molecule and still finds use as electron transporting layer in OSCs and perovskite photovoltaic applications. It is a fullerene-derivative and is usually processed to form a crystalline structure, however under certain conditions it is possible to find an amorphous morphology [123, 233]. Its sister molecule $PC_{70}BM$ readily forms an amorphous material [123].

In the scheme of section 3.1.3, PCBM forms 8 rigid sections. However, as we have noted, we will constrain it to form a single rigid section for the following simulations. This will allow us to better capture the short-range structure (by limiting dihedral rotation for the phenyl group and side chain), at the expense of the sampling rate. Taking a

single PCBM molecule, it was minimised in LAMMPS in an OPLS-AA force-field (as published by Cheung & Troisi [24]) using conjugant-gradient minimisation. This was taken as the optimised structure from which the rigid section template and MVEE was made.

5.3.1 In Comparison to Molecular Dynamics

One thousand PCBMs (each treated as a single rigid section) were then randomly positioned and orientated in a periodic box of $L_x = L_y = L_z = 300$ $\mathring{A}$ and compressed in steps by a factor $C_f = 0.99995$, ensuring each state was valid under SAMSEN collision-fitting conditions (and applying the correction as described), to a density of 1.45 g/cm^3. This is marginally above the density of Cheung &Troisi (1.44 g/cm^3). Using the lowest 12 frequency modes (m_1^{12}) the molecules were displaced at $\epsilon = 200$ K over a period of $S = 1,000$ steps per cycle for 500 cycles. The structural relaxation time was measured by fitting an exponential decay to the self-intermediate scattering function and found to be 8 cycles with $F_s(q,t)$ decaying to zero after 30 cycles. For all sampling we therefore require a larger period to separate frames to minimise correlations between states in our sample. A sampling period of 10 cycles (excluding the first 10 cycles) will be used throughout, producing a sample size of 50 states.

The radial distribution function, as measured from centre-to-centre, fullerene-centre-to-fullerene-centre and phenyl-centre-to-phenyl-centre, is shown in figure 5-11, averaging over 50 sampled states. Figure 5-11(a), the centre-to-centre $g(r)$ appears very liquid-like with a broad first peak and a broader, weaker, second neighbour peak and shows strong resemblances to the $g(r)$ presented by Neumann et al. [120] although statistical noise makes further comparison difficult. Figure 5-11(b), the fullerene-to-fullerene takes on a form closer to the unit sphere radial distributions, with features somewhat broader as one might expect given the difference in shape of the molecule and shows strong similarity to the $g(r)$ presented by Cheung & Troisi [24], albeit with a nearest neighbour peak at shorter distances and a reduced first minima. The phenyl-to-phenyl (c), also looks similar to Cheung-Troisi, with a similar sharpness of the first peak, broad leading edge and, in the case of their 300 K simulations, retains the two additional peaks between 10 and 15 $\mathring{A}$, however loses detail in the minima at approximately 7 $\mathring{A}$ and therefore more closely resembles the 400 K Cheung-Troisi $g(r)$.

[6] Reproduced from Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. Journal of Chemical Physics, 150(16):164115, 2019, with the permission of AIP Publishing

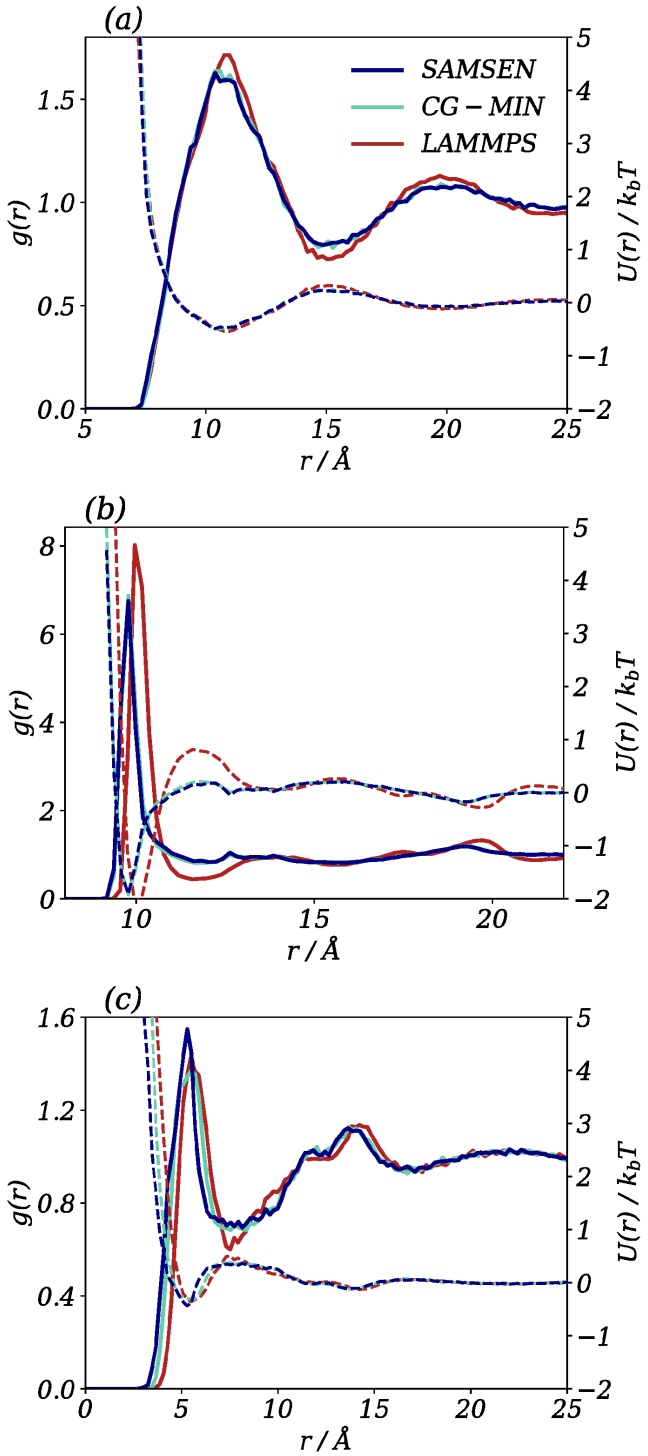

Figure 5-11: Radial distribution functions (solid) and corresponding $U(r)$ (dotted) for PCBM, measured between the (a) centre-to-centre, (b)fullerene-fullerene groups and (c) phenyl-phenyl groups for the SAMSEN outputs, the states minimised in an OPLS-AA force-field and the LAMMPS *NPT* states.[6]

Each sampled PCBM state was then minimised in LAMMPS with the same OPLS-AA force-field using conjugant-gradient minimisation (CG-MIN). The radial distribution functions, averaged over the sample, is presented in the same figure. As, it should be clear to the reader, the difference in structure between the SAMSEN and the CG-MIN states is minimal as can be seen in the centre-to-centre and the fullerene-to-fullerene $g(r)$. The similarity in structure shows that SAMSEN is able to approximate the short-range structure, such that the structures are close to those on the energy landscape produced by the OPLA-AA force-field. This statement says little about the kind of minima or the energetic distance from that minima, just that the structures produced by SAMSEN look like those in the energy landscape of a commonly-used potential in this measure. The only major difference is in (c) the phenyl-to-phenyl $g(r)$ where the first peak shifts to slightly higher separations and broadens when minimised. This difference most probably reflects the decision to make PCBM a single rigid section and that the phenyl group should be given slightly more room to bend than a single rigid section allows.

Following minimisation, each of the CG-MIN states were then used as a starting point for a molecular dynamics simulation. An NVT simulation (with a Nosé-Hoover thermostat [226, 234]) was performed over a period of 50 ps to heat the states up to 300 K. This was then followed by an NPT simulation (with an anisotropic barostat) until the pressure stabilised at 1 atm, which required around 10 ps. During this NPT simulations, the density decreased to 1.446 ± 0.002 g/cm^3. This required an average of 90 minutes per state to perform on 24 CPUs using MPI. This is in addition to SAMSEN's 6 hour average required to perform 10 cycles and produce an independent state on a single CPU (without parallelism). The $g(r)$ for each component, averaged over the sample, is also presented in figure 5-11(a-c) and is labelled 'LAMMPS'.

These LAMMPS radial distributions all show marked differences to SAMSEN and CG-MIN however they reproduce the radial distribution functions of Cheung & Troisi [24]. The most obvious differences are the increase in peak height and more pronounced minima for the three $g(r)$ compared to CG-MIN. There is also a lengthening of the nearest neighbour separation in the fullerene-to-fullerene and phenyl-to-phenyl radial distributions. Importantly, as charge transport in PCBM occurs between fullerene cages, the fullerene-to-fullerene separation was captured by performing the MD simulation from the minimised SAMSEN states. In the fullerene-to-fullerene $g(r)$ there is also an increase in nearest neighbour peak height (the 8 of LAMMPS, somewhat shallower than Cheung-Troisi which achieved just over 8) and the minima in $g(r)$ becomes pronounced. In (c), the phenyl-to-phenyl, there are only minor differences such as the slight shift to

longer separations and a sharper minima just after 7 $\mathring{A}$.

The features of the states produced by SAMSEN and then minimised and then produced by a short MD simulation, strongly resemble both Cheung & Troisi (from which the force-field was obtained) and Tummala et al. [122] in both the amplitudes and peak positions. Importantly, that fullerene-to-fullerene $g(r)$ was captured, as electron transport is predominately between these units [24]. This will give us confidence later when we measure the diffusion coefficient of these systems.

Perhaps more remarkable, however, is the similarity in the potential of mean force which is calculated as

$$U(r) = -k_b T \ln\left(g(r)\right) \tag{5.1}$$

and shown as dotted lines in figure 5-11(a-c) for the SAMSEN, CG-MIN and LAMMPS states. For the centre-to-centre $g(r)$ all three overlap almost perfectly albeit with some larger differences for the LAMMPS $U(r)$ at the 15 $\mathring{A}$ maximum (the $g(r)$ minimum). For the phenyl-to-phenyl there is a clear increase in the separation at which an effective repulsive force appears to act as one goes from SAMSEN to CG-MIN to LAMMPS. The gradient of this increase in short range potential is more or less the same however, indicating a similar broadness in $g(r)$. The same can be said for the fullerene-to-fullerene $U(r)$ too. Here the SAMSEN and CG-MIN match each other, but the LAMMPS $U(r)$ has its minima at higher separations and has a strong repulsive region around 12 $\mathring{A}$ not captured by SAMSEN and CG-MIN (like the $g(r)$). The absence of this minimum might be a result of the simplified interactions (and therefore energy landscape) expressed by the structural model (it may be that the increased penetration of the fullerene cage is reflected by a dampened effective potential at short distances), however it might also be a reflection of how the sampling of states is different between SAMSEN and LAMMPS, with MD attempting to sample in thermodynamic equilibrium and SAMSEN explicitly trying to avoid that by following the low-frequency modes.

5.3.2 Distribution of Inherent Structures

The manner in which SAMSEN samples morphologies is worth discussing in more detail. Is is clear from figure 5-11 that SAMSEN states are close enough to the minima of the energy landscape in an OPLS-AA force-field that they capture the structure of minima in that energy landscape. However, it is also clear that SAMSEN produces dif-

ferent structures to those produced after a molecular dynamics run —even one starting from a SAMSEN state. They, in large part, resemble the MD structures, but there are clear differences, even in the $g(r)$ which we have just noted.

However, it is also clear that the SAMSEN states, minimised (CG-MIN) and then evolved in molecular dynamics (LAMMPS) capture the same structure as a pure molecular dynamics simulation, by comparison to Cheung & Troisi, Neumann et al. and Tummala et al. [24, 122, 120] and others [233, 123].

This would suggest that beginning molecular dynamics simulations from the SAMSEN states and averaging without additional bias or weighting does not necessarily bias the measurement of the structures of MD states away from the equilibrium distribution. This, of course, would only be true if the distribution of states in SAMSEN was the same as the real distribution of states (sampling thermodynamically, starting from this distribution of states, one may expect to obtain the equilibrium distribution) or not adding significant bias. As had already been suggested, SAMSEN does not sample with a thermodynamic bias (as this would hinder structural relaxation), however it is unclear how exactly SAMSEN samples states and if any significant bias has been introduced.

In the same fashion as the previous simulation, another 25 sets of 1,000 PCBM molecules were placed randomly in a periodic simulation volume. The SAMSEN parameters and target density were the same as before. A SAMSEN simulation on each of the 25 starting points was performed and, following compression, and a state extracted every 10 cycles. The simulations were running simultaneously without parallelism on a 16 CPU machine and terminated after 3 weeks. 1,469 states were obtained, each separated by a number of cycles greater than than the relaxation time. These states were then minimised, as before, using conjugant-gradient minimisation in the Cheung-Troisi OPLS-AA force-field for PCBM and the energies of the states stored. A molecular dynamics simulation was then performed on the minimised states in NVT conditions. The initial velocities were set as a Gaussian representing 300 K using the LAMMPS velocity command [226]. The temperature was held at 300 K and the MD simulation progressed for 40 ps until the pressure began to stabilise. The MD states were then minimised using conjugant-gradient minimisation and the energies stored. In figure 5-12, the distribution of the potential of the SAMSEN states minimised in an OPLS-AA force-field and the minimised states following the molecular dynamics run are shown.

From figure 5-12 it can be seen that the minimised SAMSEN states are all of a much higher energy than those of the minimised states produced after performing molecular

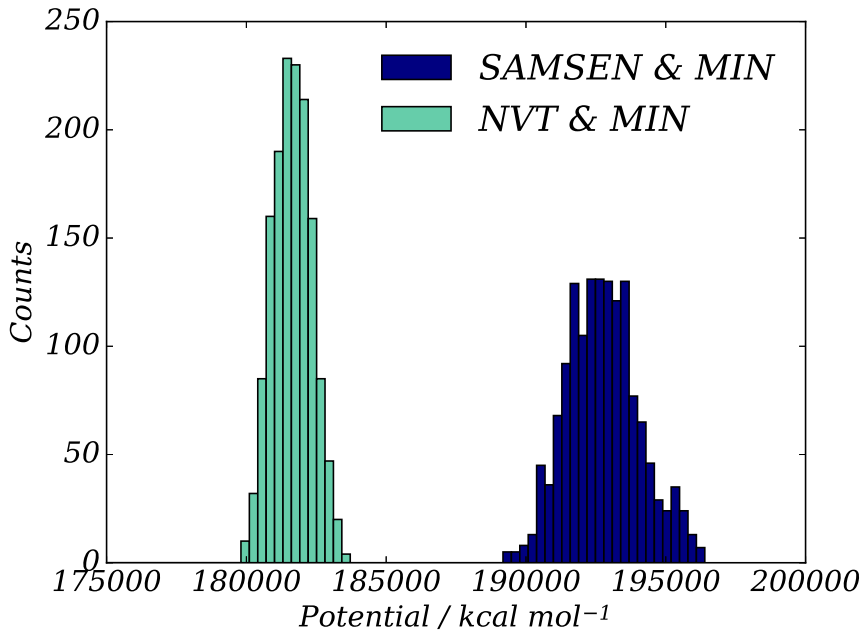

Figure 5-12: Distribution of the total OPLS-AA potential observed after performing a minimisation on the 1469 SAMSEN states and the same distribution after then running an NVT simulation on each those minimised states for 40 ps at 300 K and performing another minimisation.

dynamics at constant *NVT*. This would suggest that, in the picture of a vast energy landscape, that SAMSEN produces states which are high in the energy landscape and around very shallow minima. These shallow minima are surrounded by barriers which are low enough in energy that the systems can quickly escape their basin (potentially more than once) and find deeper minima. This picture would support the idea that SAMSEN, which assumes a smooth harmonic potential in the dynamical model and only repulsive elements in the structural model, fails to reproduce a deep attractive region followed by a second repulsive region (in the $g(r)$ or $U(r)$, for instance) with the structural model and is therefore unaware of the deep thermodynamic traps of the energy landscape. Not being tied to the energy landscape represented by the OPLS-AA force-field, SAMSEN is then able to sample many states across the landscape without being biased towards lower energy levels or spending a large amount of time in the deep minima.

The shape of the distribution of the SAMSEN minimised states is broadly Gaussian with some deviation (an increase in population) at higher energies. For glass-forming spheres in a Lennard-Jones potential, the density of inherent structures is said to be Gaussian [235, 236, 237]. This is claimed, in part, to be due to the central-limit theorem. One may therefore expect a large degree of 'Gaussianity' in the distribution of inherent structures of PCBM. Unfortunately we do not know how the addition of the phenyl group and side chain to a fullerene would modify the density of inherent structures. However, if one assumes SAMSEN does not introduce significant bias into its sampling of the distribution of inherent structures, then it would appear here that the non-Gaussian behaviour is small. Put the other way around, if the non-Gaussian behaviour really is small then SAMSEN appears to have captured the shape of the density of inherent structures, with the proviso that it has failed to capture the high-energy or low-energy tail states (this may be an inherent property of the method or simply that a much larger sample is required).

That SAMSEN does not capture the low-energy minima is an issue for direct structure prediction. One cannot simply perform a suitably-weighted average to produce a radial distribution function or an electron mobility for that material when there is not a low-energy state contributing to the average. This means SAMSEN must be limited to sampling the high energy states from which molecular dynamics will need to be performed to find thermodynamic equilibrium and sample the low-energy states. If the distribution of inherent structures produced by minimising SAMSEN states is unbiased and did represent an accurate distribution, then this would be an efficient method of creating a representative sample from which one could then produce repre-

sentative predictions of the material properties of glassy amorphous solids. However, this has not strictly been established here. All that can be said at this point is that the SAMSEN distribution is near-Gaussian and, if central-limit theorem dominates, this may approximate the PCBM distribution of inherent structures and therefore not introduce significant bias away from thermodynamic equilibrium into the distribution of MD states.

Energy and Structure

Each potential minima on this energy landscape represents a different structural configuration of the PCBM system with all atoms at their rest positions and therefore without thermal fluctuations. Among these different configurations we may expect to find different structural measurements and even charge transport properties. As discussed previously, with the full sample, or at least continuous distribution between the low-energy MD minima and SAMSEN minima, and an appropriately-weighted average we might then expect to be able to attempt to determine the bulk equilibrium properties. However this is not necessarily an option here. We are, however, able to probe the different minima and observe how their structures and other properties differ.

In figure 5-13, I have plotted the direct average radial distribution function (measured between the fullerene cages, as figure 5-11(b)) of the structures with potential minima in different ranges for the SAMSEN inherent structures. I have split the distribution into ranges of 1,000 kcal mol^{-1} (188-189, 189-190, 190-191, 192-192, 192-193, 193-194, 194-195 and 195-196 $\times 10^3$ kcal mol^{-1}), labelled as the upper limit of their range, and in figure 5-13(a) produced the average $g(r)$ in the range and in (b) produced the difference between the average $g(r)$ for the different energy ranges and the average of the 188-189 $\times 10^3$ kcal mol^{-1} range.

Across the ranges of potentials a similar radial distribution is produced, showing that the SAMSEN inherent structures do not at all capture the minima between 11 and 12 $\mathring{A}$. The nearest-neighbour peak of the lower energy potential minima is shifted to larger separations compared to the high energy minima, with fewer neighbours at the shortest ranges and a higher number past the nearest neighbour peak. The resolution on the produced $g(r)$ limits true comparison but the peak value (all in the 9.78 to 9.98 $\mathring{A}$ range) reduces with reduced energy of the potential minima. These changes are better captured in figure 5-13(b) showing the differences between the lowest energy $g(r)$ and the $g(r)$ of the range of potential minima. This better captures the shift in nearest neighbour peak (this time to closer separations for higher energy), its increase

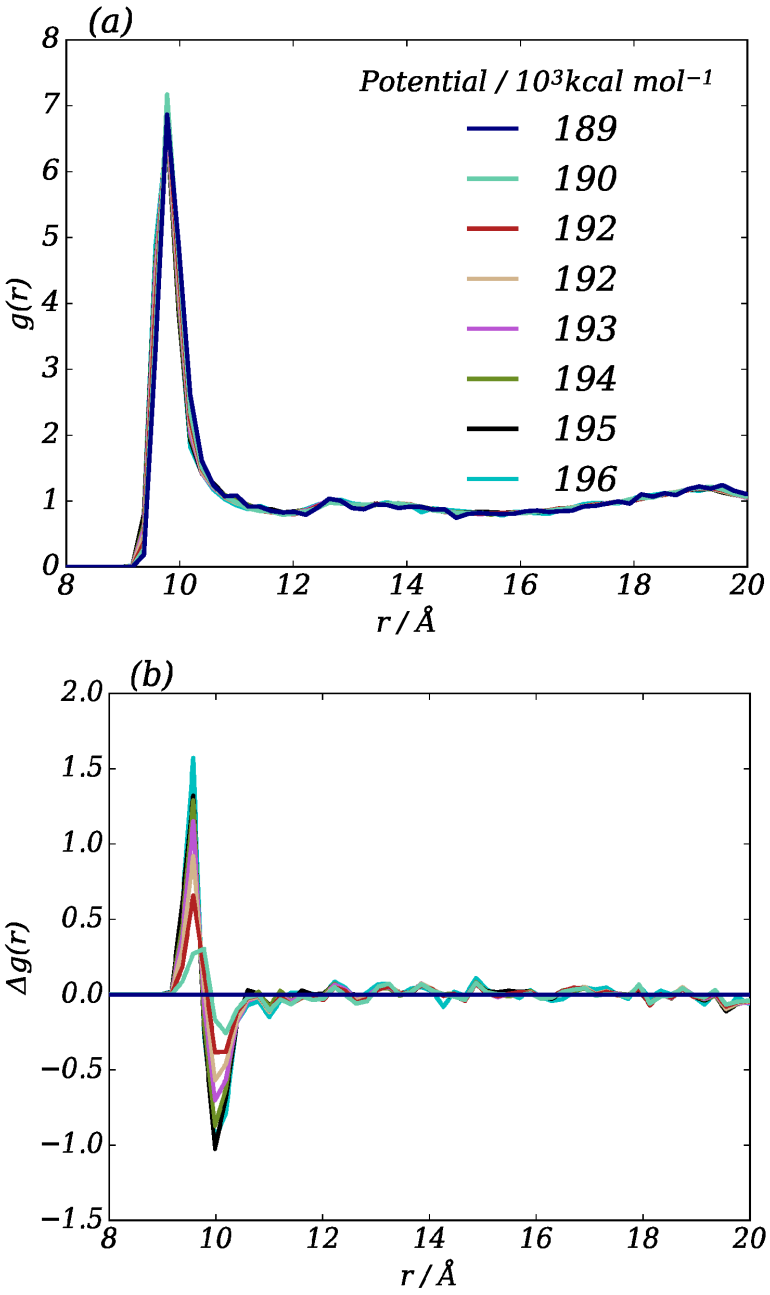

Figure 5-13: The (a) average radial distribution function for ranges of potentials of the minimised SAMSEN structures and (b) the difference between the average radial distribution function for the range of potentials and the 188,000 to 189,000 kcal mol^{-1} range

in sharpness and also the decrease in neighbours in the tail of the nearest neighbour peak. The remaining features of $g(r)$ are almost unchanging (although noise limits this comparison) with the exception of the second neighbour peak below 20 $\mathring{A}$ which shifts again to higher separations for the low energy structures. All of these changes were found to greater extents by the inherent structures produced after the brief MD simulation in figure 5-11(b) and in the results produced by Cheung & Troisi [24] - both the nearest and second nearest neighbour peak are shifted to higher separations and the minimum is also more strongly pronounced.

These subtle changes in structure correspond with the decrease in potential. Although the forces are atomic pairwise, a reduction in proximity for the fullerene cages, which contain almost three quarters of the atoms, will likely dominate the measure of potential. Assuming this, one would then expect the low energy configurations to represent larger separations of the fullerene cages.

5.3.3 Steady-State Electron Transport

Observing how the structure correlates with the potential of the minima is pertinent to studies of organic semiconductors as the electron transport of PCBM occurs predominately through the fullerene cages. In figure 5-14 the electronic wavefunction for an electron in the lowest unoccupied molecular orbital (LUMO) with the lower orbitals occupied is shown. This wavefunction was calculated at the DFT level in CP2K by Dibyajyoti Ghosh formerly at the University of Bath, now Los Alamos National Laboratory. It shows clearly the delocalisation of the wavefunction across the fullerene cage and an absence of the wavefunction in the phenyl group and side-chain. As the PCBM configurations across the distribution have variable separation of fullerene cages and the states with the most separated fullerene cages are the lowest energy states, then one might expect a change in charge mobility or even the diffusion constant as inherent structures of different potentials are sampled.

Charge hopping is significantly dependent on the separation between localised sites - in the Miller-Abrahams hopping expression it is exponentially decaying with distance, in Marcus theory it depends on the coupling of electronic orbitals which has been shown to also have an exponential decay with separations [24, 238] (and the dependence of other properties such as reorganisation energy have also been correlated with distance [40]). Although Cheung-Troisi continue to explore the delocalisation of charges, finding a localisation length of around 50 $\mathring{A}$ and the electron delocalised over distances that are comparable to their simulation cell length, we will continue to assume charges hop

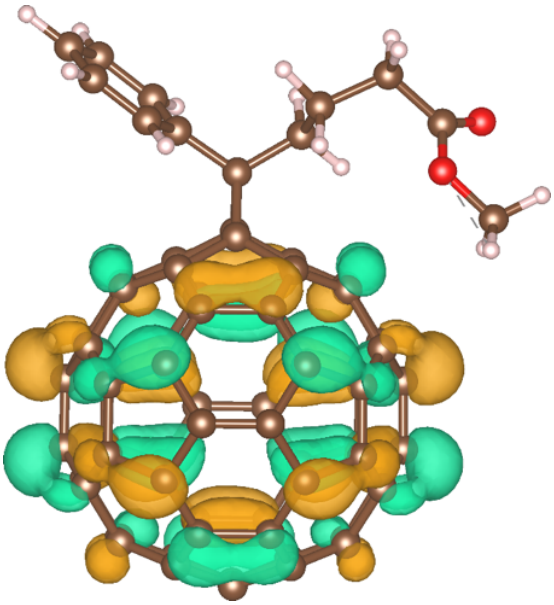

Figure 5-14: The electronic wavefunction of an electron in the lowest unoccupied molecular orbital of PCBM

between the centre of mass of the fullerene cages. Other researchers have simulated transport this way using a Miller-Abrahams hopping rate with an inverse localisation of $\alpha = 2$ nm^{-1} [74, 239] and, without access to facilities that would enable calculations of the transfer integrals for each of the configurations we are unable to use a Marcus or more quantum mechanical approaches. However, one may still expect to draw similar conclusions to those presented here for systems where charge transport is well defined by hopping between localised sites and where low energy configurations are those with larger separations between hopping sites.

Using the transfer matrix method described in section 3.4.2, each of the 1469 inherent structures of both the SAMSEN and NVT states introduced in the previous section were used as the morphology upon which a charge transport calculation was performed. In each state, the centre of mass of each of the fullerene sections of PCBM were obtained and used as the coordinates of the hopping sites. A list of all neighbours within a cut-off radius of 20 $\mathring{A}$ was obtained. The fullerene site energies were picked from a Gaussian distribution with width $\sigma = 0.1$ eV (the value obtained from Athanasopoulos et al. and Kimber [74, 239]). As we are only considering single-species transport —here electrons—the absolute energy level is of little consequence. For the Miller-Abrahams rate expression $\alpha = 2$ nm^{-1} and $v_0 = 1$ ps [74, 239] and the temperature was set to 300 K. Between all neighbouring sites (within cut-off) the rates were calculated for the given energy differences between sites and the probabilities at $t = 10^{-15}$ seconds entered into

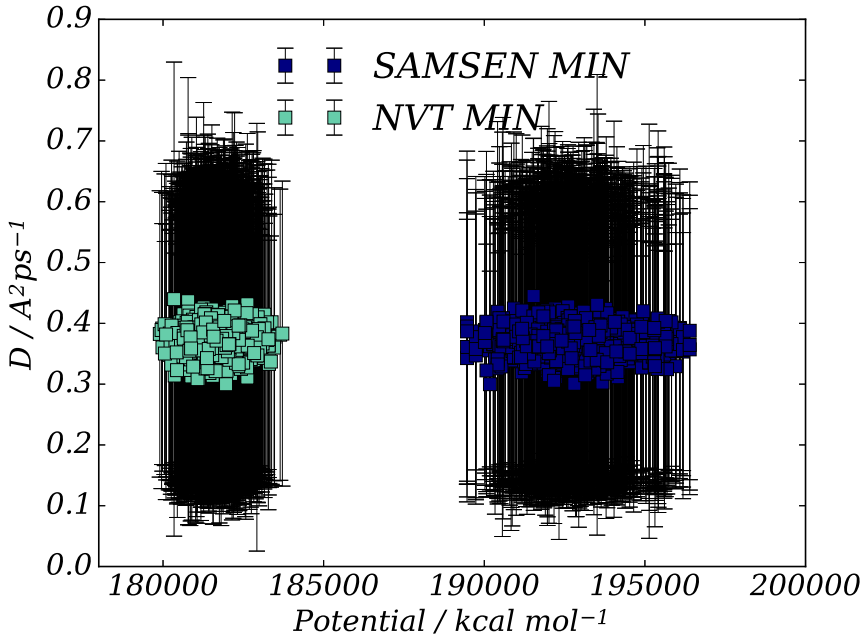

Figure 5-15: The diffusion coefficient for the inherent structures produced by conjugant gradient minimisation of the sample of states produced by SAMSEN and the conjugant gradient minimisation of those states after performing an NVT molecular dynamics simulation at 300 K.

the transfer matrix. The matrix was diagonalised and the eigenvalue corresponding to the steady state found. In the event of multiple eigenvalues represented numerically as 1, the last eigenvalue in the range was used. The corresponding eigenvector was then tested for steady-state conditions by performing a matrix multiplication with the transfer matrix. If no change was observed, this was then considered the steady-state occupation probabilities of the given state. The mean squared displacement in time, t, was then determined from the probability of hopping to a neighbour at square distance, r^2, weighted by the probability of occupancy, which leads directly to the diffusion constant, D. This process was repeated 200 times picking a new set of energies from a Gaussian distribution each time to obtain an average for the state. This was performed for all SAMSEN and NVT inherent structures and the diffusion constant is plotted against the potential of the minima in figure 5-15. The uncertainty is represented as the standard deviation of the diffusion constant for the 200 energy distributions of that state.

In both samples of inherent structures, there is a narrow distribution of the average diffusion coefficients for each morphology, themselves representing a much broader dis-

tribution of calculated diffusion coefficients for each distribution of site energies. For the SAMSEN distribution there is a very small negative skew at -8.38×10^{-7} $\mathring{A}^2$ ps^{-1}/kcal mol^{-1} (significantly smaller than the associated uncertainty in a linear fit) and for the NVT simulation there is a very small positive skew at 5.07×10^{-7} $\mathring{A}^2$ ps^{-1}/kcal mol^{-1} (again smaller than the associated uncertainty). Over the entire range of energies, combining both distributions, there is even less of a dependence of the diffusion constant on potential of the minimised configuration at -1.53×10^{-7} $\mathring{A}^2$ ps^{-1}/kcal mol^{-1}. This would suggest a 0.0023 $\mathring{A}^2$ ps^{-1} change across the entire range of detected minima in this energy landscape.

This lack of change can also be seen in figure 5-16 where the distribution of the average diffusion coefficient across all states minimised after SAMSEN and after NVT is shown, across the potential energy range there is little change to the measured diffusion. Perhaps there is a change in the width of the distribution, with the minimised NVT states producing a slightly narrower distribution of D than the minimised SAMSEN states suggests however it, for the most part, any change appears negligible. From this analysis, it would appear that the diffusion coefficient is a constant, regardless of the potential energy of the minimised structure and therefore, in a localised hopping transport model, the material produced in non-equilibirum conditions produces similar properties to that producing an equilibrium state.

In the SAMSEN minimised states there are several data points at the same energy, indicating that a longer relaxation period was required. The number of unique energy basins sampled was 367. However, the NVT simulation contains no duplicates. This would suggest that the barriers to escape the SAMSEN minima are low in energy and many in number - which would reflect the idea of the SAMSEN states being high energy minima but perhaps suggest that the inherent structures sampled are not potential minima within much larger 'superbasins'. As discussed by Goedecker [142], methods without 'a history' (without discouraging the revisiting of previously sampled states) are prone to repeat sampling. It is possible that the average relaxation time (inferred from $F_s(q,t)$) was a poor representation for that configuration (a longer time was required) or that, after relaxing, the modes displaced the system back to a previous basin. These additional states should be removed from any future weighted-average of properties of the structure to prevent unintended bias.

It is perhaps remarkable that, as the structure is changing (the neighbours moving to higher separations as energy is decreased) that the charge diffusion coefficient is seemingly unaffected, uncorrelated and independent of the energy of the inherent structure. This may be a reflection of the large delocalisation and the small inverse localisation

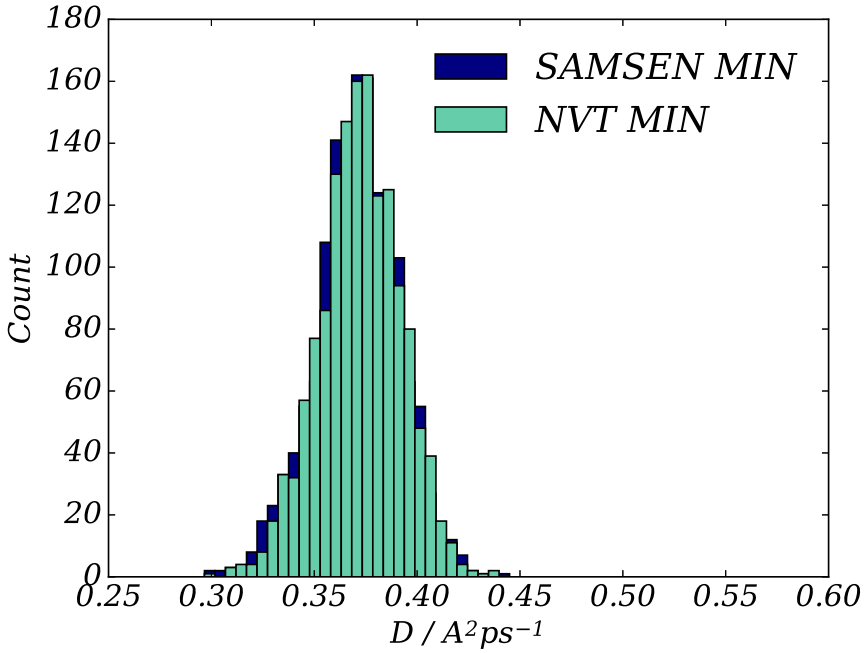

Figure 5-16: Distributions of the average diffusion coefficient for the minimised SAMSEN and SAMSEN-NVT states.

length in the Miller-Abrahams rates, which would limit the impact of small changes in separation. A different story might be told if delocalisation or a semi-band transport model applied. It could, however, be a reflection of the small system size. The wide distribution of diffusion coefficients for each state is perhaps dominated by this - the distribution in energies and the correlation in site energies has a much larger effect on smaller systems [43, 45] and overestimates of the diffusion constant also occur in such circumstances [45]. We will consider finite system size in more detail in chapter 7.

5.4 Summary

In this chapter, SAMSEN was applied to three small molecules in order to assess the structural model's ability to produce structures in comparison to those of other simulation methods. From the C_{60} simulations we saw that SAMSEN continues to produce amorphous structures and that the displacements of the dynamical model which follow the low-frequency modes are highly efficient pathways to achieve structural relaxation and fast sampling rates of states in the energy landscape of amorphous molecular systems. We later saw that for PCBM, the states were generally high in energy and close

to the structures of the high energy minima and not those of the equilibrium states which accounts for the differences in structural measurements between MD states and SAMSEN states.

α-NPD was also used to investigate the rigid section approximations, comparing the rigid section groupings produced by the algorithm of section 3.1.3 to both coarse-grained and atomistic MD simulations of α-NPD. This is solely an assessment of how well the structures are approximated by the assumed rigidity of the dihedral bonds between the cyclic regions. It would appear that the method of section 3.1.3 under estimates the rigidity of the dihedral forces and that merging some rigid sections (restricting the motion of the napthyl groups) may produce the best structural representation of α-NPD morphologies. It may also be wise, going forward, to introduce dihedral constraints between rigid sections (although that will not be performed in this work and, as we shall see in a moment with P3HT, there are instances where such constraints are inappropriate). For sampling, the fastest relaxation time was produced with the rigid section groupings created from the algorithm described in section 3.1.3 and, as we saw with PCBM, it is possible to use molecular dynamics to find the structures representative of equilibrium after a SAMSEN simulation, in which case it may be unnecessary to constrain dihedral angles at this point.

For PCBM, the samples of minimised states produced from SAMSEN morphologies as well as the minimised MD states (themselves produced from the SAMSEN morphologies) were used as the structures for performing charge transport simulations. Here the fullerene cages act as the transport centres and electrons hopped between them (although in some cases charges in reality delocalise [24]) using the transfer matrix described in section 3.4.2 with a Miller-Abrahams model using parameters from literature [74, 239]. Despite the large differences in potential energy of the morphologies and the associated structural differences (particularly that of the nearest neighbour distances), very little change was observed in the electron diffusion constant between the SAMSEN high energy minima and the MD low energy minima.

While this section shows that large sample sizes and states that span a large number of minima in the energy landscape can be produced by SAMSEN, it should also be emphasised that large system sizes can also be studied. Figure 5-17 shows such a system of 60,000 PCBM molecules in a system with dimensions 25 nm $\times$ 25 nm $\times$ 100 nm. In the one rigid section approximation of PCBM a system of thin-film thicknesses can therefore be modelled and, with enough computer memory, the low-frequency modes can still be used to sample new states. These states can then be used for a much shorter MD simulation (compared to a full MD simulation). In this way SAMSEN provides a

new method for sampling states out-of and in-equilibirum and on much larger scales than often employed in modern computational studies.

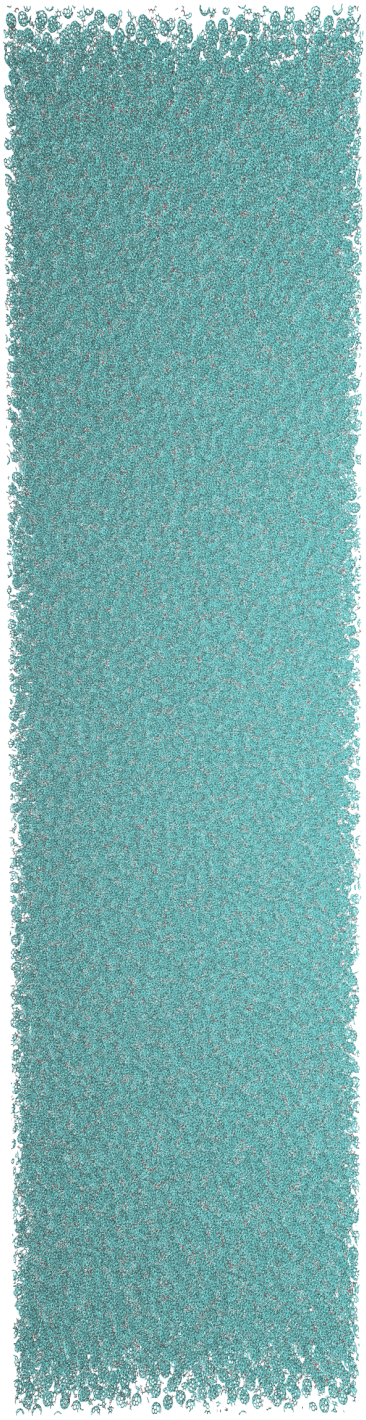

Figure 5-17: A VMD visualisation of a SAMSEN PCBM morphology with dimensions 25 nm × 25 nm × 100 nm and containing 60,000 molecules (5.28 million atoms).

Chapter 6

Structure and Vibrational Modes of Polymer-Fullerene Blends

Polymer-fullerene blends have been at the core of OPV research for many years [240, 109, 118]. The bulk hetero-junction, for instance, requires a hole and an electron transporting material to be blended and retain enough correlation in structure (large enough hole-transporting domains and interconnects) so that the charges can be successfully extracted. Blends of a hole-transporting polymer, P3HT, and electron-transporting fullerene, PCBM, perhaps, the most widely studied. [241, 242, 133, 109, 134, 19, 57, 243, 114, 240, 122, 113, 35, 118].

The addition of fullerenes and their derivatives into a host polymer system has been shown to alter the dynamics and structure of the system. Pertinently, this has been observed in P3HT, with the polymer backbone shown to be wrapping around the fullerene cage of PCBM. This has been argued to frustrate the vibrations of the polymer, reducing the relaxation time and restricting the orientation of the polymer backbone [118].

This chapter will look at both of these effects in an attempt to gain further insight into SAMSEN's behaviour as well as the structural, vibrational properties of P3HT. The following chapter will then study exciton transport in both amorphous and crystalline phases of pure P3HT.

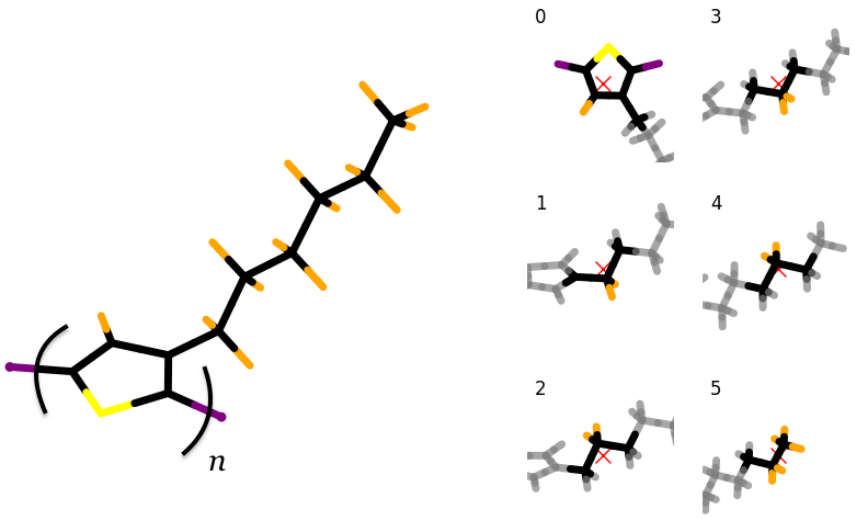

Figure 6-1: P3HT and the rigid sections produced using the algorithm in section 3.1.3 with their geometric centres marked by red crosses. The centre of the thiophene section is offset from the centre of the pentagon.

6.1 Structures

6.1.1 P3HT

Poly(3-hexylthiophene) (P3HT) is a hole-transporting semiconducting polymer used widely in OPV research. The repeat unit consists of a thiophene pentagonal-ring with a hexyl side-chain at the base of the thiophene pentagon. Polymerisation occurs between the atoms bonded to the two carbon atoms nearest the sulphur at the top of the pentagon. The bond between two repeat units is therefore not aligned with the axis between repeat units and this creates an offset between the repeat units, forcing the polymer to adopt either a twisted position or a head-tail position - where the tip of the pentagon is pointing in the opposite direction (a 180° rotation). In reality a mixture occurs and bending of the polymer chain is accompanied by a rotation away from the 180° conformation [57]. Despite the head-tail pattern repeating over two monomer units, here a repeat unit will be referred to as a single monomer.

In the scheme of section 3.1.3, P3HT forms 6 rigid sections. One contains the thiophene ring (as well as the first atom in the hexyl side-chain and also the two polymerisation

markers, indicated in purple in figure 6-1) and the remaining contain the hexyl side-chain, each with 3 carbon atoms (the first hexyl section containing the adjoining carbon in the thiophene ring). When polymerisation occurs the neighbouring thiophene rings will be connected by a single bond, with the rigid section from the neighbouring thiophene ring also containing the connecting atom. The structure of figure 6-1 will be used for the entirety of this chapter and was produced by optimising a P3HT 6-mer in a regioregular head-tail configuration, terminated with methyl groups, at the DFT level in the CP2K package (this calculation was performed by Dibyajyoti Ghosh) and then removing all but one of the central repeat units and placing markers where the adjoining atoms were. This optimised structure therefore represents the P3HT repeat unit in a polymer rather than the monomer.

6.1.2 Incorporation of PCBM

We shall begin by studying the P3HT structure for a range of PCBM concentrations, including pure P3HT. The blends used here are 0, 5, 10, 20, 30, 40, 50 wt % (percent by weight) of PCBM. Performing a range of SAMSEN simulations and analysing the structure of the morphology, we can then compare how the polymer structure changes upon addition of PCBM and progress towards studying how the dynamical model reacts to the difference in structure.

For each of the blends, a random input was generated in a periodic cubic volume of sides 300 $\mathring{A}$ (a starting density of between 0.06 and 0.12 g/cm^3). The polymers were placed first, with the next repeat unit rotated by 180° degrees about the axis connecting the polymerisation markers, and without any further rotations. This was continued to create a straight P3HT regioregular 20-mer in a head-tail configuration and repeated for each polymer. After the polymer chains were generated, the PCBM was then placed by finding random positions and orientations. Here, PCBM was treated as a molecule containing the 8 rigid sections that the scheme of section 3.1.3 produces (to be consistent with P3HT as discussed in section 5.2). If the ellipsoids of any rigid section intercepted with another MVEE outside of that molecule or, in the case of P3HT, any ellipsoids intercepted with another that were not in the same repeat unit or neighbouring repeat units (within the same polymer), then both the molecule or the entire polymer chain was deleted and a new random position and orientation was attempted.

The number of P3HT repeat units was held constant at 6,000 (300 chains of 20-mers) while the amount of PCBM was modified. As it is known that the addition of PCBM

wt %	P3HT	PCBM	Density / g/cm^{-3}	$L_x = L_y = L_z$ / $\mathring{A}$
0	6000	0	1.075	115.539
5	6000	59	1.089	117.057
10	6000	125	1.104	118.701
20	6000	280	1.134	122.426
30	6000	481	1.167	126.908
40	6000	748	1.201	132.431
50	6000	1122	1.237	139.458

Table 6.1: The number of P3HT repeat units and PCBM molecules used in the P3HT-PCBM blends. The corresponding system density and dimensions after compression are also provided.

modifies the film density [244, 118], with pure PCBM at a density of 1.45 g/cm^3 (as in the previous chapter) and pure P3HT in the amorphous phase has a density of approximately 1.1 g/cm^3, the new densities for each of the blend needed to be determined. The pure P3HT (1.075g/cm^3 [118]), the 20% mixture [118, 244] and the pure PCBM [24, 123] were known. The rest of the densities were determined by calculating the volume per molecule at known densities and determining the new total volume at the new mixing ratios [245]. Table 6.1 provides values for the number of molecules, the density and the box length for each of the simulations.

Each of the blends were compressed every step with $C_f = 0.99995$ until the specified density was reached, ensuring no conditions of a valid SAMSEN state were broken with $\chi_c = 0.6$ $\mathring{A}$ and $\chi_f = 0.3$ $\mathring{A}$. Upon reaching the target density, the system was displaced using the low-frequency modes in the manner described in chapter 3 with $\alpha = 1.2$, m_1^{24}, $S = 1,000$, $\epsilon = 200$ K for 30 cycles.

A visualisation of one of the blends is shown in figure 6-2, showing the P3HT backbone (the sulphur and the two carbons that bond to neighbouring repeat units) and the PCBM surrounding it. In this image you can observe how the bends in the P3HT chain are accompanied by rotation of the repeat units away from the low-energy configuration as has previously been observed [57].

Radial Distribution Function

Figure 6-3 shows the radial distribution function for the various P3HT-PCBM blends at their respective densities averaged over samples separated by 3 cycles and measured between the centre of the thiophene ring sections as marked in figure 6-1. This is therefore just a measure of the P3HT backbone structure in their respective blends.

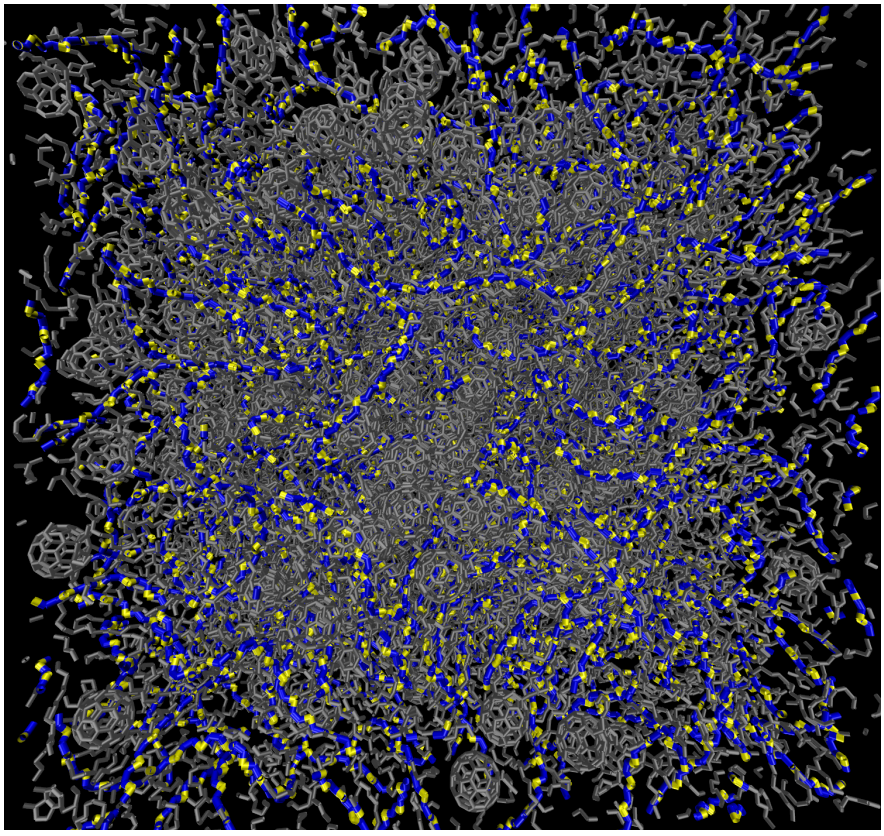

Figure 6-2: VMD generated image of a P3HT-PCBM blend only showing the PCBM (grey) and the P3HT sulphur (yellow) and neighbouring carbon (blue) atoms. Note how the twists in the polymer chain are accompanied by a rotation of the repeat unit.

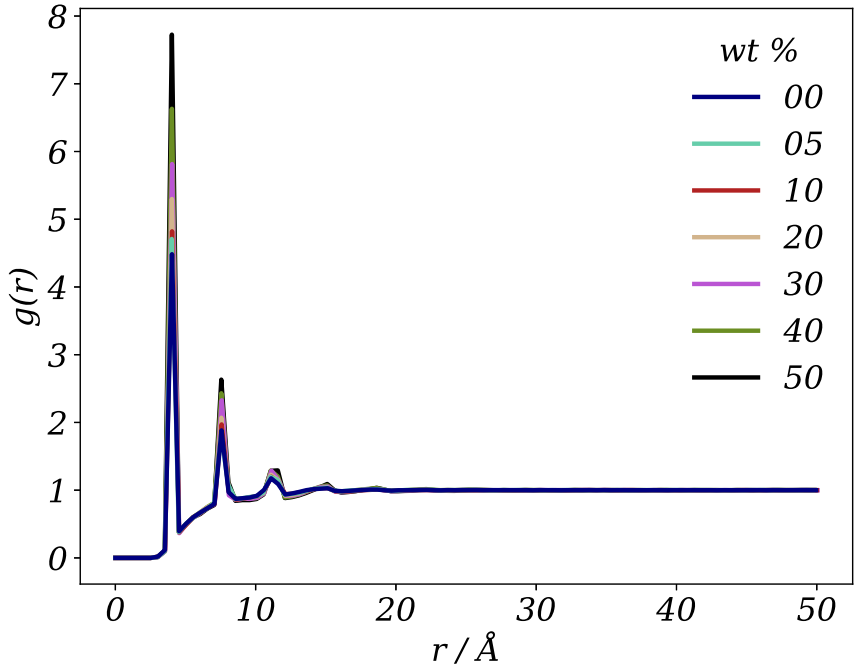

Figure 6-3: Radial distribution function, $g(r)$, measured between the P3HT-backbone units in the different P3HT-PCBM blends. The exact location from which this measurement is made is indicated by the red cross in figure 6-1 panel 0 (the geometric centre of the thiophene rigid section) and is between all P3HT backbone units in the system.

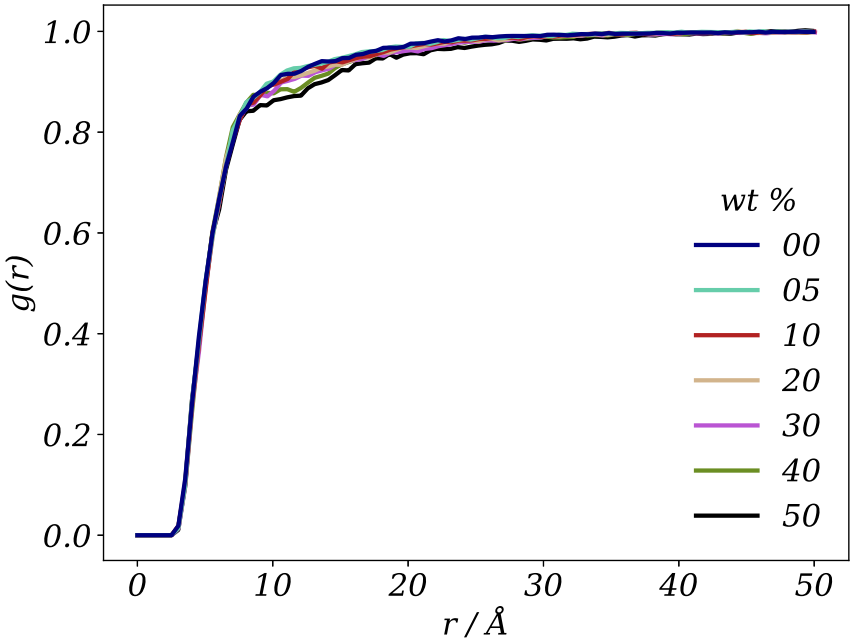

Figure 6-4: Radial distribution function, $g(r)$, between P3HT-backbone units that are not a member of the same polymer. This measure is made for each of the different P3HT-PCBM blends. The exact location from which the measurement is made is indicated by the red cross in figure 6-1 panel 0 and is between a backbone unit and all other backbone units in different polymers and normalised accordingly.

This $g(r)$ is dominated by the neighbours in the same polymer chain, with separations which become irregular and 'fade-in' to the background density after 3 (perhaps 4) repeat units down the chain. This gives a length-scale to the persistence (the characteristic distance over which the polymer bends, reducing correlation in the orientation of the monomers in the chain) of the polymer which appears to be of greater lengths in higher concentrations of PCBM. This is indicated by the sharpness of the peaks and, at low concentrations, the slight shift to closer separations (a reduced distance indicating increased curvature). In the background, there appears to be little structure with the system average density (of the subset of the system of P3HT thiophene rings) being acquired at around 12 to 15 $\mathring{A}$ separation.

In figure 6-4, I have removed the internal contributions to the $g(r)$ of figure 6-3 and re-normalised accordingly for the new sample so that at large r (for a large enough system), $g(r) = 1$. Figure 6-4 shows more clearly how the neighbouring polymer chains are positioning themselves in relation to any given thiophene rigid section centre. Unfortunately, there are few structural features. This due to there only normally being

a weak peak (6.5 $\mathring{A}$) and minimum (8 $\mathring{A}$) —at least in coarse-grained simulations [116, 113, 109] —but also because the rigid section centre is offset from the centre of the thiophene ring and the extra degrees of freedom obscure detail in this measure (this was seen in the $g(r)$ for the PCBM centre-to-centre compared to fullerene-to-fullerene in figure 5-11, for instance) and because SAMSEN tends to generate structures with slightly muted peaks and minima.

Looking at the work of To and Adams [57] as well as Jones [116] these particular radial distribution functions look closer to one of a 600 K P3HT simulation. In the work of Jones which performs a CGMD simulation and an atomistic simulation of 72 P3HT decamers under NPT conditions, the atomistic $g(r)$ measured between the thiophene rings is approximately zero until approximately 5 $\mathring{A}$ before rising and plateauing between 6 and 8 $\mathring{A}$ with a subsequent shoulder at 10 $\mathring{A}$ and a smaller one at 15 $\mathring{A}$ while the CG $g(r)$ takes a similar shape with much less pronounced features (most noticeably, the plateau is only a shoulder). To and Adams show the same measurement for a P3HT system of 12 polymers with a degree of polymerisation of 45 at 600 K in a CGMD and atomistic MD simulation. The atomistic radial distribution here has a similar shape to the CGMD simulation of Jones, while the CGMD appears a smoother version, with $g(r)$ rising from a value of 0 as early as 3 $\mathring{A}$ and only showing a plateau in the region of 12 to 17 $\mathring{A}$ (with a $g(r)$ at 15 $\mathring{A}$) of approximately 0.9). The $g(r)$ for the simulations in this work (shown in figure 6-4) show very few features for the purest P3HT blends with only the 40% and 50% blends capturing a shoulder between 8 and 12 $\mathring{A}$ with a slow rise continuing until 15 and 17 $\mathring{A}$ respectively where the $g(r)$ begins to match the featureless $g(r)$ of the purer blends. It would appear they most closely match the coarse-grained treatment by To and Adams but still reflects a much more disordered P3HT structure than both To and Adams and Jones while the purest blends appear highly disordered and do not capture the features of other work.

It would therefore again appear that SAMSEN is navigating the high portion of the energy landscape and I would suggest that, following a brief NVT simulation, the 300 K structure might again be obtained. The system may, however, struggle to achieve features associated with crystallinity or π-π stacking and be trapped in a disordered state as electrostatics were not considered during the SAMSEN simulation.

Across the range of PCBM concentrations, we find that there is little change in the external $g(r)$ except for a slight reduction in counts at all r (as PCBM now occasionally dilutes the system) and also for the 40 and 50 wt % mixtures where there is a pronounced absence of thiophene rings in the 8 to 15 $\mathring{A}$ range and, for the 30, 40 and 50 wt % mixtures, there is a reduced count for another 15 $\mathring{A}$ or so. These two ratios are

well above the miscibility limit of PCBM in P3HT where one would now expect to see crystalline aggregates of PCBM [118, 246, 247]. In this situation it may be, at least in part, a more strongly pronounced absence of thiophene neighbours contributing to the reduced structure rather than any aggregates in the SAMSEN simulation, which were not observed for any of the blends and would require $g(r) > 1$ at some $r < \frac{1}{2}L_{x,y,z}$.

We can again check numerically for aggregation of PCBM by looking at the PCBM-PCBM $g(r)$ for the fullerene centres (this measure is different to that of section 5.3 as here the measurement is made between the centre of the rigid sections which contains atoms outside the fullerene cage, so the centre is just above that of the fullerene section). In figure 6-5, I have plotted that partial $g(r)$ for the P3HT-PCBM blends above and including 5% PCBM by weight. The 5% mixture is very noisy - there are only 59 molecules in each sample so there is large variability. There does appear to be a second neighbour peak, but it takes very few occurrences to cause such a large peak. The remaining blends do not appear to show this behaviour and show a reduced peak amplitude at 10 $\mathring{A}$ as the statistics start to improve and all tend towards a value of 1 by around 30 $\mathring{A}$. Importantly, a value of $g(r) \approx 1$ at this distance (rather than $g(r) > 1$ and converging on to 1 at half the box-size), means there are no significant density fluctuations (such as clustering, phase separation or aggregation) for this type of particle (PCBM fullerenes), including for 40 and 50 wt %. This may suggest that the absence of external thiophene-thiophene counts in figure 6-4 is most probably due to the larger population of PCBM (a dilution of the P3HT) rather than any aggregation of PCBM. The lack of thiophene-neighbours may, however, also relate to the increased persistence length.

Persistence Length

The persistence length of the polymer chains can be more accurately determined than by inspecting the decay of the internal peaks of the $g(r)$. The persistence length is the correlation length of the tangent vectors along the chain. In a worm-like chain model this function can be continuous [248]. In the case of P3HT the tangent of interest is that between the thiophene rings along the backbone, best represented by the tangent at the centre of the thiophene section. With neighbouring repeat units offset from that tangent due to the angle formed between repeat units and, as we have not actively constrained the dihedral angle, the more appropriate polymer model is that of the freely-rotating chain. This is a model for polymers where neighbouring monomer units are offset at a constant angle, but are able to rotate about their shared bond, the full

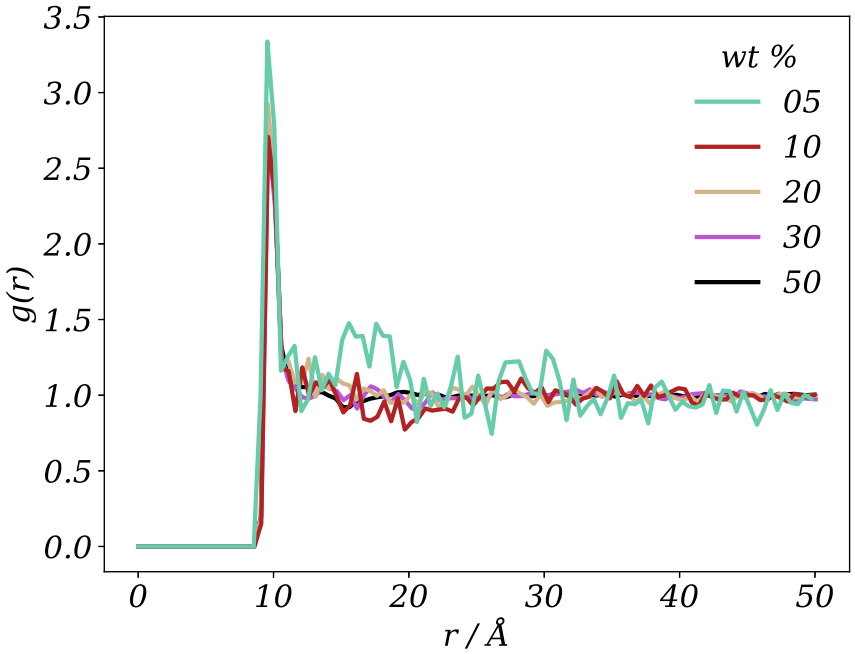

Figure 6-5: Radial distribution function, $g(r)$, measured between the PCBM fullerene rigid section centres in the different P3HT-PCBM blends.

path of that rotation drawing out a cone [249]. In the case of P3HT that rotation can occur (although it requires a bend in the chain) and the cone is drawn out by the centre of the thiophene ring.

The persistence length, L_p, in this model is given by

$$\langle \cos(\theta_{ij}) \rangle = \exp\left(\frac{-l(j-i)}{L_p}\right) \tag{6.1}$$

where l is the monomer-monomer repeat length and θ the angle between tangent vectors at monomers i and j. For neighbouring monomers in the freely-rotating chain this should be the angle of the cone drawn out by the neighbouring monomer. The persistence of the P3HT backbone, measured between the centre of the thiophene rigid sections and excluding the first and last thiophene section in each polymer chain, is shown in figure 6-6.

For each of the blends, the decay in correlation of the tangent vectors begins by following the expected exponential decay. However, approaching half the chain length ($j-i \approx 8$), the correlation is increased in each of the blends, before then falling and decaying once more. The chains here are short and the correlation doesn't completely diminish, but

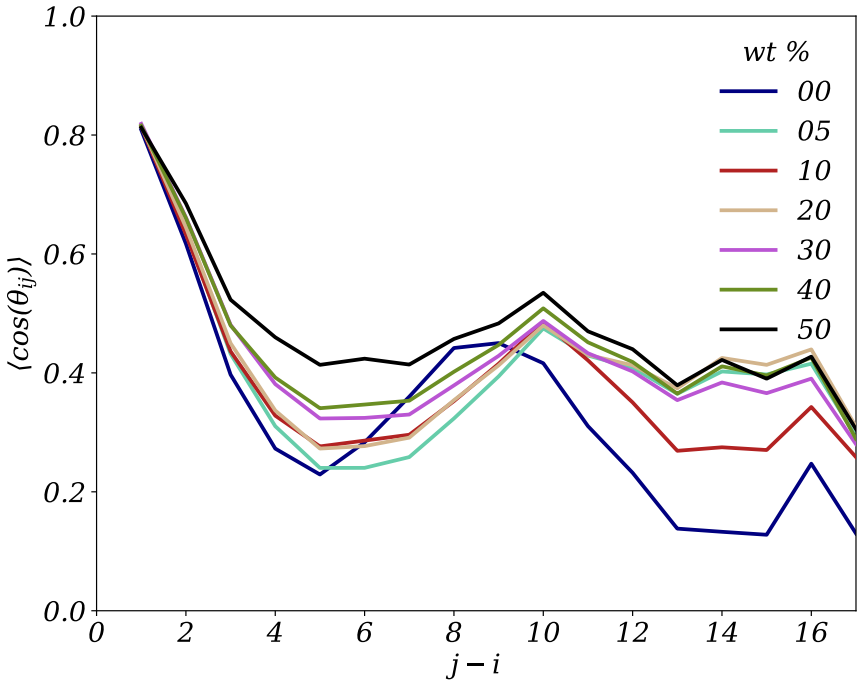

Figure 6-6: Tangent correlation function assuming a free-rotating chain model polymer for the P3HT backbone (thiophene rigid sections) in the different P3HT-PCBM blends.

wt %	Persistence / monomers	Persistence / $\mathring{A}$
0	3.493	13.971
5	3.745	14.978
10	3.879	15.516
20	3.958	15.832
30	4.366	17.464
40	4.445	17.781
50	5.164	20.654

Table 6.2: The persistence lengths of the P3HT backbone in each of the P3HT-PCBM blends, measured by an exponential decay to the early portion of the tangent correlation function.

a significant final reduction in correlation occurs after distances of 12 to 17 monomer units. The angle made between the tangent of a monomer and the tangent of the nearest neighbouring monomer unit is, on average, 35.9° in all of the simulations.

An increase in correlation far down the chain was an unexpected result. It is possible that the decorrelation at short distances is a result of the lack of dihedral potential and the structural model is very quickly able to locally rotate the repeat units which, as discussed previously, forces the chain to bend. However, it would appear that while there is a large degree of bending at short distances (perhaps due to the lack of dihedral constraints), the chains themselves are not yet disordered at intermediate distances. The cause of this is unknown: it could be that a sufficient number of cycles following compression have not passed and chains still partially resemble the initial straight chains. This is unlikely to be the sole cause given the the large differences in tangent correlation between the blend ratios even at similar box size and, as we shall see in a moment, similar relaxation times which were very short. It could also be a consistent curling of the polymer chains which have been occasionally observed qualitatively in visualisations of the morphologies.

It also appears that systems with higher blending ratios of PCBM maintain more directionally correlated chains and increased persistence even at short distances. The persistence length of the fast decay (fitting an exponential decay of the form in equation 6.1 to the first 5 neighbouring monomers) for each of the blends is shown in table 6.2 with the persistence length displayed in terms of the number of monomers and the distance along the chain using the first peak of the internal contributions to the $g(r)$ (≈ 4 $\mathring{A}$).

The change in persistence length as the PCBM ratio is increased is in keeping with the reduction in neighbours observed in the external $g(r)$ measured between between the

thiophene rigid section centres in different chains in figure 6-4. The largest reduction in neighbours (in the 50 wt % $g(r)$) coincides with a significant increase in chain persistence. As the packing efficiency of straighter objects is significantly reduced [250, 251, 252] one would expect a reduction in P3HT neighbours and either a compensatory increase in local PCBM concentration or voids to form. However, as we have picked the system density based upon the volume fractions of each component, I would suggest that vacant regions of the simulation box are unlikely - they were not observed, despite great efforts to measure them.

Wrapping

Guilbert et al. [118] observed that the P3HT thiophene rings tend to wrap around the PCBM dopants such that the face of the thiophene pentagon lies planar to one of the faces of the PCBM fullerene cage. They go on to argue that this constrained position around the fullerene leads to the mechanical frustration of the P3HT backbone which they observe as a reduction in relaxation time through quasi-elastic neutron scattering (QENS) measurements.

In the system of P3HT-PCBM blends, we see an increase in PCBM proximity to P3HT chains as the PCBM concentration is increased, however, using the order parameter of Guilbert et al. (a similar one is also used by To and Adams [57]) we find no increase in co-faciality at different mixing ratios. We find that the systems have slightly more orientational order at short distances than a hot P3HT system of To and Adams [57] but less orientational order at longer distances than Guilbert et al. [118]. At short distances the 360K results of Guilbert et al. are captured well, but it appears SAMSEN's structural model cannot maintain orientational correlation at larger distances and misses the additional alignment present at $12\mathring{A}$. This can be seen figure 6-7 which shows $\langle 3\cos^2(\theta) - 1\rangle$ as a function of r (a direct comparison to figure 4a of Guilbert et al. [118]). Here θ is the angle formed between two vectors, the first being the vector connecting the centre of the fullerene cage of a PCBM molecule and the second being the normal vector to the plane of the thiophene ring in a P3HT repeat unit and $\langle 3cos^2(\theta) - 1\rangle$ is an orientational order parameter which produces a value of 0 if there is no orientational preference, 2 if a flat side of the PCBM fullerene cage and the thiophene ring and co-facial and -1 if they are perpenducilar to each other.

The SAMSEN P3HT structures are more disordered than those of their counterparts, showing weaker peaks in structural measurements such as $g(r)$ and appearing a much hotter system. It would seem that the most blended systems have straighter P3HT

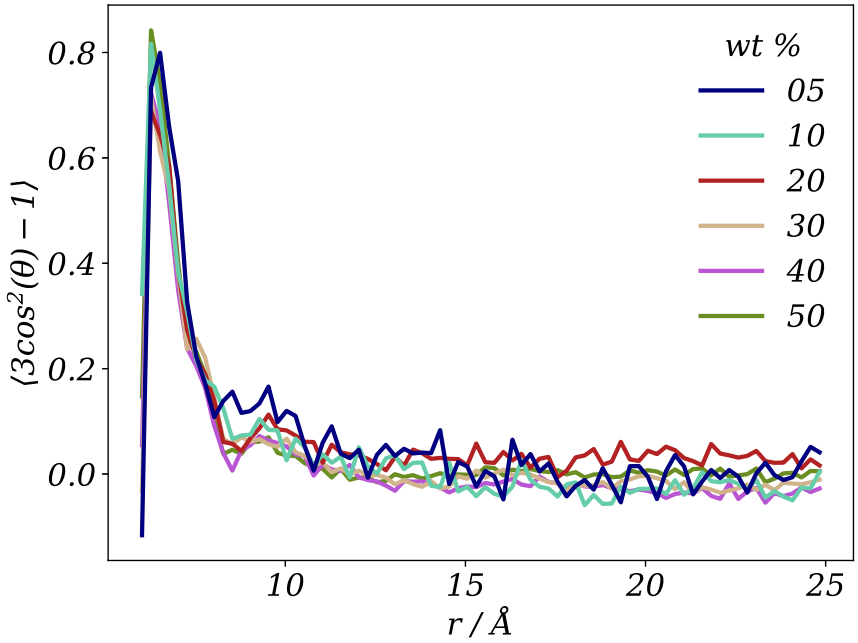

Figure 6-7: Orientational order parameter measuring the alignment of the thiophene ring to the faces of the fullerene cage for each of the P3HT-PCBM blends.

chains, decreasing the number of P3HT thiophene rings in the local area but increasing the proximity of PCBM to the polymer. Despite this, the alignment of PCBM with the thiophene rings is not enhanced and is found to be similar to Guilbert et al. at short distances for all P3HT-PCBM blends.

6.2 Frustration of P3HT

Having quantified the structure of the P3HT-PCBM blends, we can now begin to examine the vibrational modes of the system and the SAMSEN dynamical model in the context of polymeric systems. We shall attempt to observe the frustration of the P3HT backbone, as shown by a prolonged relaxation time observed by Guilbert et al., and proceed to investigate if the nature of the vibrational modes change accordingly.

6.2.1 Diffusion

Beginning with the relaxation time of the P3HT backbone, the self-intermediate scattering function of the thiophene rings, measuring the correlation in position over an

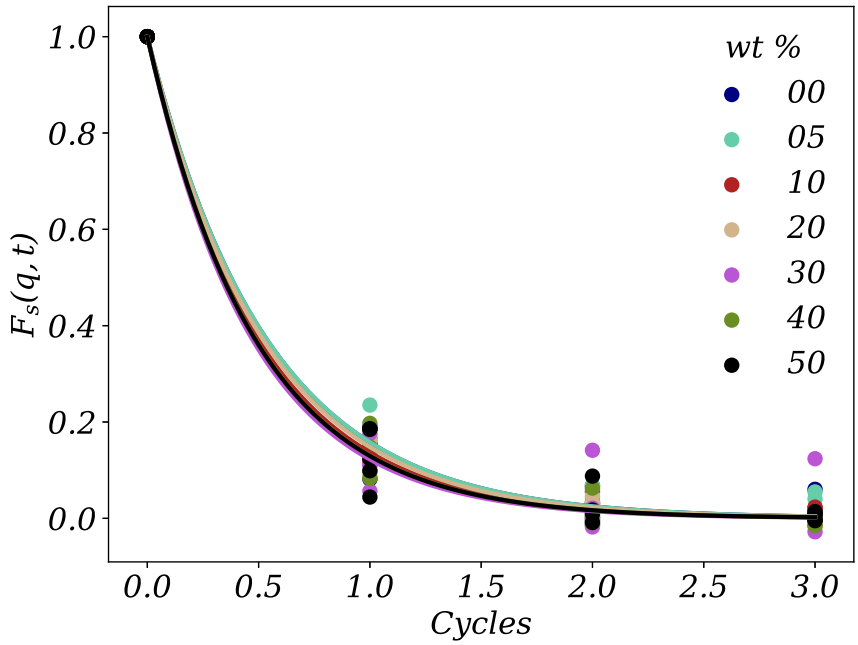

Figure 6-8: Self-intermediate scattering function, $F_s(q,t)$, of the P3HT-backbone units in the different P3HT-PCBM blends. The raw data has been provided along with a fit using an exponential decay.

inverse distance $q = \pi/\sqrt{(a)^2 + (b)^2 + (c)^2}$ where a, b and c are the length of the minimum volume enclosing ellipsoid axes (for a unit radius sphere, $|a| = |b| = |c| = 1$), is shown in figure 6-8. The raw data is shown as a scatter plot and an exponential decay has been fit to the data points. Again, there appears to be a single step relaxation with the characteristic relaxation time of approximately 0.6 cycles for the thiophene rings in each of the blends. There is no clear dependence of the relaxation on blend mixing ratio and the variability in relaxation time is dominated by uncertainty in $F_s(q,t)$.

It should again be stated that this is the relaxation time is within the SAMSEN model. The physical relaxation process which is driven by a vast distribution of modes and their relative contribution to relaxation will, in part, be determined by the frequency, with the dynamics playing out on much shorter timescales. The harmonic displacements are also only strictly valid over short distances. However, if one accepts a harmonic treatment can be used on these structures (by assuming proximity to the minima in a potential acting between coarsened rigid sections rather than atoms and can be described at short displacements by a harmonic approximation) and that the modes describe a propensity to move, then it would appear that the propensity of the low-frequency modes to relax the structure is unaltered between the different P3HT-PCBM

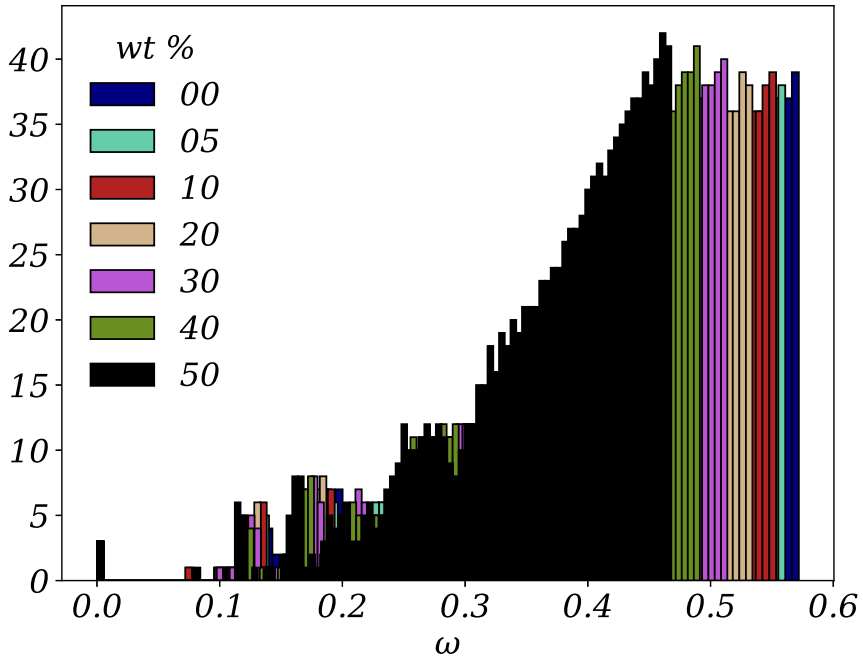

Figure 6-9: Distributions of the lowest 1,200 vibrational modes for the different P3HT-PCBM blends by weight-percent of PCBM averaged over 20 states separated by 2 SAMSEN cycles.

mixing ratios.

6.2.2 Vibrational Modes

Looking at the frequencies of the vibrational modes themselves, we can now examine the timescale of the low-frequency mode relaxation pathways. In figure 6-9, the distribution of the frequencies of the lowest 1,200 vibrational modes for the different P3HT-PCBM blends are displayed. The frequencies have been sampled every 2 cycle periods.

Between the different mixing ratios we see that the distributions of the vibrational modes of the system shift in frequency to lower frequencies as the concentration of PCBM is increased. The band-like structure is maintained for the distributions, although the band gaps do begin to narrow. The shift affects the entirety of the spectrum, including the lowest 24 modes (m_1^{24}) which were used here to relax the system.

As the relaxation time (in terms of the number of cycles) was approximately equal between the different systems, a decrease in frequency of the modes implies a decrease in frequency of the relaxation event. This means that as the pure P3HT system is

doped by increasing amounts of PCBM, the vibrational modes slow and the relaxation time increases. Looking at the quasi-elastic neutron scattering results of Guilbert et al. [118], a shift in relaxation time to longer timescales is seen for both the system as a whole and each of the components at 360K as P3HT is blended with PCBM. Capturing this result is perhaps unsurprising - we have a simulated system which, for a large part, captures the structure of an amorphous P3HT system, including local orientation, and, as we expect the structure to determine the vibrational modes, we also would expect to capture a similar shift in vibrational modes if the assumptions underlying the dynamical model were valid. We are, however, unable to calculate the full spectrum for the samples in reasonable computational timescales and the validity is limited by our approximations (i.e. harmonic potential without restoring forces, $k = 1$ kg ps^{-1}, $r_{ij} = r_0$, $\alpha = 1.1$, etc.). We also have not performed a molecular dynamics simulation, so should assume we are considering the modes of the structures that are near the high-energy minima on the potential energy landscape.

However, it would appear that the increased linearity of P3HT chains, not at increased wrapping around the PCBM, coincides with a reduction in the relaxation time. The frustration of the P3HT backbone would appear to be caused by proximity to PCBM, still adopting a partially co-facial position, but with less curvature of the chain and not an increase in the strength of wrapping. The argument presented here is, instead, that the increased concentration of PCBM increases the local density around the polymer chain which creates greater structural arrest and slows relaxation.

This idea is reinforced by looking at the character of the vibrational modes, shown in figure 6-10. There is no difference in the average local dot product or the phase quotient of the P3HT thiophene rings: the collectivity of the polymer backbone in terms of just the direction of displacements and the overall coherence of the vibration is the same, regardless of blending ratio. However, there are some differences in the relative contribution of the thiophene rigid sections to the participation ratio with higher PCBM concentrations limiting the amplitude of displacement of the polymer backbone. This change in contribution to the participation ratio, however, corresponds to an increase in the number of nodes in the elastic network (the polymer background contribution is reduced as the number of nodes increases with increasing PCBM ratio).

In summary, SAMSEN's dynamical model finds that the propensity for the low-frequency modes to relax the structure (produced by the SAMSEN structural model) and the collectiveness of those modes is unaltered by altering the P3HT-PCBM mixing ratio.

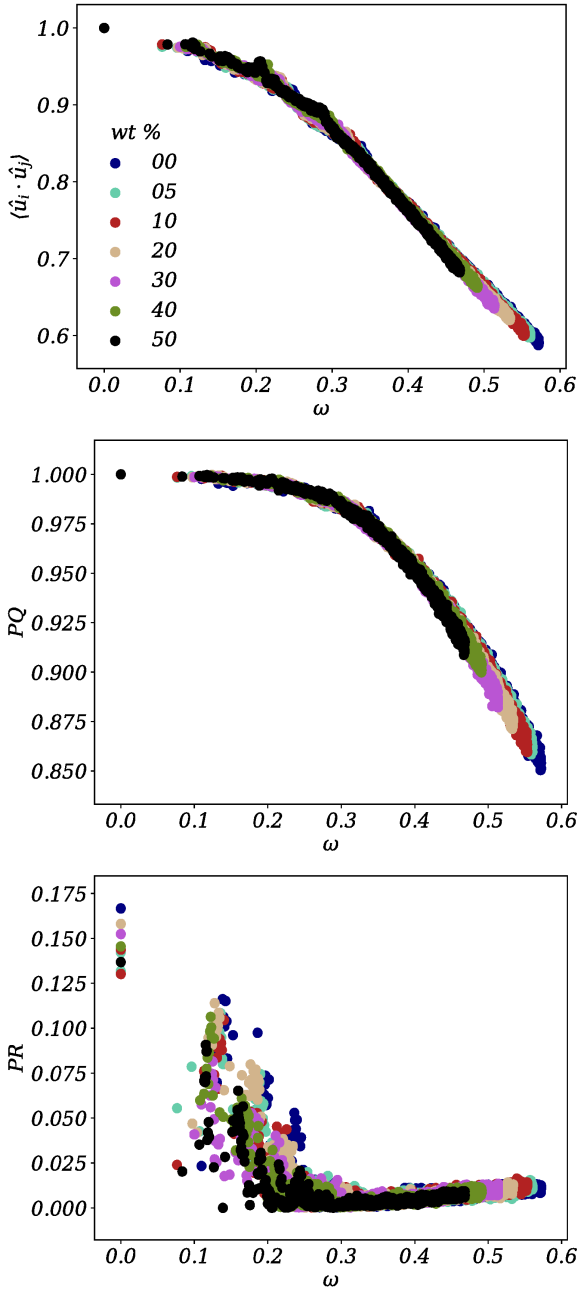

Figure 6-10: The local average dot product, the phase quotient and the participation ratio of the P3HT backbone units in the different P3HT-PCBM blends.

6.3 Summary

This chapter was devoted to the application of SAMSEN to polymeric systems. Using the SAMSEN structural model, which does not constrain the dihedral angle between rigid sections, morphologies of poly(3-hexylthiophene) (P3HT) were generated in the pure amorphous phase and in blends with PCBM. P3HT was an ideal candidate for use in SAMSEN, having a weak torsional potential. The structure of the morphologies were found to suffer from the same limitations that the small molecule systems did: that the structures appeared to be of a higher temperature than was intended. Once again we expect that SAMSEN does not bias its sampling of states thermodynamically to low energy configurations (resulting in the appearance of higher temperature structures). In this model, the only structural difference observed as the PCBM concentration is increased is that the polymer chains appear straighter, or at least, the tangents along the backbone are more correlated. These less bent chains appear to allow more PCBM to occupy the volume surrounding the P3HT thiophene units, however no increase in the tendency for the P3HT thiophene rings to align with the face of the PCBM fullerene cages is observed.

The SAMSEN dynamical model applied to P3HT was also found to achieve a computationally fast structural relaxation time. Again, the low frequency modes showed highly collective motion (albeit with narrower gaps between bands) and specifically showed high collectivity of the thiophene backbones within each mode. The dynamical model was also able to demonstrate that there is an increase in relaxation timescales as more PCBM is added to the P3HT system, as found by Guilbert et al. [118]. However the cause did not appear to be increased wrapping of the polymer backbone around the PCBM fullerene cages as had been suggested, instead showing similar local structures and collective dynamics at all PCBM concentrations (including a system containing no PCBM). It would appear that the reasoning that increased polymer chain wrapping is frustrating the polymer backbone is not supported by the simulation results presented here. Instead, the increased persistence length and increased packing density (the only observed structural changes in this analysis) corresponds to shift in the vibrational frequencies of the system to lower frequencies (longer timescales). Translating the relaxation time in cycles to relaxation time on a relative scale ($1/\omega$), it would therefore appear that the increased PCBM concentration and P3HT persistence are the cause of the reduced relaxation time.

It should however be reiterated that the lack of control over the torsional angles as well as the neglect of any electrostatic forces will have had an impact upon the simulations.

Most notably, effects such as π-π stacking will not have been encouraged by this model in a way they might have been under a conventional force-field applied in a molecular dynamics or Monte Carlo simulation. This is perhaps likely to have been a cause of the lack of alignment between the thiophene rings in P3HT and the fullerene cage of PCBM. Similarly, any form of crystallinity or medium-to-long range order in either the P3HT or the PCBM will have been muted in this simulation compared to other methods such as MD or MC. As the dihedral potential is not considered at all by the SAMSEN simulation, the polymer may be expected to appear much more disordered as neighbouring repeat units can rotate (relatively) freely about their shared bonds. This may also help to explain why the system appeared at a higher temperature than intended and the persistence length shorter than expected.

Chapter 7

Excitons in P3HT Nanoparticles

P3HT has been the cornerstone of OPV research for the last two decades but has slowly been displaced by other hole transporting polymers and donor-acceptor co-polymers. In recent years, P3HT has found use in organic nanoparticle-based OPV devices which, due to significantly lower manufacturing costs of the particular process, is approaching commercial viability [253, 254, 18].

The group at the Centre for Organic Electronics, University of Newcastle, Australia, fabricate their P3HT nanoparticles using water as their solvent. They can perform a similar process for PCBM and the fabricated device consist of a blend of P3HT and PCBM nanoparticles which are then heated and sintered - creating a pure core and a blended interface between the particles. By varying the heating process they can alter the amount of blending that occurs and, by varying the concentration of sodium dodecyl sulfate (SDS), they can modify the size of the nanoparticles [18]. This modifies both the structure of the material and transport properties such as the exciton dissociation efficiency. However, during previous experimental and computational work on exciton transport in these nanoparticles, Holmes et al. [18] noticed extremely large exciton dissociation efficiencies which could only explained by a very large Förster radius - larger than the group had previously determined it to be experimentally [80].

The previous computational work employed a lattice model. The nanoparticle films were imaged by scanning electron microscope (SEM) and using a Hough transform, the shape of the nanoparticles were captured and, accordingly, the type of nanoparticle (the PCBM nanoparticles tend to be much smoother and rounder, while the P3HT nanoparticles are more jagged with distinct faces). A lattice with 3 nm spacings was overlaid and, if a nanoparticle enclosed the lattice point, the lattice was then designated

as that material. If no nanoparticle enclosed the lattice then that point was removed. A kinetic Monte Carlo simulation was performed to estimate the exciton dissociation efficiency - an exciton separated into an electron and hole when reached a boundary of two nanoparticles. Using this model, the high exciton dissociation efficiency could be only explained by large diffusion length of approximately 50 nm [18] (and a correspondingly large Förster radius). This model, however, does not take into account position, orientation and any correlations into account, nor account for the difference in structure between the amorphous and crystalline sub-domains - both of which are present in P3HT nanoparticles and P3HT films [18, 244, 255].

During a collaboration and secondment to the Centre for Organic Electronics the following work was performed to take morphology into account. In this chapter, the amorphous/crystalline structure of the nanoparticles will be characterised by comparing to an amorphous P3HT system produced by SAMSEN as well as a perfect crystal based upon the DFT optimised P3HT 6-mer and X-ray diffraction (XRD) measurements. These structures will also be used as the base morphology for exciton diffusion length calculations using the transfer matrix method described in section 3.4.1.

7.1 Creating P3HT Structures

Amorphous Morphologies

The amorphous morphologies used in this section were created in a similar manner to that in the previous section. The chosen number of polymers were created at random positions and orientations, replicating the chain in the straight line so that they are the desired number of monomers long. These are then compressed at a rate of $C_f = 0.99995$ each step, ensuring the conditions of a valid SAMSEN state are maintained with thresholds set to $\chi_c = 0.6$ $\mathring{A}$ and $\chi_f = 0.3$ $\mathring{A}$.

Due to the size of systems being studied, the dynamical model was not employed in this section. The increased number of nodes increases the size of random-access memory required by the ARPACK++ implementation to diagonalise the Hessian matrix - the factorisation process undermines any memory reduction achieved by sparse-matrix representation. At more than 70,000 rigid sections, the requirements are larger than the computing resources available for this project (32 GB). If a less memory-demanding implementation of diagonalisation is available, this will enable the modes to be employed on larger systems, and I encourage its use. Preliminary investigation would suggest

No. of Polymers	Chain Length / monomers	Box Length / $\AA$	Density / g/cm^3
50	10	50	1.106
399	10	100	1.103
1345	10	150	1.102
3187	10	200	1.101

Table 7.1: The specifications of the amorphous P3HT systems used in the transport calculations. Larger systems of 6225, 10756 and 25496 were generated but were too large to perform the matrix diagonalisation upon

a package such as the Multifrontal Massively Parallel Sparse Direct Solver (MUMPS) version 5.2.0 [256] onwards may solve the factorisation memory issue.

Table 7.1 provides the specification of the morphologies produced after compression which will be used in the transport calculations that follow.

Nanoparticle

In an attempt to replicate the nanoparticle morphology, a system of 30,000 P3HT repeat units (300 polymers, each 100 monomers long, terminated with hydrogens) in a box of dimensions $L_x = L_y = L_z = 120$ nm without periodic boundary conditions was created. Atoms which reach the boundary of the simulation volume would be displaced so that the coordinate beyond the box boundary becomes equal to the box boundary - however this did not occur. The system was compressed towards the centre of the box (rather than the box edges) by multiplying the difference in coordinates between the given atom and the centre of the box by 0.00005 and subtracting this value from the given coordinate. The box dimensions remained unchanged. This was continued until the collision and mismatch correction could not establish a valid SAMSEN state after 50,000 attempts. At this point the multiplication factor was lowered to 0.00001 and the simulation continued for 5 days. The final state was then taken to be the nanoparticle morphology. This simulation was performed on the University of Bath's High Performance Computing Facility, Balena.

Producing a nanoparticle morphology, rather than relying on those with periodic boundary conditions, allows us to also consider the effect of density profile within the nanoparticle. In the morphology produced here, the centre is significantly denser at approximately 1.35 to 1.40 g/cm^3 and the system becomes sparser near the edges reducing to around 1.2 g/cm^3.

The final nanoparticle is roughly 20 nm in diameter and mostly spherical although

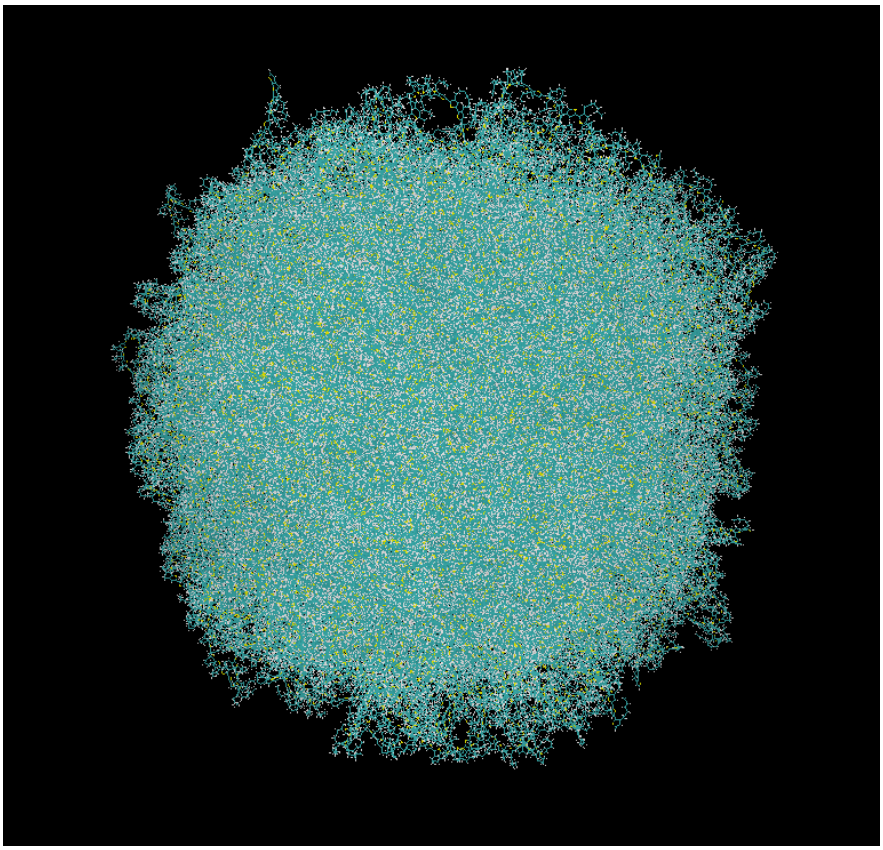

Figure 7-1: VMD generated image of the P3HT nanoparticle morphology.

some straighter faces (artefacts from the initial setup) remain in the final system. A visualisation of the final nanoparticle produced using VMD [257] is shown in figure 7-1.

Crystal

Two P3HT monomer units, in head-tail configuration as well as XRD data of the P3HT nanoparticles (performed by Levi Tegg, University of Newcastle, New South Wales, Australia) were used to create the basis of the P3HT crystal. The optimised DFT structure and its repeat head-tail structure determined the length of one of the unit cell axes. The XRD data, showed that the stacking distance (between two neighbouring thiophene rings in adjacent chains) was 3.75 $\mathring{A}$ and that the final axes in the unit cell (the length between chains in the direction of the side chains) was 16.5 $\mathring{A}$. Taking the original two monomers (the unit cell), these were replicated in each of the axes for a chosen number of unit cells. The structure was created such that the chains continued unbroken over the periodic boundary and no terminating atoms were added.

No. of Unit Cells	No. of Monomers	$L_x/\mathring{A}$	$L_y/\mathring{A}$	$L_z/\mathring{A}$	Density / g/cm^3
10^3	2,000	165.0	37.5	77.7677	1.15
15^3	6,750	247.5	56.25	116.457	1.15
20^3	16,000	330.0	75.0	115.275	1.15
25^3	31,250	412.5	93.75	194.094	1.15

Table 7.2: The specifications of the crystalline P3HT systems used in the transport calculations in this section with structures based upon XRD data from P3HT nanoparticles

This structure therefore represents a cold, perfect crystal. The structures generated lack any disorder and these will be the crystal structures used in the following sections. To thermalise the crystal a molecular dynamics run at the desired temperature could be performed, however this was not done in this work. The sizes of crystals used in this section are detailed in table 7.2.

7.2 XRD of P3HT Nanoparticles

Having produced amorphous, crystal and nanoparticle morphologies, we can begin to characterise them and compare their structure to those of the real P3HT nanoparticles, fabricated at the Centre for Organic Electronics, University of Newcastle, Australia.

X-ray diffraction (XRD) measurements of the structure of the nanoparticles were also performed on four different sized P3HT nanoparticles (controlled by SDS concentration) suspended in water and at low density (constituting around 0.05% of the solution) such that interactions with other nanoparticles could be ignored. The XRD measurement was performed using a Phillips X'Pert MPD XRD operated in a Bragg-Brentano 2θ reflection geometry with a Cu Kα anode at wavelength, $\lambda = 1.541$ $\mathring{A}$. A beam mask was used to maintain a constant irradiated area of 15 nm $\times$ 15 nm.

The P3HT nanoparticle (NP) samples (with water and SDS) were placed in 26 mm diameter 3D-printed PLA cups and left unsealed. The samples were rotated at 0.25 revolutions per second in the sample holder. The XRD measurement was performed over 40 minutes. During the series of measurements, much of the water in the samples evaporated. To attempt to maintain a roughly constant water-level, when the water level lowered, additional water was added to the sample holder to return it to the correct level - this assessment and correction was performed manually. A background measurement of a pure water sample was also taken. The diffraction patterns for the water sample was subtracted from the pattern of the samples.

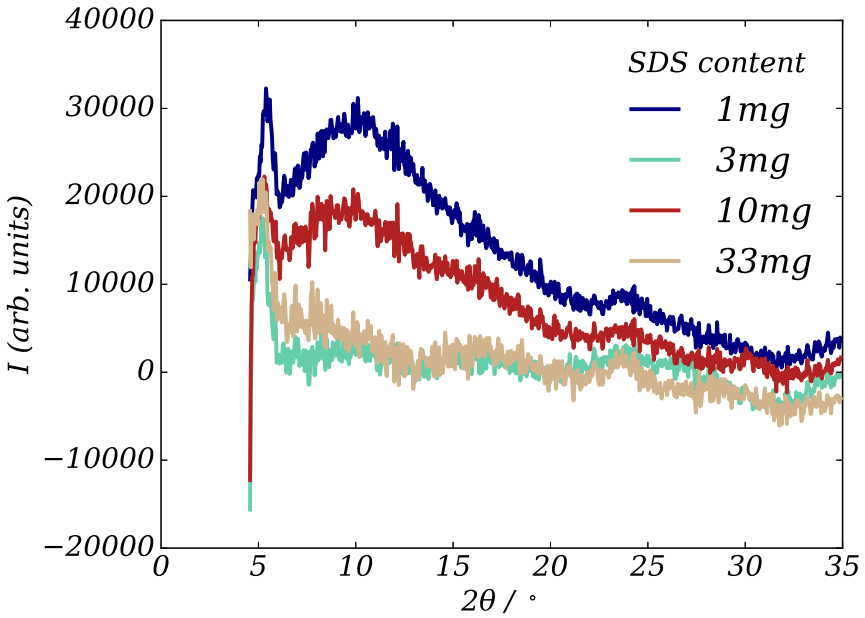

Figure 7-2: XRD scattering intensities for the P3HT nanoparticles, fabricated with differing SDS content.

The XRD data, presented in figure 7-2 shows few, if any, trends that we could use to infer information about the effect of SDS concentration, and therefore nanoparticle size, on the structure of the P3HT nanoparticles. To begin with, there is a broad peak between 6° and 20° in both the 1 mg and 10 mg SDS spectra which are not present in the 2 mg or 33 mg spectra. So it cannot be said that SDS concentration alone influences this peak. There are also very few other discernible features apart from the 5° (100) lamella stacking peak which has a similar broadness for each of the nanoparticles and shifts to slightly lower angles with increasing SDS, and the 3.75 $\mathring{A}$ (010) π-π stacking peak which follows no discernible trend in terms of amplitude or width.

The spectra also have some noteworthy issues. The large negative intensities at the lowest angles for the 10 mg and 2 mg and around 30° for the 2 mg and 33 mg SDS spectra are only possible because of the background subtraction. However, the spectra are not negative for all SDS concentrations. This effect comes about because the water in the PLA cups evaporated during the measurement. This means that, throughout the spectra, there are slightly different amounts of water contributing to the pattern and subtracting the XRD measurement of pure water creates these artefacts in the spectra. The scan was performed from low-angle to high-angle so, these explain the negative intensities at high angle. Ignoring any regions of negative intensity, the 2 mg and 33 mg spectra look very similar to other reported XRD spectra of P3HT [258,

259, 260, 261, 262]. The broad shoulder in the 1 mg and 10 mg spectra is not present in other reports. This may be caused by the SDS (although the lack of a shoulder at 33 mg SDS perhaps discredits this suggestion) or by faster-than-expected evaporation affecting the whole pattern. This uncertainty will affect some of the conclusions that may be drawn from this chapter.

7.2.1 Comparison with SAMSEN Morphology

The intensities produced by coherent diffraction between planes in an powder X-ray diffraction measurement can be calculated from the Debye scattering formula

$$I(Q) = \sum_{i}^{N} \sum_{j}^{N} f_i(Q) f_j(Q) \frac{\sin(Qr_{ij})}{Qr_{ij}} \tag{7.1}$$

where $f_i(Q)$ is the atomic scattering factor for the given element and has angular dependence (which we will now omit from the notation for brevity) [263]. This equation can be modified to separate the incoherent scattering (self-scattering) components

$$I(Q) = N\langle f^2 \rangle + \sum_{i}^{N} \sum_{j,j\neq i}^{N} f_i f_j \frac{\sin(Qr_{ij})}{Qr_{ij}} \tag{7.2}$$

and is suitable for calculation from simulated atomistic morphologies. The atomic scattering factors (the angular dependence of intensities produced by a single atom of a given element) are generally inferred from experiment or density-functional theory calculations and can be looked up in suitable data tables. Here we shall use the data from the International Tables for Crystallography (2006) [264] where the fit to the angular dependence has one form at low angles and another at high angles.

For low angle ($0 < \sin(\theta)/\lambda < 2\mathring{A}^{-1}$) atomic scattering factors, this equation (and appropriate coefficients) provide a good representation

$$f(sin(\theta)/\lambda) = \sum_{i=1}^{4} a_i \exp(-b_i \sin^2(\theta)/\lambda^2) + c \tag{7.3}$$

while at high angles ($2\mathring{A}^{-1} < \sin(\theta)/\lambda < 6\mathring{A}^{-1}$) this equation proves a better representation

	H	C	S
a_1	0.489918	2.31000	6.90530
b_1	20.6593	20.8439	1.46790
a_2	0.262003	1.02000	5.20340
b_2	7.74039	0.2075	22.2151
a_3	0.196767	1.58860	1.43790
b_3	49.5519	0.568700	0.253600
a_4	0.049879	0.865000	1.58630
b_4	2.20159	51.6512	56.1720
c	0.001305	0.215600	0.866900

Table 7.3: Coefficients used in the approximation to the atomic scattering factors for each element

$$\ln\left(f(sin(\theta)/\lambda)\right) = \sum_{i=0}^{3} a_i(\sin(\theta)/\lambda)^i \,. \tag{7.4}$$

Using values from table 7.3 in equation 7.3 and another set [264] for the high angle equation 7.4, the X-ray diffraction patterns of the pure P3HT SAMSEN morphology used in the previous chapter (6,000 monomers under periodic boundary conditions) and the SAMSEN nanoparticle morphology were calculated. For the nanoparticle (NP), there was no averaging as there was only a single frame, and the pattern for the SAMSEN periodic boundary conditions (PBC) morphology was generated using atomic positions from a single frame. The XRD pattern for the perfect crystal with 25 $\times$ 25 $\times$ 25 unit cells, introduced in section 7.1 was also calculated. The diffraction patterns for all three morphologies, using $\lambda = 1.541$ $\mathring{A}$ and scanning the range 1° to 35° in 0.01° steps, is presented in figure 7-3.

Each of the SAMSEN XRD patterns are flat for much of the range (particularly at double angles above 15°). The peak at 5° (16.5 $\mathring{A}$, associated with (100) lamella stacking [261, 262]) —the only peak consistently found across the literature [258, 259, 260] —is present but is very weak. At lower angles, the intensity increases greatly. Between the NP and PBC morphologies, there are little differences at high angles, but the 5° peak is broadened and shifted to lower angles and the rise in intensity at smaller angles is more profound. The crystalline morphology is dominated by a sinusoidal signal at all angles in the probed range, however captures the 16.5 $\mathring{A}$ peak and also captures additional intensity at $\approx$ 24° (3.75 $\mathring{A}$, associated with (010) π-π stacking [261, 262]). The crystalline spectrum also captures the 8.8 $\mathring{A}$ (200) peak at 10° and the 5.5 $\mathring{A}$ (111) reflection associated with a two-layer unit cell [262] as well as additional reflections.

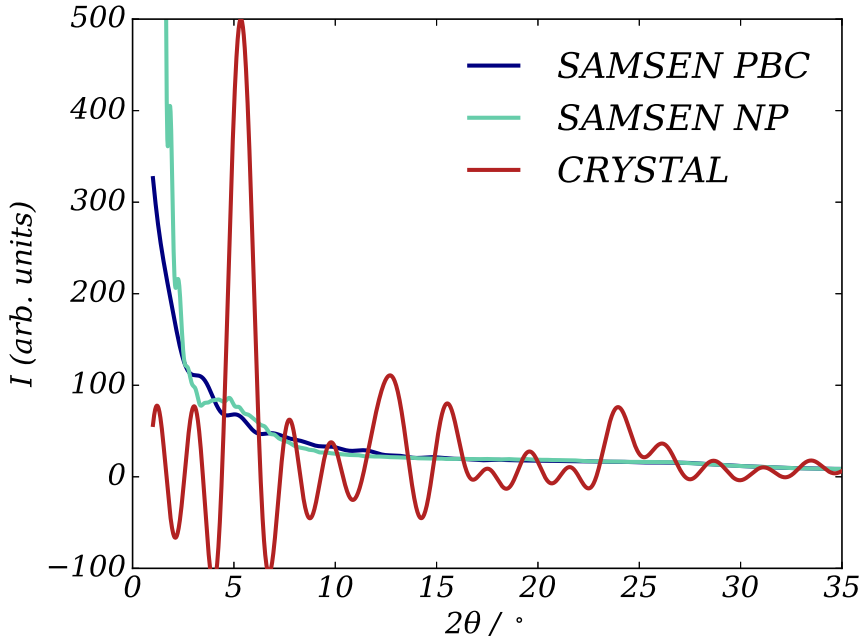

Figure 7-3: The XRD patterns calculated for a pure P3HT crystal, the SAMSEN amorphous system (6,000 monomers used in section 6.1) as well as the SAMSEN P3HT nanoparticle attempt. The peak around 5° corresponds to the lamella stacking between chains.

These additional peaks, born from the order in the local structure (and potentially some computational artefacts), are all missing from the SAMSEN spectra.

Once again it would appear that SAMSEN shows muted structural details compared to the literature, even in the case of amorphous spectra. In this case the absence of any peak at short distances (higher angles) and the diminished long range (low angle) structure is also due to the structural model not incorporating electrostatics or including torsional control between repeat units. This may aid the sampling of states, but limits SAMSENs ability to predict structures without the assistance of a short MD run, which was not performed here. The SAMSEN structures are highly amorphous, while the literature data is often a mix of amorphous and semi-crystalline morphologies which would act to amplify the 5° peak. This places strict limits on the usefulness of such structural comparison with the nanoparticles or P3HT films. However, in this particular instance, we can use the knowledge that our system is amorphous and lacking in local order (and semi-crystallinity) to our advantage.

Bearing this in mind, we can now take both the amorphous SAMSEN structure which we assume to be entirely devoid of local order (including a lack of $\pi - \pi$ and lamella stacking) and combine its diffraction pattern with that of the 'perfect' crystal. By altering the mixing ratios of the spectra we can estimate the crystallinity or, more accurately, the crystalline-to-amorphous ratio of the P3HT nanoparticles. If the XRD spectra, experimentally measured, can be recreated by a simple mixture of an amorphous and crystalline XRD pattern we could begin to comment on the likelihood of finding any given monomer in a crystalline or amorphous environment.

Using the XRD spectra of the P3HT nanoparticles, fabricated with different SDS concentrations, discussed in section 7.2 and shown in figure 7-2, as the optimal spectra, a genetic algorithm was employed to try and find the 'fittest' (most optimal) combination of the SAMSEN amorphous and 'ideal' crystalline spectra. Genetic algorithms are a computational technique which, effectively, perform a minimisation of some function based upon a corresponding cost function. Here, we are trying to minimise the distance between a guess spectrum and a target spectrum: the measured XRD data.

The genetic algorithm was initialised by creating a population of 1,000 random combinations of amorphous and crystalline contributions which were normalised by the number of monomers contributing to the respective spectra. The algorithm proceeded as follows. At each step towards convergence, the population would be tested for fitness by producing a spectrum adding the absolute contributions of each of the amorphous and crystalline spectra by measuring the mean squared error between it and the target

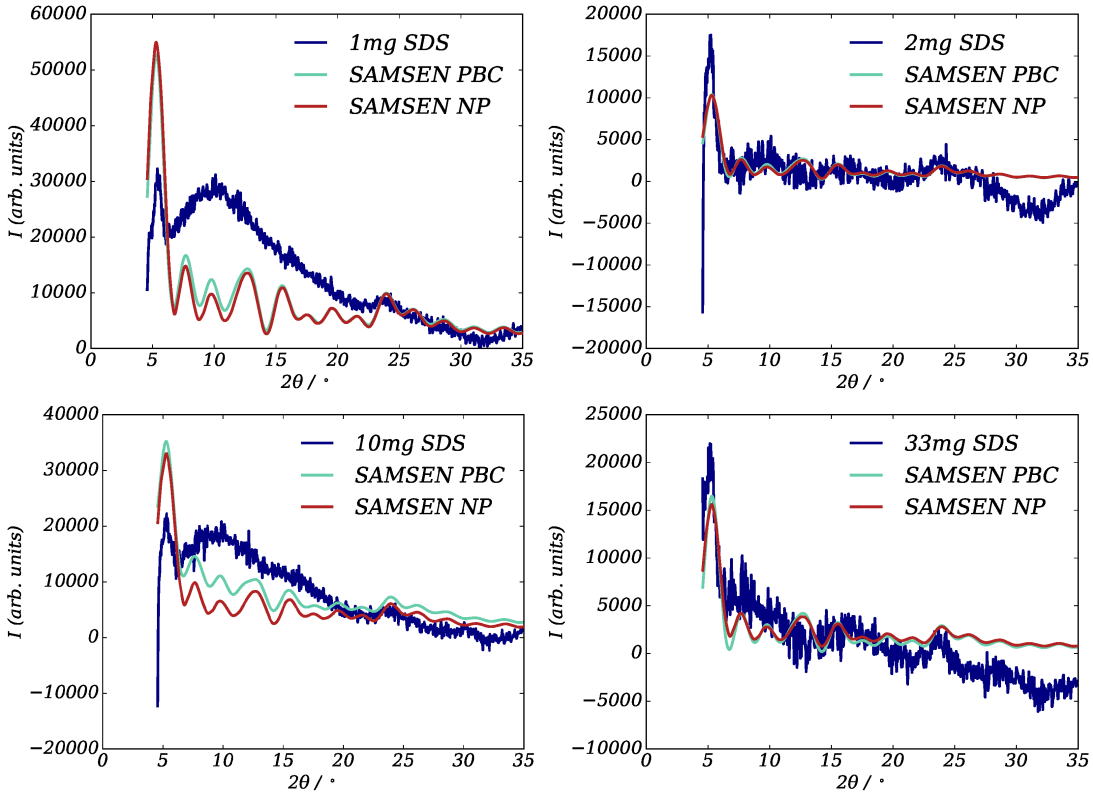

Figure 7-4: XRD scattering intensities for the P3HT nanoparticles compared to the XRD pattern calculated for a mixture of a pure crystalline and SAMSEN amorphous system as well as a pure P3HT crystal and the SAMSEN nanoparticle.

spectrum (the P3HT nanoparticle XRD data). The population are then ranked, with the lowest mean squared error placed first. The bottom (least fit) 35% of the population 'die' and are replaced by: 5% produced by exclusive breeding of the top 10% of the population (the average of both of the absolute spectra contributions is used for the new contributions for their offspring), 20% produced by breeding randomly within the top 40% of the population, and 10% filled by immigration into the system (a new random selection of spectra contributions were generated). Randomly across this new population, 2% were given a mutation of one of their absolute contributions (a new random number was generated). This ranking, reforming of the population and mutation, continued until convergence which was taken to be when 80% of the population (before mutation) had absolute contributions within 10^{-9} of the 'fittest' member of the population. This method was applied for the mixing of both the SAMSEN PBC and 'ideal' crystalline systems and the SAMSEN NP and 'ideal' crystalline systems to each of the XRD patterns for the different SDS concentrations during the nanoparticle fabrication. The resulting 'fittest' spectra are shown in figure 7-4 and the relative contributions to the spectra for the SAMSEN PBC and the crystalline-to-amorphous ratios are displayed in table 7.4.

The fits of the combined simulation spectra to the XRD measurements are varied in quality. Fits to the spectra of the P3HT nanoparticles fabricated with 1 mg and 10 mg SDS concentrations are poor at many angles, except at the 25° peak - they generally overestimate the intensity of the 5° peak and miss the broad shoulder between 6° and 20°. This is true for both the SAMSEN PBC and SAMSEN NP fitted spectra, although these two deviate from one another across the range of the broad shoulder. In the case of the 2 mg and 33 mg SDS concentrations, the fits are much better with close agreement across the range, with the exception of a slightly reduced 5° peak and missing the decrease in intensity (below 0) at around 30°. These 2 mg and 33 mg spectra also more closely match the literature (at $2\theta < 20°$) which provides more confidence in the ratio of crystalline-to-amorphous produced by the genetic algorithm.

As we have noted already, the XRD data is not perfect and perhaps distorted by the measurement being performed with the nanoparticles suspended in water —a measurement we believe has not been reported before. The broad shoulder could not be attributed to any one reflection and is not observed in the literature data. The drop in intensity at high angles is also a result of the water subtraction which was difficult to perform reliably as the amount of water in the samples varied throughout the measurements and needed to be replenished. This, along with interference, from the SDS which coat the nanoparticles, and other surface effects may have altered the XRD mea-

SDS content	Amorphous Contribution	Crystalline Contribution
1 mg	0.466	0.533
2 mg	0.366	0.633
10 mg	0.653	0.347
33 mg	0.330	0.670

Table 7.4: The relative contributions to the 'fittest' XRD spectra - the combination of the SAMSEN PBC and 0 K crystal XRD patterns that produced the closest fit to the experimental XRD spectra of the P3HT nanoparticles (after background subtraction).

surement in a way which could not be replicated in these simulations. The SAMSEN NP morphology, which attempted to recreate some of these effects (most notably the non-uniform density profile), provided no improvement in comparison to the measured spectra. Similarly the SAMSEN PBC morphologies themselves were not ideal (with an MD simulation not performed the morphology was likely to be a high energy state).

Given these limitations on both the structures, the fit and the XRD data itself, I present the ratio of amorphous and crystalline components of the P3HT XRD data in table 7.4. Across the range, there is no clear dependency of crystallinity on SDS concentration (and therefore nanoparticle size), as there is not for the XRD spectra themselves. However, if one just looks at the 2 mg and 33 mg spectra, which conform more strongly to the literature data and provide the closer fit, we find very little change in the internal crystallinity despite the increased size of nanoparticle, with roughly one third of the nanoparticle being amorphous in nature and the remaining two thirds being crystalline. There is perhaps a slight increase in the crystalline contribution to the XRD as we move from 2 mg to 33 mg of SDS during fabrication, but with so much uncertainty (particularly over the water content, which limits the quality of the fit at high angles) it can be hard to say this for sure. This is true of the spectra for the same reasons.

There are, of course, obvious limitations of this fitting approach to determine the crystal/amorphous content of the P3HT nanoparticles. The distribution of SAMSEN structures sampled is, as we have previously shown with PCBM, unlikely to be at thermodynamic equilibrium for a 300 K system. Similarly, the crystal presents a zero-temperature structure and not a structure at room temperature. This means the 5° peak (the others too) is artificially enhanced and less crystallinity is required to reproduce the experimental spectrum that a 300 K structure might. Conversely, the SAMSEN PBC system potentially represents far too high a temperature, with its XRD features diminished, and an increase in amorphous content is required. The ratios

presented here are likely to overestimate the amorphous content and underestimate the crystalline proportion of the real P3HT nanoparticles. Therefore, both of these structures would first need to be put in molecular dynamics (or any thermodynamic molecular simulation method) and equilibirated to 300 K. The combination of these spectra, still taking the average over the now-equilibirated SAMSEN spectra, would then be able to determine a more accurate crystal-to-amorphous ratio. However, being able to separate the structures into a highly amorphous component and pure crystal component enables such a study which methods, prone to finding mixed domains, may not.

7.3 Exciton Diffusion in P3HT

Having quantified the structure of the SAMSEN morphologies and commented on its validity with respect to experimental measurements, we can now begin to consider the charge transport properties of this material and the P3HT nanoparticles themselves. In this section we will focus on exciton diffusion, making comparison to previous computational work on these nanoparticles and values of the diffusion length inferred from experiment.

The Dipole Orientation Factor

In the SAMSEN morphologies of pure P3HT, the average dipole orientation factor is $\langle\kappa^2\rangle = 2/3$ indicating that the SAMSEN morphology has this amorphous, orientationally uncorrelated structure, previously described. The distribution of κ^2 for the 6,000 monomer morphology is shown in figure 7-5 and, statically, reproduces the distribution of a dynamically rotating system [81]. This distribution is for all pairs of monomers in the system and the large κ^2 values occur on all length-scales and dominate the exciton transport if a static dipole model is considered.

In principle, the relative dipole orientations should always be taken into account when considering dense amorphous systems where the relaxation time is long and the material below its melting temperature or, under the right preparation, under its glass transition temperature. This was noted by Förster in 1948 [78]. However, in order to draw comparison to the recent literature, calculating the relative orientation of static dipoles is going to be ignored and instead a system average value will be used. The κ^2 value of Herz et al. [84, 85] will be used in both the amorphous and crystalline system.

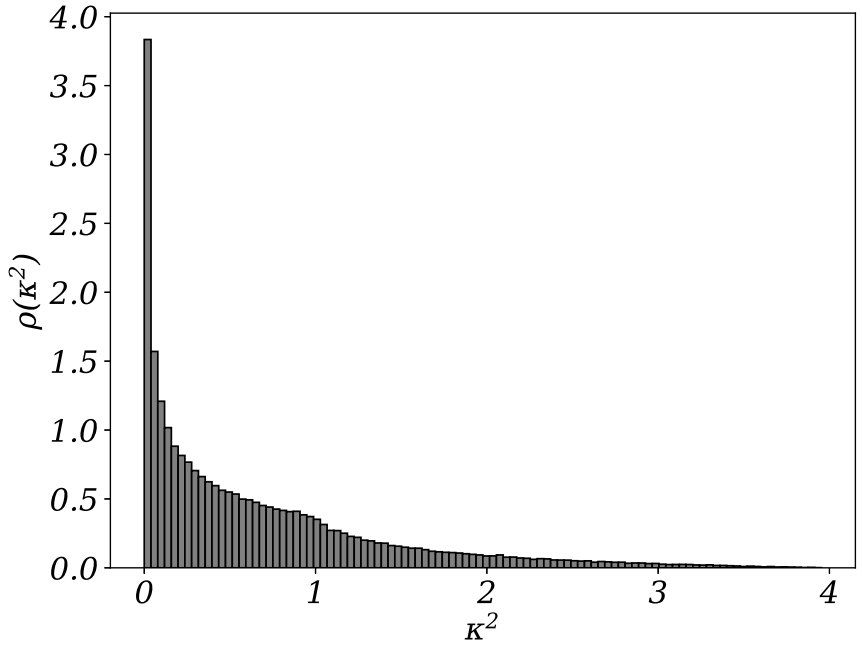

Figure 7-5: The distribution of the dipole orientation factor for all pairs of thiophene sites in the SAMSEN P3HT system. The average value of this distribution corresponds to the dynamic average value due to the high disorder in the system ($\langle\kappa^2\rangle = 2/3$

However, we expect that alignment of dipoles does occur and we expect that these pairs will have the largest exciton transfer efficiencies. Therefore, the same simulations will also be repeated at a $\kappa^2 = 4$ which represents the 'best case' scenario in terms of dipole alignment. These will then represent the lower and upper bounds of the expected exciton transport quantities such as the exciton diffusion length, produced by the transport model on these morphologies.

Estimate of the Förster Radius

To begin, we must first extract the transfer parameters. In work by Feron et al. [80], simulations of exciton diffusion and dissociation were performed for values of the exciton lifetime, $\tau = 0.9$ ns [80] and the energetic disorder, $\sigma = 0.06$ eV. Lattice-based simulations were performed using these parameters with a lattice constant of $a = 1$ nm and assuming a system-wide $\kappa^2 = 2/3$ at a variety of Förster radii. This was repeated until simulations produced a diffusion length of 8.5 nm from literature [215] at a Förster radius of $R_0 = 2.3$ nm.

Unfortunately, it is not ideal to re-apply the Förster radius to a different system.

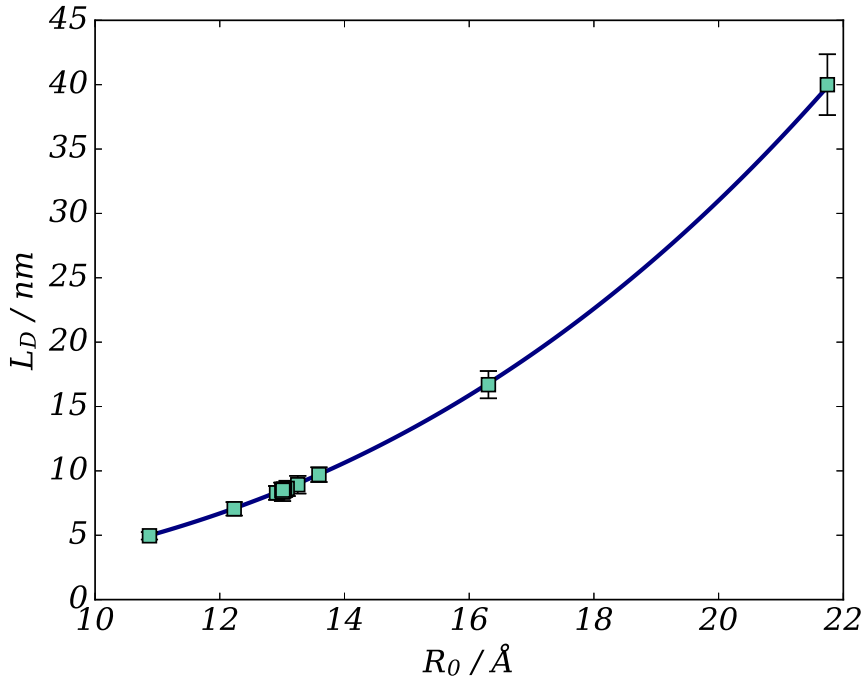

Figure 7-6: Exciton diffusion length, L_D, as a function of the Förster Radius, R_0. A binary search was performed to find the value of $R_0 = 13.76$ Å at $L_D = 8.5$ nm.

The lattice spacing will alter the calculated diffusion length [79] (that was the test for finding the original Förster radius itself in work by Holmes et al. [18]). Instead, we shall follow a similar procedure here. The diffusion length, L_D, will be calculated using the transport model described in section 3.4.2 for a selection of Förster radii and the value of R_0 that produces the correct exciton 1D diffusion length (which shall be assumed to be $L_D = 8.5$ nm as reported by Feron et al. and Shaw et al. [80, 215]) will be taken forward. To determine L_D, equation 3.29 with $Z = 1$ shall be used. Finding the correct R_0 was performed using a binary tree search, starting with an initial range of Förster radius between 40 Å and 1 Å. The midpoint of the range represents the test point. If the value of L_D found was smaller than desired, the R_0 halfway between the two extremes would become the new upper limit, if L_D was too large then this R_0 would become the lower limit. This continued until convergence. The results of this search are shown in figure 7-6 and the expected relation [78, 79] $L_D \propto {R_0}^3$ is found.

The value of R_0 to be tested at each step will be assumed to have reflected the Herz et al. estimate of κ^2 rather than the dynamic model's $\kappa^2 = 2/3$. This point is important to note when we shall swap to and from $\kappa^2 = 4$ as this produces a larger R_0 in the aligned dipole system than it would had we assumed a $\kappa^2 = 2/3$ value. For the unaligned

dipoles, therefore, $R_0 = 13.01$ $\mathring{A}$ and for the fully-aligned system $R_0 = 18.55$ $\mathring{A}$.

7.3.1 Ordered vs. Disordered

With the parameters for the Förster rate chosen and the morphologies generated, calculations of the exciton diffusion length can begin. Here we shall look at the SAMSEN amorphous PBC system (which we now believe represents roughly one third of the content of P3HT nanoparticles) and the perfect crystal (which represents roughly two thirds of the nanoparticle). We also accept the limitation of the SAMSEN states, namely that they are high in the energy landscape and unlikely to be found at equilibrium, and also the limitation of a perfect crystal which, at 300 K, would be much more disordered.

Using the morphologies of various sizes (crystal: table 7.2, amorphous: table 7.1, the transfer matrix method described in section 3.4.2 and the parameters just established, the transfer rates for each pair of monomers (between the centre of the thiophene rigid section, marked with a red cross in figure 6-1) were calculated over a time period of 1 fs using the value of R_0 that assumed $\kappa^2 = (0.845\sqrt{2/3})^2$ and an upgraded R_0 which was recalculated as $\kappa^2 \to 4$. The transfer matrix was populated with these rates and diagonalised, taking the eigenvector with an associated eigenvalue of 1 to be the steady state distribution. With knowledge of both the probability of occupancy and probability of transfer, the diffusion constant could then be calculated and converted to a 1D exciton diffusion length. This was performed for a range of cut-off radii, r_c and repeated 100 times taking site energies randomly from a Gaussian of width σ. The result for the amorphous system with $\kappa^2 = (0.845\sqrt{2/3})^2, R_0 = 13.01$ $\mathring{A}$ is shown in figure 7-7(a) and the result for $\kappa^2 = 4, R_0 = 18.55$ $\mathring{A}$ is shown in figure 7-7(b). The result for the crystalline system is shown in figure 7-8 with $\kappa^2 = (0.845\sqrt{2/3})^2, R_0 = 13.01$ $\mathring{A}$ in panel (a) and $\kappa^2 = 4, R_0 = 18.55$ $\mathring{A}$ in panel (b).

For the amorphous system with unaligned dipoles ($R_0 = 13.01$ $\mathring{A}$) it becomes clear that short cut-off radii hinders the exciton diffusion and a plateau in L_D is only achieved when the cut-off is increased to a much larger value. For the smaller systems (5 nm and 10 nm) there is also large uncertainty and also, again, restricted diffusion. The small systems also begin to show spurious behaviour when the cut-off radius reaches half the box dimension, although this is perhaps not unexpected. When system sizes are sufficiently large (15 nm and 20 nm) the diffusion is no longer hindered by the size of the simulation cell and a good estimate can be obtained. For the amorphous system with unaligned dipoles $\kappa^2 = (0.845\sqrt{2/3})^2, R_0 = 13.01$ $\mathring{A}$, this full analysis

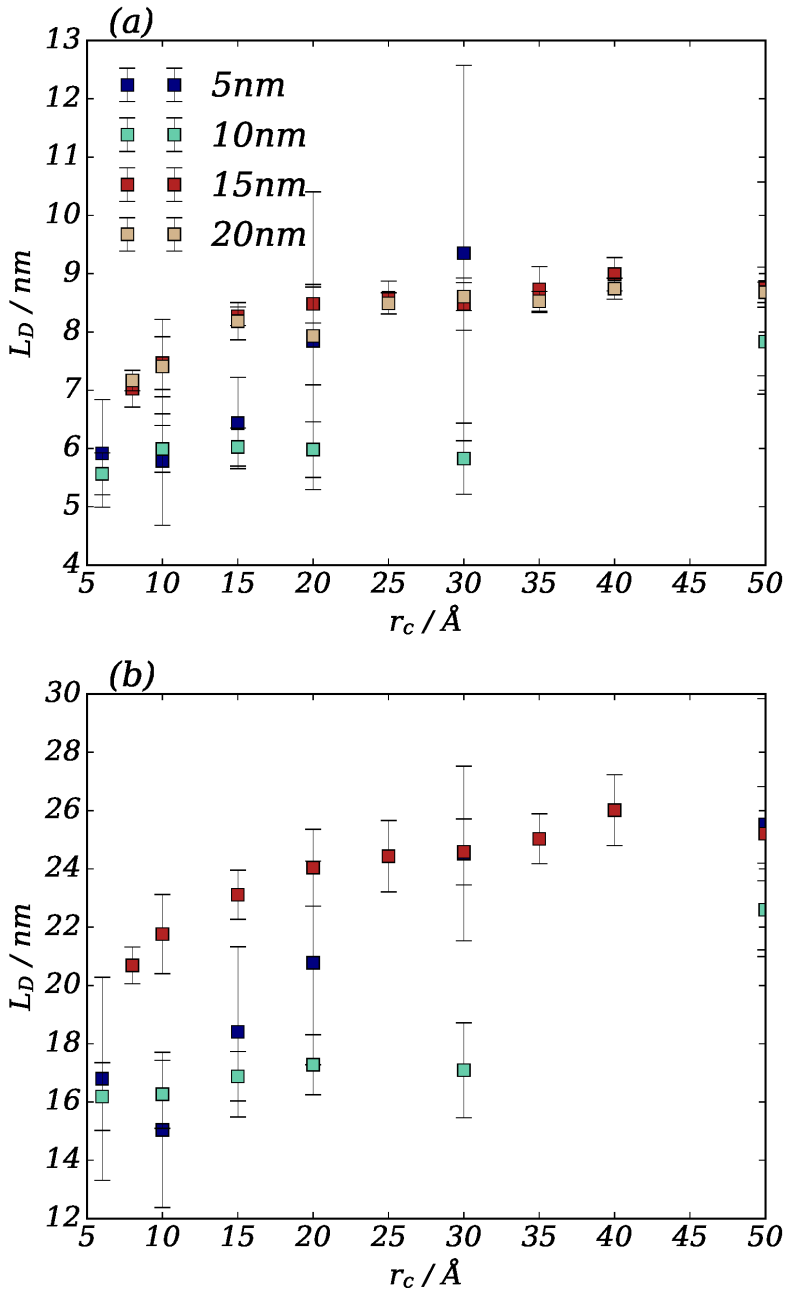

Figure 7-7: Diffusion length for the various amorphous system sizes as a function of the cut-off radius with (a) a dynamic dipole model ($\kappa^2 = 2/3$) and (b) a static dipole model with all dipoles aligned ($\kappa^2 = 4$). This last case is unrealistic in a disordered system such as this.

predicts an exciton diffusion length of $L_D = 8.8$ nm —deviating slightly from the value against which R_0 was calibrated. Upgrading the Förster radius to correspond to a $\kappa^2 = 4$ ($R_0 = 18.55$ $\mathring{A}$) produces similar results. Above half the cut-off the smaller systems produce spurious results and show large uncertainties throughout. The larger systems show a plateau at around $L_D = 25$ nm which this analysis suggests is the 1D exciton diffusion length for a structurally amorphous system with fully aligned electronic dipoles.

These values of the exciton diffusion length, assuming a system of disordered dipoles or a system of complete dipole alignment in a morphology of amorphous structure, of course are much smaller than the predicted 50 nm suggested by Holmes et al. for P3HT nanoparticles [18]. Of course, it should be noted that this second picture may be unphysical. This would suggest that the amorphous domains of the P3HT nanoparticles cannot dominate the transport properties. That they form only a third of the nanoparticle (from fitting SAMSEN and crystalline XRD patterns to experimentally measured ones in the previous section), means that this may indeed be the case and that exciton transport is dominated by crystalline regions and longer-range transport between them. This is expected for electrons and holes [110] and also for resonance transfer [81] due to the orientational order.

Looking at the diffusion lengths for the crystalline morphology in figure 7-8, we find that there is an increase in the diffusion length as the order is increased. This may be, in part, explained by a decrease in the average separation of monomers in the crystalline structure compared to the amorphous morphology which is understood to increase the diffusion length ($L_D \propto \sqrt{D} \propto 1/a^2$ [78]). In these simulations the uncertainty (here assuming solely from the sampling of the energy distributions) is initially much higher but decreases as the system size is increased. There is also a much less obvious dependence on r_c, even as low as $r_c = 8$ $\mathring{A}$ which would suggest that this transport is dominated by transfer in the $\pi - \pi$ stacking direction and along the polymer chain, with less contribution from hops in the lamella stacking direction. The values of L_D do begin to stabilise at cut-offs where these hops become accessible however, with all system sizes (with perhaps the exception of the $5 \times 5 \times 5$ unit cell system) producing similar diffusion lengths. For $\kappa^2 = (0.845\sqrt{2/3})^2$, the crystalline morphology produces an L_D of ≈ 11.7 nm and for a fully aligned $\kappa^2 = 4$, $L_D = 33.7$ nm.

These 1D diffusion lengths are again much smaller than Holmes et al., including the fully aligned dipole model which was an initially surprising find. The crystalline system makes up a large portion of the nanoparticles and it was expected that exciton transport in the crystalline domains would be the dominant pathway. With a 0 K crystal

structure, amplifying the transfer rates by assuming all dipoles are fully aligned and assuming the excitons are isolated, this simulation was anticipated to be an overestimate of the true exciton diffusion length. That is is smaller can only mean that there are aspects of exciton transport that are missing from this model and is therefore limiting the calculated exciton diffusion lengths.

One promising candidate is exciton delocalisation. The exciton is not a point-object and in a polymer it is known to delocalise along many segments of the polymer chain [56, 116, 265]. In P3HT, the weak torsional potential breaks the $\pi - \pi$ conjugation and, in some instances, localises the exciton [265, 56]. However, the extent of localisation can still be over many repeat units. Crystalline domains are also expected to increase the delocalisation and exciton lifetime, with the $\pi - \pi$ stacking enabling delocalisation between chains [266]. In this circumstances, the point-dipole model begins to break down and a new model may be required to describe transport of the exciton through the material - both between the chains and along the polymer backbone. Considerations of these effects are beyond the scope of this work, however.

In the diffusion length calculations above there was a clear impact of system size on the uncertainty in L_D. For small systems, the uncertainty was profound, but is drastically narrowed in the case of the large systems. In figure 7-9, we can inspect this more closely. At a cut-off radius $r_c = 30$ $\mathring{A}$, the exciton diffusion length for both the amorphous and crystalline system is shown as a function of system size (in terms of the number of monomers) in panel (a) and the uncertainty for each average (the 100 repeats picking site energies from a Gaussian) against system size in panel (b). Figure 7-9(a) demonstrates this reduction in uncertainty incredibly clearly. In small systems, the exact distribution of site energies can become the dominant contribution to the diffusion length and produces a much wider distribution. This can be rationalised as small correlations in energy having a much larger impact as they cover a large portion of the box, while in large systems, correlations that are large enough to impact the diffusion length become highly improbable [45]. For large systems, the distribution of energies also more reliably resembles a Gaussian and is more resilient to statistical noise [43].

These finite size effects are an important feature of charge transport modelling. It has been suggested that a too small system size means the tail of the Gaussian distribution is not sampled correctly. Since, the low energy end is most populated by charges, missing these tails states mean that charge mobility in a simulation can be artificially enhanced if a system is too small [43]. Similarly it has been shown that if a lattice is too small, then the chance of finding a percolation pathway (a path of very fast transport

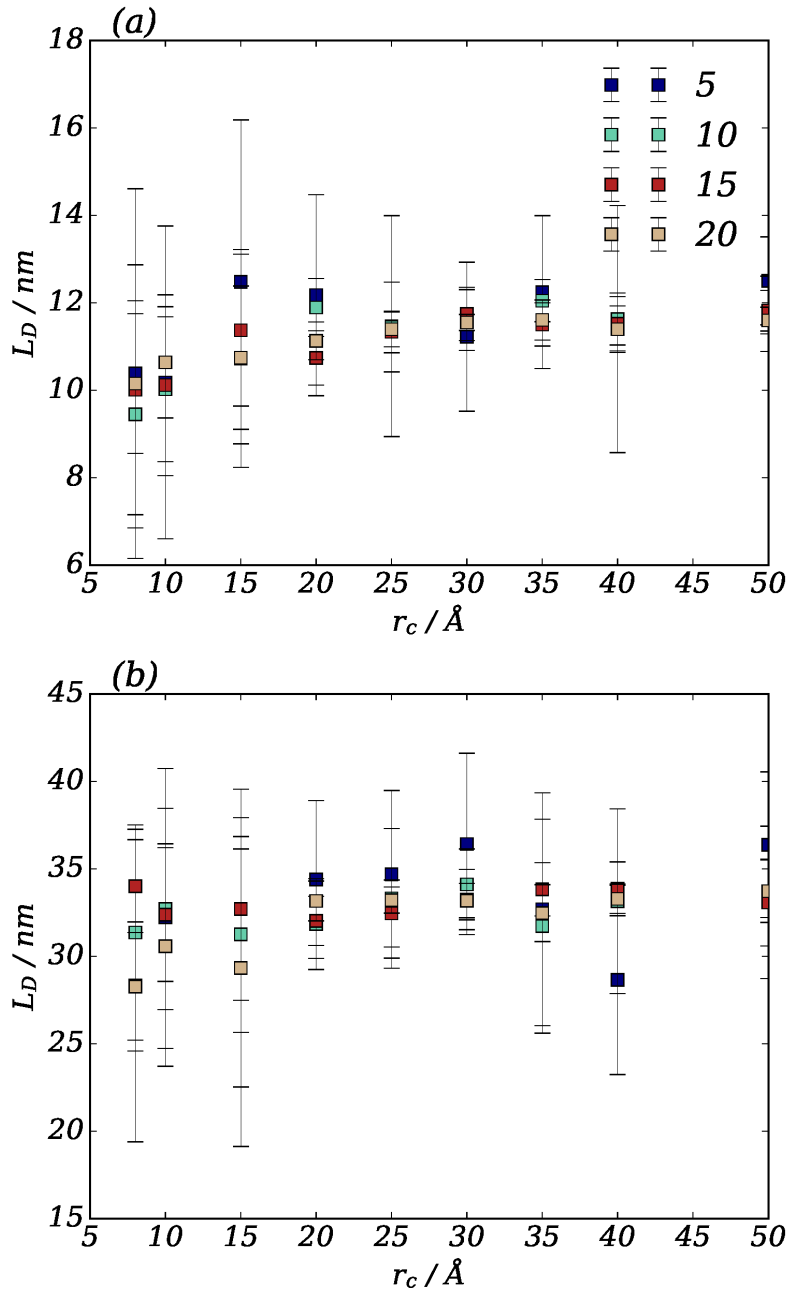

Figure 7-8: Diffusion length for the P3HT crystals as a function of the cut-off radius with (a) a dynamic dipole model ($\kappa^2 = 2/3$) and (b) a static dipole model with all dipoles aligned ($\kappa^2 = 4$). This last case represents a crystal with no thermal structural disorder.

spanning the simulation cell) for a given distribution of energy is much higher and this also inflates reported charge mobilities [45]. In the simulations presented here, it would appear that the crystal is much more resilient to percolation problems (perhaps because Förster transfer is longer range) while in the amorphous system morphology dependence is strong below 10^4 monomers. It would seem that we have reached systems where the uncertainty is small enough and the diffusion length result constant enough that we have overcome much of the finite size effects due to the morphology itself and sampling the distribution of site energies.

In this analysis we have not considered the sampling of morphologies as we did with PCBM. This too would need to be considered and transport simulations performed over a wide range of morphologies, especially in smaller systems as shown in figure 7-9(b) (where the standard deviation is 1/3 of the result itself) and where the importance of doing so is demonstrated in figure 5-16, showing that the distribution of the average diffusion constant for each PCBM morphology is rather broad.

7.4 Summary

In this chapter the structure of the SAMSEN-produced amorphous P3HT morphology, a generated 'ideal' P3HT crystal and an amorphous SAMSEN-generated P3HT nanoparticle morphology were compared to experimental measurements of P3HT nanoparticles suspended in water. Characterised by XRD measurements, the nanoparticles showed a mixed amorphous-crystalline character, the proportions of which were slightly modified by the size of the nanoparticles. Using the simulated structures and performing a computer XRD measurement upon them, an estimate of the crystalline-amorphous ratio was made with a very small decrease in amorphous content for smaller nanoparticles.

The amorphous pure P3HT system and the crystalline P3HT system were then used as the basis for exciton diffusion calculations and provided a comparison to both experimental and computational work on P3HT nanoparticles. Exciton diffusion was significantly increased in the crystal phase and increased even more strongly when the dipole alignment was considered. However, even assuming perfect dipole alignment exciton diffusion length calculated was lower than one which would explain the high exciton dissociation efficiency found in the nanoparticles. We therefore expect that a Förster point-to-point transfer model is inappropriate and other effects such as exciton delocalisation or faster internal transfer mechanisms are present within the P3HT

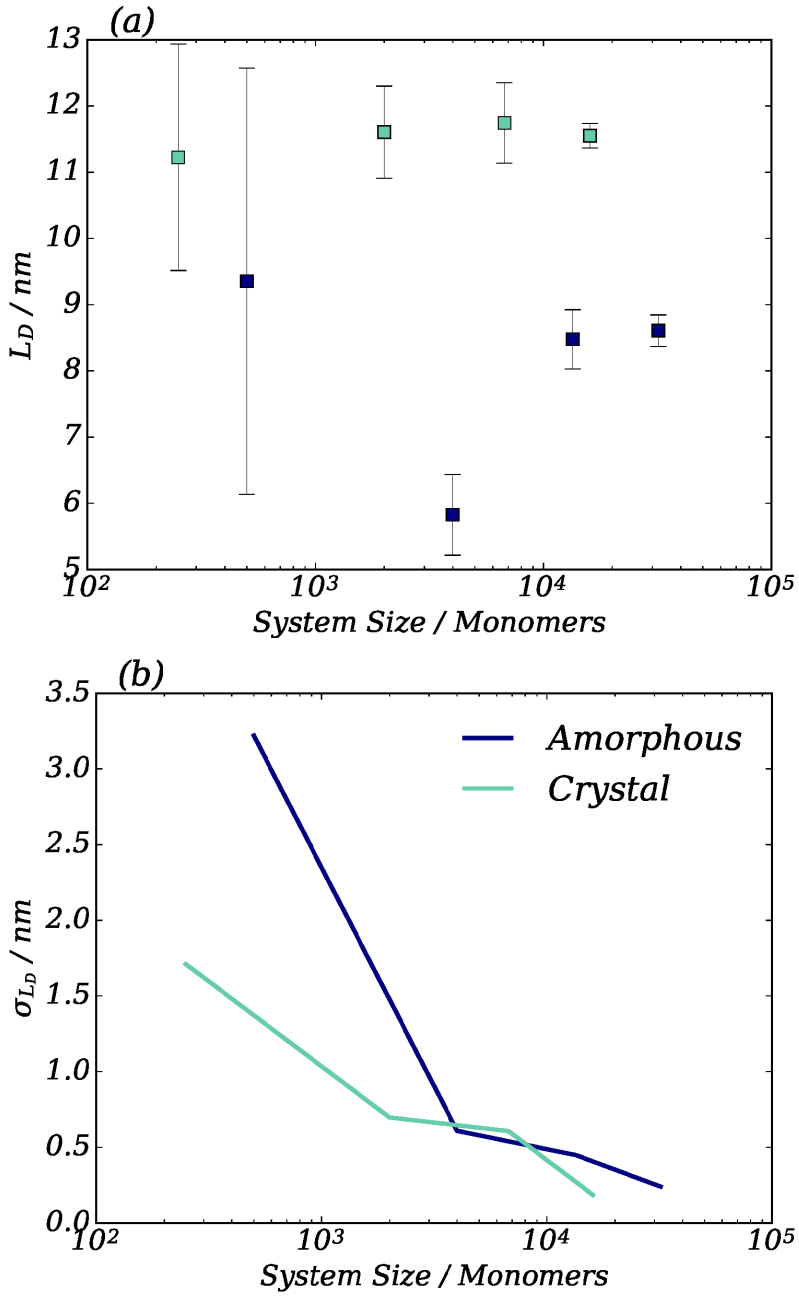

Figure 7-9: (a) the diffusion length and (b) the standard deviation in the diffusion length for excitons in the crystalline and SAMSEN amorphous sytems at a selection of system sizes.

nanoparticles and would need to be considered in further work on exciton transport modelling. This is true of charge transport in polymers in general and also of delocalised charges in small molecules systems too and research on this topic is on-going [9, 26, 11, 56, 265].

Chapter 8

IDT-BT: the 'Disorder-free' Polymer

Indacenodithiophene-co-benzothiadiazole (IDT-BT) is an organic semiconducting polymer which has attracted attention in recent years, in particular as a candidate OFET material. This is, in part, due to a seemingly low energetic disorder which is believed to enhance hole transport through the morphology [106]. This has even been found to be true in the amorphous phase and IDT-BT has therefore been described as a 'disorder-free' polymer [106], although it also forms a well defined crystalline phase [106, 267]. IDT-BT is also an interesting material for environmental reasons as it can be used in production processes with non-halogenated solvents [268].

In this short chapter, the charge transport, specifically the hole diffusion, of an IDT-BT morphology shall be calculated using a SAMSEN morphology and a Marcus theory hopping model between chains. The typical intra-molecular transport rates (within the polymer) shall be estimated by comparison to experimental results.

8.1 Creating Morphologies

The IDT-BT repeat unit was initially computer-generated in Avogadro [225] from a SMILES formula for the molecule, with methyl groups placed at the end of the repeat unit. An energy minimisation was then performed in the General Amber Force-Field using conjugant-gradient minimisation in Avogadro until convergence of the potential energy of the structure. This was repeated with the side chains in varying positions,

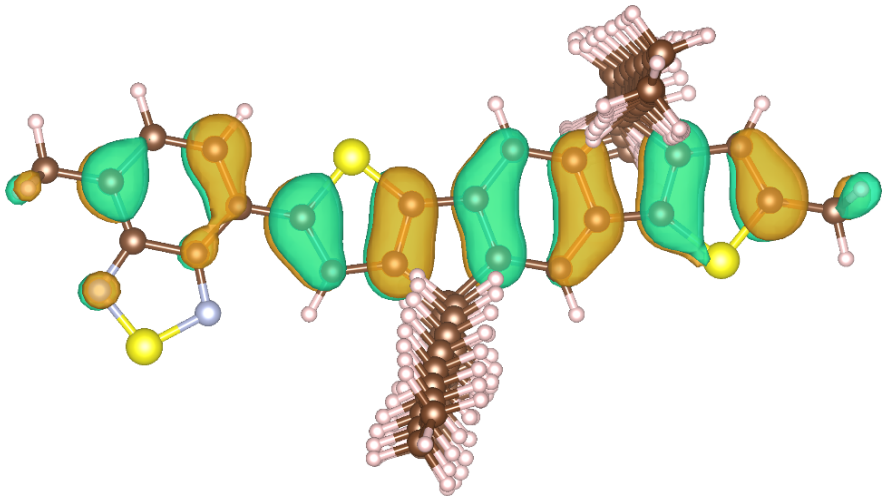

Figure 8-1: The wavefunction of the HOMO in an IDT-BT monomer (terminated with methyl groups) produced in CP2K.

with two minimised structures found. These were then geometry optimised at the density functional theory level in CP2K (performed by Dibyajyoti Ghosh), and two low-energy structures were confirmed to exist. The lowest energy structure was used as the likely optimal configuration for the SAMSEN simulations and is shown in figure 8-1 with the associated HOMO wavefunction. The terminating methyl groups were then removed and replaced by the polymerisation markers. This structure was then converted to rigid section templates using the method in section 3.1.3 and a set of MVEEs created for the repeat unit. The rigid sections for the IDT and BT sections are shown in figure 8-2 with the side chains omitted (the rigid sections of the side chains, being an alkyl chain, are mostly three carbon atoms plus two hydrogens connecting the central atom, with the exception of the ends of the chain). There are a total of 66 rigid sections per repeat unit.

100 IDT-BT chains, each with 10 repeat units in a heal-tail arrangement were placed randomly in a box volume of 400 $\mathring{A} \times 400\ \mathring{A} \times 400\ \mathring{A}$ with periodic boundary conditions. The values of χ_c and χ_f from the previous chapter (0.6 $\mathring{A}$ and 0.3 $\mathring{A}$, respectively) were retained. Compression was performed using a $C_f = 0.99995$ every step. At each step, the rigid sections were given a small random displacement of up to 0.05 $\mathring{A}$ in each of the Cartesian co-ordinates to prevent jamming (practically speaking, the atoms were displaced as a group according to their rigid section membership, taking an average value if they were members of multiple rigid sections). At each step, after compression and after random displacement, the collision-fitting conditions were checked and enforced as described in chapter 3. The compression was continued until the target density of

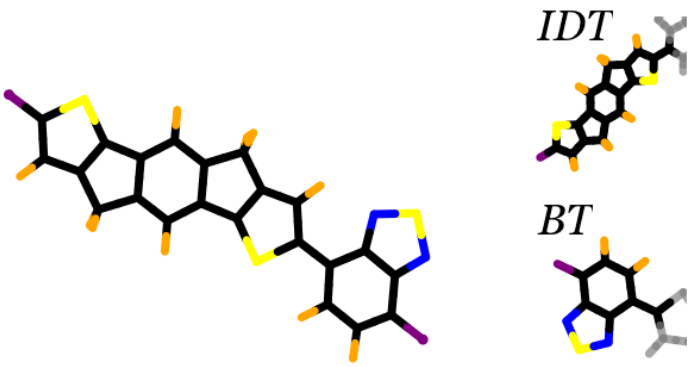

Figure 8-2: The rigid sections for the IDT and BT units in IDT-BT. The purple atoms indicate the polymerisation markers. The side-chains and their 64 rigid sections have been omitted for clarity.

0.84 g/cm^3 was reached. This density was chosen as it was the density determined by molecular dynamics simulations performed by collaborators at the University of Mons. Due to the number of rigid sections in the system, it was once again impractical to meet the memory requirements of the Hessian and calculate the eigenmodes. Random displacements were continued for 10,000 steps in an attempt to obtain a morphology away from the initially compressed state.

Applying the random motion in a similar way to the original FRODA algorithm was an essential step for obtaining the target density. Without it, the polymer chains would be unable to relax back to their template structures while also maintaining the collision conditions. In 21 other attempts at this compression process without random displacement, all failed to reach the target density, producing a density distribution shown in figure 8-3. The polymer repeat units can be considered a pair of long and short rigid rods that can flex by a small angle controlled by χ_f at their connecting points (with 16 additional short rods for each of the the side-chains). Knowing that high-aspect ratio objects have a lower-maxmum packing fraction [250, 251] explains why the density in this system is so low, it may also explain why there is a large variation in densities obtained by SAMSEN. This, along with the lack of dihedral angle control between rigid sections, should be investigated in future work. While, there may indeed be a range of densities on the microscale in the physical IDT-BT films, such a claim about that distribution cannot yet be made here using the SAMSEN structural model alone. If such a range of densities did exist however, it would appear that SAMSEN

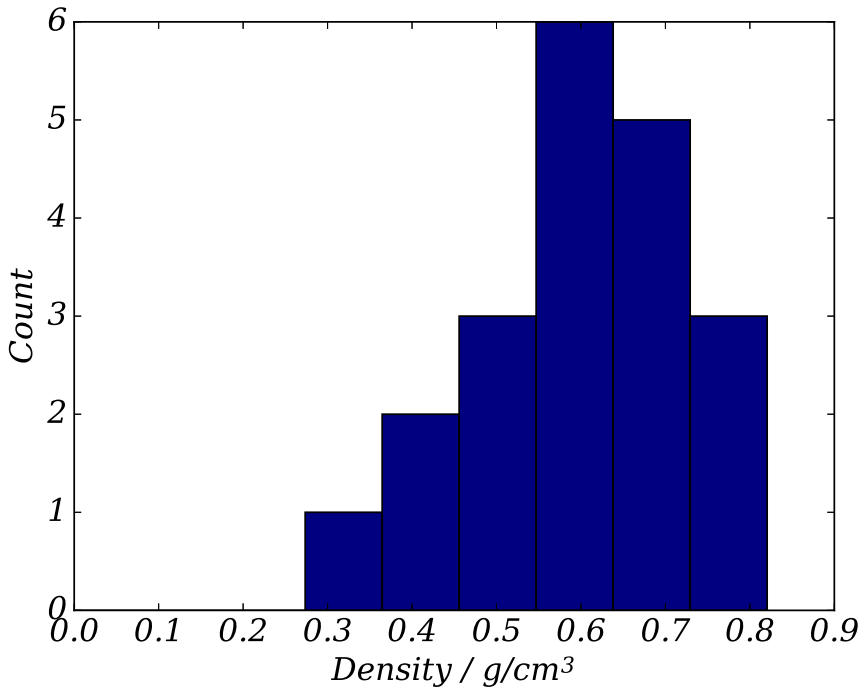

Figure 8-3: Densities at which the collision-mismatch conditions of the SAMSEN structural model could no longer be enforced during 21 compressions of IDT-BT. Only one, not included in this figure but used for the remainder of this chapter, reached the target density of 0.84 g/cm^3.

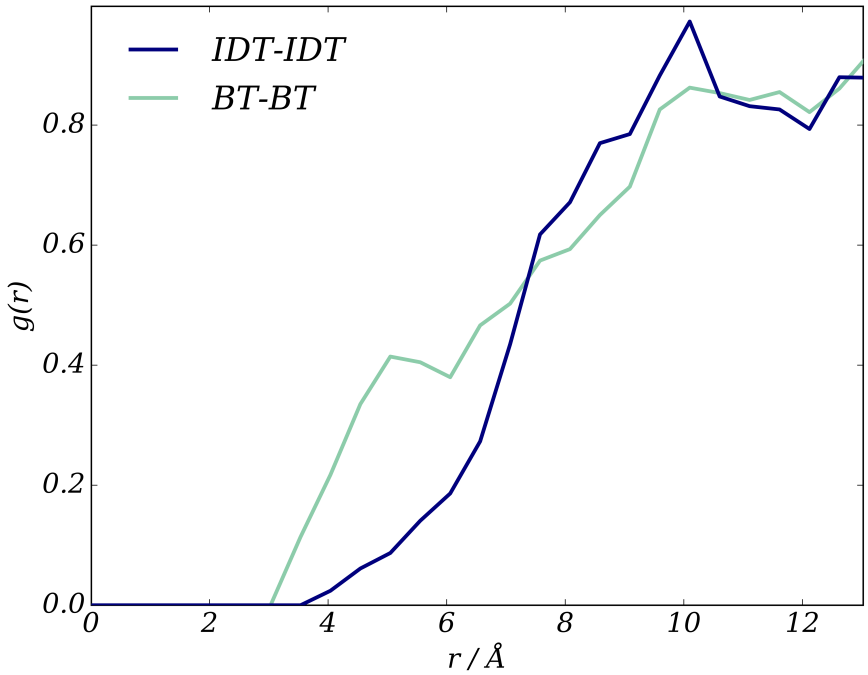

Figure 8-4: The radial distribution function between the geometric centre of the IDT rigid sections and other IDT rigid sections in separate polymer chains and BT rigid sections and BT rigid sections in other polymer chains in the distance range covering the closest approach. These close contact points will come to dominate charge transport between the chains.

once again finds the high-energy states which in this case may represent non-equilibirum states produced by rapid quenching. Assuming the MD density represents a system closer to equilibirum, the results in the next section shall be performed on the 0.84 g/cm^3 system created by adding random motion during compression and continuing for 10,000 steps afterwards.

8.2 Hole Diffusion in IDT-BT

IDT-BT is a hole-transporting polymer and the HOMO for the repeat unit (with methyl termination) is shown in figure 8-1. There is strong delocalisation of the HOMO wavefunction across the entire length of the conjugated regions in both the indacenodithiophene (IDT) and benzothiadiazole (BT) sections. In the polymer, this delocalisation extends across repeat units, in part due to the low structural and energetic disorder of the polymer in the bulk [106, 96]. Hole transport along the polymer backbone is

therefore considered to be very fast and is assumed to dominate and be the cause of the high-hole mobilities in experimental measurements [106].

Between polymer chains, however, charge transport is far more restrictive. Due to the size of the four extended side chains that, in their low-energy configuration at least, span out of the plane of the conjugated region, the distance of closest approach between any two neighbouring repeat units is highly constrained. As can be seen in figure 8-4 which shows the $g(r)$ of external IDT and BT units (only considering neighbours in different polymer chains). This shows that there are very few conjugated sections in different polymer chains in proximity to a given IDT or BT portion of a repeat unit. Even at distances as large as 1.2 nm, there is still (effectively) a lack of neighbouring repeat units for charges to hop between than a measure like the density might suggest. Were a potential of mean-force generated for these interactions, it would suggest that there is a small effective repulsive potential at these distances and imply some exclusion in the local volume. At the shortest distances (below 6 $\mathring{A}$), it would appear that the BT sections are able to approach one another to a much greater extent than the IDT sections which are bonded to the side-chains. I would therefore expect that overlaps in the wavefunction in the BT portion of the repeat unit would dominate hole transport.

8.2.1 Internal and External Transport Rates

To calculate the transport rates between IDT-BT repeat units, a quantum-chemical approach was taken. The IDT-BT morphology was passed to Yoann Oliver at the University of Mons who employed the ADF package [90, 269, 270] using a dimer fragment approach (based upon two repeat units to capture the head-tail conformation). Here, the orbitals of a pair are expressed as a linear combination of the molecular orbitals of the fragments fragment, calculated by solving the Kohn-Sham equations. From this the site energies (in principle) and the transfer integrals can be computed. Here, the transfer integrals for repeat units containing atoms within 6 $\mathring{A}$ of atoms in neighbouring units in different polymer chains were calculated. With this information, as well as an estimation of the reorganisation energy taken from previous work (here, we also use the same estimate of the energetic disorder due to static contributions as the width of the Gaussian) [96], we then have all the information required to calculate the Marcus hopping rate for transfers between polymer chains.

This quantum-chemistry approach, however, does not tell us about the internal transport rate within a polymer chain. Previous reports suggests that this dominates [106], so simply choosing a high internal transport rate (as has been previously suggested

[19]) will not prove a satisfactory estimate that leads us to a valid estimate of a charge transport property within this material. While other methods for calculating the intramolecular transport rates exist (some of which require further DFT calculations and/or parameterisation) [9, 11], determination of the exact mechanism for intra-chain transport is perhaps outside the scope of this work.

Instead an approach was taken where a uniform intra-chain (internal) transport rate was used between nearest neighbour repeat units within the chain. This rate was varied over several orders of magnitude around the average Marcus theory rates for inter-chain (external) hops, allowing the dependence of the diffusion constant on the intramolecular rate to be studied. In this approach, only by comparison to experimental measurements of the charge mobility, could an internal transfer rate be estimated.

8.2.2 Estimating the Internal Transfer Rate

Using the IDT-BT morphology generated in section 8.1 and the associated transfer integrals within the 6 $\mathring{A}$ nearest atom cut-off as determined by colleagues, a series of calculations of the steady-state occupation probabilities of holes and the diffusion constant were performed using the master-equation transfer matrix method described in section 3.4.2. The site energies were approximated by a Gaussian distribution with a width of $\sigma = 33$ meV and a reorganisation energy of $\lambda = 0.25$ eV taken from colleagues' previous work [96, 106]. The sites were defined as the centre-of-mass of the repeat units (when excluding the side-chains) and the hopping was modelled as occurring between these points.

An internal (intra-chain) rate constant for hopping to neighbouring repeat units in the chain, K_I, was chosen and used for 100 sets of site energies taken from the specified Gaussian distribution. The steady-state was calculated by finding the eigenvalue corresponding to an unchanging distribution of probability of occupancy and the diffusion constant calculated. K_I was then doubled and the same calculation performed. This process continued until the timescales of the external and internal hops were too disparate for the probabilistic transfer matrix method to operate - numerically the external rates became negligible and, in that time window of $t = 1$ fs, the probability of transfer reached approximately 1. This complete transfer ($\max(T_{ij}) \approx 1$) simulation and associated D value was removed from the dataset.

The remaining values of the diffusion constant for a K_I in the range of 0.001 and 524 ps^{-1} are shown in figure 8-5. For reference, the average Marcus transfer rate

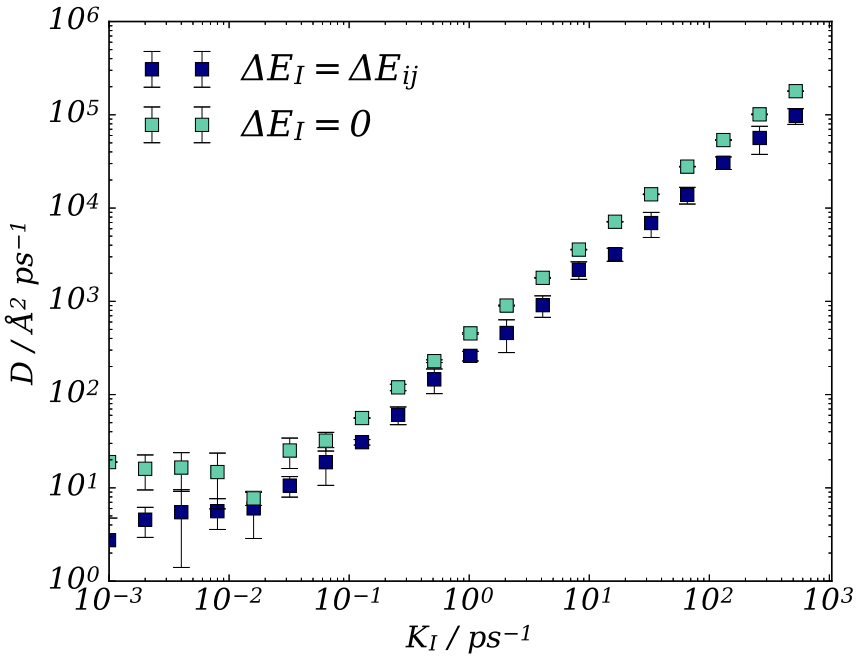

Figure 8-5: Diffusion constant for the IDT-BT system as a function of the internal rate constant for transport between neighbouring repeat units within a polymer chain, K_I. The (external) hops between chains were calculated using quantum chemistry package, ADF, by Joann Olivier at the University of Mons and incorporated the effects of the relative positions and orientations of the monomers in neighbouring chains. The cut-off radius for external hops was 6 Å between any atom in either monomer unit. This calculation is performed for a transport assuming energetic barriers have no effect to intra-chain charge transport and performed again assuming it has an Arrhenius-like effect on the intra-chain transport rates.

between chains was approximately 0.5 ps^{-1}. It is when the K_I reaches a similar value that the diffusion constant becomes dominated by the internal hopping rate. At lower values of K_I, the transport is dominated by external hops with a diffusion constant of approximately 18.89 $\pm$ 0.01 $\mathring{A}^2$ ps^{-1} if one assumes no perturbation to the internal transport rate due to conformational changes altering the site energy, but as low as 2.7 $\pm$ 2.0 $\mathring{A}$ ps^{-1} when an Arrhenius barrier to hopping intra-chain transport is assumed. These two models produces values of D that appear to begin to converge as a low K_I increases towards $K_I \approx \langle K_{Marcus} \rangle$, separate, and hold similar rates of increase towards much greater values of K_I, with the barrier-free transport producing diffusion constants slightly less than double the barrier-hopping transport model.

It is worth once again emphasising that the transport models used here (barrier-free and barrier-hopping hole transport) are not necessarily the correct models. It is likely that they both make approximations which the physical picture does not possess. For instance, it is unlikely that the transport rates are uniform throughout the film and that some energetic disorder [96, 106], bending [106], break in conjugation [56] or nuclear motion [9, 92] may occur and slow hole transport within a single polymer chain. However, the other extreme of assuming charges are localised on a single repeat unit and hop between chains with an Arrhenius-like barrier is probably an underestimate of the physical transfer rates between sections. A semi-band-like or more quantum mechanical approach should be considered due to the low structural disorder potentially enabling greater delocalisation of charges [9, 26, 11, 267]. I present these two results as the likely extreme cases and with an expectation that the physically representative diffusion constant and intra-molecular transport rate lie somewhere in between.

One can, however, estimate the diffusion constant from experimental data, from which one can then estimate the likely range of an effective intra-chain transport rate in a given transport model. The Einstein-Smoluchowski relation provides the ratio between the diffusion constant, D, and the charge mobility, μ and is given as

$$D = \mu k_B T \tag{8.1}$$

where $k_B T$ is the Boltzmann temperature which is $\approx$ 0.026 eV at 300 K. Using the experimental measurement for mobility of 1.5 cm^2 / V s [271, 267], this implies an experimental diffusion constant of $D \approx 390$ $\mathring{A}^2$ ps^{-1}. Looking to figure 8-5, a diffusion constant at this value occurs at a K_I of approximately 0.83 ps^{-1} for the barrier-free transport and 1.4 ps^{-1} for the Arrhenius-like hopping models. Unfortunately, results

for the effective intra-chain transport rate or even the intra-chain charge mobility for IDT-BT do not appear to be published in the literature to the author's knowledge, hindering comparison and evaluation of the approach taken here.

8.3 Summary

In this short chapter, the hole diffusion of IDT-BT was studied. After requiring random displacements to be added to the SAMSEN structural model during compression, a morphology of IDT-BT was created at 0.84 g/cm^3. Calculating the transfer integrals and using a Marcus theory description of transport between chains, the dependence of the material diffusion constant upon the intra-chain transfer rates was assessed in barrier-free and Arrhenius charge-hopping models.

The barrier-free intra-chain transport rate approach produces diffusion constants approximately two times greater than that of the localised hopping model in the regime where intra-chain transport dominates. Comparing to experimental values of mobility and inferring the experimental diffusion constant would suggest that, in an amorphous IDT-BT system, the energetic disorder reduces the rate of intra-chain transport by up to 0.6 ps^{-1} if the physical transport could be described as hopping. It is, however, likely not a pure hopping model with Arrhenius rates that describes the physical reality, nor a fully band-like disorder-free model either. Future work should consider applying quantum mechanical methods or methods such as flexible surface hopping [9] to calculate the intra-chain transport rates in this 'disorder-free' material.

Chapter 9

Conclusion

9.1 SAMSEN

The model presented in chapter 3 and demonstrated with hard spheres in chapter 4 has some clear benefits for improving sampling and some undesirable drawbacks for generating structures of organic semiconductors which should be discussed in closing.

9.1.1 The Structural and Dynamical Models

SAMSEN is a method which is very adept at taking high-density packings or amorphous systems and finding new configurations in different basins in the potential energy landscape. The low-frequency modes are, indeed, highly collective pathways and therefore the structural hindrance to relaxation is limited, enabling computationally-fast sampling of states. Using high-amplitude displacements and low-frequency modes, the dynamical model within SAMSEN is able to propel the nodes in its elastic network far away from their starting point such that, in as little as one or two mode displacements, a structurally independent state can be achieved. After minimisation in a class 1 all-atom force-field such as OPLS [227] or the GAFF [121], a large sample of inherent structures can be found using modest computational resources. In the case of PCBM this was 720 CPU minutes per independent state, with a roughly even split between time spent equilibrating the state in LAMMPS and finding a new state in SAMSEN.

The structural model of SAMSEN also plays a role. It limits the states that are created by the low-frequency modes to those that are not extreme in energy (in a Lennard-Jones potential) by enforcing collisions when atomic spheres reach a maximum

overlap distance, the collision threshold χ_c. This restriction on the closest approach, however, does not appear to interfere with the relaxation time (here the number of cycles before positional decorrelation of molecules in the SAMSEN state occurs). The collision condition, combined with the fitting condition, is able to approximately convey the shape of the molecule and produce structures which are similar to that of the minima in an OPLS force-field (in the case of PCBM). The fitting condition, where no atoms can stray beyond mismatch threshold, χ_f, from their rigid section template is able to constrain the bond lengths and angles within the rigid sections but also between them, such that the molecular structure is maintained. In the case of P3HT, the angular constraints and lack of dihedral angle constraints between repeat units enabled behaviour captured in coarse-grained simulations such as head-tail flipping and the associated bending in the polymer backbone.

Structures

Each of the SAMSEN structures produced have shown reduced structural features compared to those in other published works. It was also common to find the peaks of the $g(r)$ to be at slightly smaller separations than in other data and even when compared against the molecular dynamics simulations performed upon the minimised SAMSEN states (we found this under constant NPT conditions, but later found it under constant NVT conditions too). These structures are therefore only approximate representations of those expected to be produced by a thermodynamic simulation with a class 1 force-field and, instead, higher energy structures are found. Clearly the structural model, which does not take into account the specific atomic environments (which alter forces) and ignores electrostatics and does not sample thermodynamically, has its shortcomings as a model for structure prediction. However, this loss in structure can be regained by performing an MD simulation on the state afterwards. This is not ideal, but the shortened simulation time in terms of maintaining an individual molecular structure and not interfering with the dynamical model's ability to generate a large sample of independent states means this structural model brings great benefit.

Perhaps an interesting feature of the structural model is the proximity of the SAMSEN states to the minimised states. The structural model appears to be able to capture something about the shape of the molecules and, importantly, the excluded volumes of the molecules. This appears to place the structures close to something representative of the real energy landscape. As we saw on a large sample of PCBM states, these configurations appear to be relatively high energy states which were approximately

Gaussian distributed. It may be that, by only placing an upper limit on the energy of overlap and separation and by preventing the sampling of the most extreme states (high overlap, large bond stretches), we are not creating significant thermodynamic bias and the sampling of states appears to reflect something close to the likely density of states of the inherent structures. Of course, without knowing the density of states of the minimised PCBM structures, this assertion can only be made weakly assuming the central limit theorem and the fullerene-portion of the molecular structure dominates the potential energy landscale and the density of states of inherent structures of PCBM resembles the density of states of inherent structures of spheres.

Modes

The vibrational modes calculated using the harmonic approximation assuming the current particle separations are the equilibrium separations is applicable in some circumstances: a dense amorphous packing of spheres considering only physical contacts. However, the highly collective motion that we used here was away from such a regime. By increasing the interaction cut-off distance, the spectrum of vibrational modes changes. In chapter 4 this method revealed a rigidity transition and, at higher cut-offs, a return to fast dynamics and the appearance of bands of vibrational modes. These bands have appeared in all simulations with a sufficiently large value of α.

Studying the collective dynamics of the vibrational modes through measures such as the participation ratio, the average local dot product and the phase quotient, it was established that the lowest frequency modes are the most collective and between bands this collectivity can significantly reduce. Computationally, the fastest structural relaxations occured when following the lowest band of modes and this was identified as the optimal choice for efficiently generating a series of independent structures. Across each range of modes studied, the self-intermediate scattering function followed a one-step exponential decay indicating that they did not suffer from structural arrest which we attributed to the high collectivity of the low frequency modes. This high collectivity and sampling efficiency was found to be true for spheres using radial cut-offs but also for small molecules, with the interactions determined by ellipsoid intercepts, and with small molecules broken into rigid sections, each with their own enclosing ellipsoids.

9.1.2 Spheres, Jamming and the Vibrational Modes

Looking at the approximately hard-sphere glass structures produced in chapter 4, it would appear that the SAMSEN structural model is able to find dense, highly-packed states, in this case, approaching jamming through athermal compression alone. As we saw later when studying small molecules, the SAMSEN structures appear similar to those of quenched states. We then used this fact as our starting point for the study of the vibrations.

Creating an elastic network of relaxed harmonic springs with equal spring constant between neighbouring spheres, the vibrational modes of the sphere system were studied during compression up to the maximum achievable density. During this compression the timescale of the vibrational modes altered significantly, slowing down as density was increased, finding an inflection point at the Maxwell criterion for rigity, and, if α was set to values greater than unity, the dynamics were then found to speed up again at higher densities. At this point, the vibrational modes split into distinct bands. By altering α the occurance of these bands and the packing fraction at which mechanical rigidity was achieved could also be altered. This allows us to both ensure that we can find a region where bands exist by varying α and also to estimate the the maximum packing fraction SAMSEN can achieve (representing the $\alpha \to 1$ limit) which was found to be $\phi = 0.63$ which is on the upper end of the suspected range of Kauzmann densities and just below the jamming density predicted by others.

Another benefit of changing α was that it allowed us to find the more collective modes. The soft-modes around the boson peak were not found to be highly collective and were potentially more localised. By moving to higher α, these soft/excess modes, usually found near mechanical stability, could be avoided and the highly collective modes where the bands emerge could then be used in the dynamical model. The lowest-frequency vibrational band were found to have the most collective mode displacements and, when moving the densely packed spheres along them, were found to produce the fastest structural relaxation. The rate of structural relaxation could also be increased by increasing the maximum amplitude of displacement through ϵ.

Conventiently, changing ω and ϵ was found to have little-to-no effect on the measures of the structure such as the radial distribution function $g(r)$. Conversely, changing the parameters of the structural model, χ_c was also found to have little-to-no effect on the rate at which the dynamical model allowed the system to undergo structural relaxation. The modes were unchanging in frequency as χ_c (and therefore the maximum overlap of the spheres) was modified. During this variation, it was observed that the

distribution of overlaps was Maxwell-Boltzmann distributed and the average overlap and the implied temperature of the system) could be linearly controlled for χ_c in the range of 0.001 to 0.4 times the sphere radius. This separation in structural temperature and dynamical temperature allowed us to exploit the rapid structural relaxation of the low-frequency modes without changing our approximation of the strucutres of spheres, small molecules and polymers.

9.2 Small Molecules

9.2.1 C60, α-NPD and PCBM

Knowing the response of the SAMSEN structural and dynamical models to changing the parameters, ϕ, α, ω, ϵ, χ, we proceeded to apply the model to small systems of semiconducting small molecules. Using C_{60} as a sphere-like case study, we compressed the system and, as the density increased, saw that the mode response was similar to the spheres with a slowing down, an inflection at stability, and an increase in the frequencies at higher densities, with the appearance of bands of modes possessing similar structure. We found computationally quick structural relaxation that could be controlled by ω and ϵ without altering the measured structure of the morphology and again demonstrated the benefit of using the collective modes. The structure of C_{60} states generated was highly amorphous —extirely lacking in the multiple peaks in the $g(r)$ that would be expected for this usually crystalline material.

OLED material, α-NPD, was then studied in three different rigid section groupings designed to allow us to inspect the effect of increasing the internal degrees of freedom and assess the impact of turning off control of the dihedral angles. A range of radial distributions measured between different pairs of points within the molecular structure was produced in 5-7. These showed that dihedral angle control has a strong impact on the structure throughout the system and that, in the case of α-NPD, determination of which sections are rigid requires more than simple bonding information. Each of the structures failed to reproduce certain details and it appears that the 8-body model, the most free (and closest to the method of 3.1.3), with alterations such as the merger of the anime-centres and the napthyl groups, may be the closest representation of the freedoms within α-NPD.

Energetic methods such as FIRST or analysis of force-fields inferred from DFT simulation are methods worth exploring to determine which sections are rigid. It may also

be wise to consider limiting the dihedral motion through dihedral templates. Both of these approaches would allow SAMSEN to better represent the likely states in which one may expect to find α-NPD at a given temperature. However, to create an accurate measure of structure, a thermodynamic bias will need to be introduced into the simulation at some point.

Considering PCBM as a single rigid section, a series of morphologies were created through structural relaxation enabled by the dynamical model and used as inputs for conjugant-gradient minimisation and molecular dynamics simulations in an OPLS all-atom force-field. Comparison was made between all three as well as the work of Cheung & Troisi's (2010) and it was found the MD simulation recreated their structure. The comparison between the SAMSEN structures was more interesting, however. It appeared that the SAMSEN states lacked some of the structure (such as the repulsive region after the nearest neighbour peak in the fullerene-fullerene radial distribution function) and the nearest neighbour separations for the fullerene cages were at higher separations. However, in comparison to the minimised structures, there was very little difference observed in terms of the centre-to-centre and fullerene-to-fullerene $g(r)$. The only change observed during minimisation was due to the single rigid section approximation, which prevented the broadening of the phenyl-phenyl structural features.

9.2.2 Minimisation, Potenial Energies and the Diffusion Constant

The minimisation of the SAMSEN PCBM states in an OPLA-AA force field revealed the similarity in structure of the SAMSEN states and those of the minimised states. However, it also showed that the minimisation did not reproduce the the sharper structural features of the MD simulation, as one might expect, and therefore raised a question about what the distribution of SAMSEN minimised states represented.

Performing a series of simulations from 25 random initial configurations and obtaining states separated by at least one characteristic structural relaxation time, a distribution of 1469 SAMSEN PCBM states were obtained. Minimising these states, a broad distribution of potential minima were obtained (some of which were duplicates). However, performing an MD simulation under NVT conditions and performing a subsequent minimisation showed that the distribution of SAMSEN minima was not the equilibrium distribution and instead represented a higher energy set of minima than those found from the MD simulation. Looking in detail at the structures, the low energy end of the distribution represented larger separations in the nearest neighbour peak than the high energy end of the distribution, explaining why the SAMSEN states did not accurately

reproduce the MD or Cheung & Troisi measurements.

It was noteworthy that the distribution of SAMSEN minima was of near-Gaussian shape as this may represent a distribution with very low thermodynamic bias (the density of states for Lennard-Jones spheres is Gaussian and the central limit theorem plays a significant role). That the MD states (having been created from the SAMSEN minima) managed to recreate the Cheung & Troisi structures may suggest that the distribution is indeed of very low bias or that the low-energy amorphous structures are similar. It may also suggest that the SAMSEN minima were generally very shallow basins, with small barriers to escape into other configurational basins that could be achieved by the short MD simulation. All of these points may suggest that SAMSEN is a useful tool for sampling non-equilibrium states of organic semiconductors or creating an equilibrium distribution of states by performing a short MD simulation on the SAMSEN states.

For those interested in organic semiconductors, studying the charge transport properties of a material is quintessential. Using the large sample of SAMSEN minimised states, as well as the minimsed MD states, the electron diffusion constant was determined for each configuration using a Miller-Abrahams hopping model in a master equation transfer matrix approach. Across all states in the sampled regions of the potential energy landscape, it was found that the diffusion constant was almost entirely unaffected by the potential or the structure, of the minimised states.

This raises some important questions. Are charge transport properties such as the diffusion constant really independent of the potential and small scale changes in separation or is a more quantum-mechanically accurate model required? In this particular case, is the energetic Gaussian disorder model in the charge transport model dominating the results (as it might be expected to [45]) such that the structural variation has little effect or, in a system where electron transport is expected to be semi-band-like has the structural variation been under-represented by the transport model? An investigation into the change in electron delocalisation length or effective mass across the potential energy landscape should be considered where practical. Other approaches such as using quantum-chemistry packages ADF [90] or VOTCA [89] to calculate the LUMO energies, transfer integrals and charge transport simulation, still in a hopping model, may be able to provide additional insight into properties such as the energetic disorder for different states in the energy landscape and higher accuracy diffusion constants. It would also be interesting to see if other properties such as the charge mobility show a different response.

9.3 Polymers

9.3.1 Structure and Vibrational Modes in P3HT-PCBM Blends

Turning our attention to polymers, SAMSEN was applied to systems of P3HT and P3HT-PCBM blends to further test the structural and dynamical model, drawing comparisons to coarse-grained and fully-atomistic molecular dynamics simulations and experimental measurements.

Frustration of P3HT

Morphologies of P3HT were generated using the now-established simulation parameters at seven different mixing ratios, all achieving structural relaxation of the thiophene backbone within one cycle. The structures produced were once again amorphous, showing very few features in structural measurements and showing no sign of clustering or phase separation of the particles. Structural measurements were mostly dominated by structure internal to the polymer and were characterised by short persistence lengths of the polymer chains, with systems having higher PCBM concentrations showing increased persistence of the P3HT backbone and all showing a slight preference for face-to-face alignment of the thiophene ring to the fullerene cages.

Calculating the collective properties of the thiophene units within the vibrational modes, it was found there was no difference in angular or phase collectivity as PCBM mixing ratio was increased and only absolute changes to the mode participation, with the thiophene rigid sections receiving proportionally the same amount of mode displacement. This explains the similarity in relaxation rate (in terms of the number of SAMSEN cycles) between the simulations. The only difference between the vibrational modes were in the frequencies themselves, which shifted to lower frequencies as the concentration of PCBM increased, as was also reported in quasi-elastic neutron scattering experiments [118]. However, seeing no changes to the strength of the alignment of thiophene finds to fullerene cages, the data presented here does not suggest that wrapping around the PCBM causes frustration of the P3HT backbone. Instead these simulations would suggest that an increased persistence and an increased density are the only causes of the change in timescale of the relaxation of the P3HT backbone.

9.3.2 P3HT as Amorphous, Crystalline and Nanoparticle Morphologies

Looking further at the structure of SAMSEN P3HT morphologies, an X-ray diffraction pattern was generated for an amorphous morphology in periodic boundary conditions, an amorphous nanoparticle-like morphology (generated by applying a central compression force without periodic boundaries), and a zero-temperature crystal structure generated using the DFT-optimised structure and XRD measurements of P3HT nanoparticles. Experimental nanoparticle XRD patterns, for nanoparticles of varying size, were measured while the nanoparticles were suspended in water which introduced some error into the series of measurements. The SAMSEN structures were indeed highly disordered, lacking the $\pi - \pi$ and lamella stacking peaks. However, by using a genetic algorithm to find optimal mixing ratios of the amorphous and crystalline structures, an estimate was made of the crystallinity of the P3HT nanoparticles which was in the region of expected values [244]. The size of the P3HT nanoparticles appeared to do little to change the crystallinity in this analysis.

Exciton Diffusion in P3HT

Studying exciton transport in amorphous and crystalline P3HT through a transfer matrix approach and considering the orientation of the dipoles, it would appear that this model does not sufficiently capture the physics within the material and therefore produces too low an exciton diffusion length to explain the high dissociation efficiency in P3HT nanoparticles. The amorphous system produced a diffusion length of 8.8 nm (as calibrated) and a crystalline system with a diffusion length of 11.7 nm in a dynamic dipole approximation and 33.7 nm in a fully alligned static dipole simulation. This is approximately 16.3 nm short of the expected result based upon previous work by Holmes et al. [18]. This would suggest that other behaviour such as faster transport within a polymer chain (this was observed in the Förster model, but it is clearly insufficient) or delocalisation between chains (caused by the $\pi - \pi$ stacking in the crystalline and semi-crystalline phases [266]) should be included in further simulations.

To improve the transport calculations and to try and understand the high diffusion length in the P3HT nanoparticles there are some steps which could be taken in subsequent simulations. The, perhaps most difficult, would be to generate a P3HT structure that represents that of the real nanoparticle. This would mean a system large enough to capture the size of the crystalline domains, with the minimum size inferred from

the XRD data (using the Scherrer equation) and to also capture the semi-crystalline and amorphous domains in between. The morphology should have a crystalline-to-amorphous ratio similar to the one presented here, using paracrystallinity as a proxy or performing structure factor or diffraction pattern calculations to confirm. Producing a system such as this and on this size is only practical in a coarse-grained MD simulation, still requiring access to supercomputer or other high-performance computing facilities. With this representative system (or systems) an exciton transport calculation should then be performed which takes into account the delocalisation of the exciton —a point-to-point resonance transfer model will be insufficient, as demonstrated here for the fully aligned crystal.

9.3.3 IDT-BT

The molecular structure of IDT-BT represents an interesting problem for SAMSEN due to its high dihedral and angular potentials between repeat units and the presence of its extended side chains. It also represents a practical problem, as the number of succesful attempts at reaching the chosen density (achieved by MD) was very low, as was shown in figure 8-3 and instead a distribution of densities was produced. It would be worthwhile remeasuring this density in larger MD simulations using a general force field or determining an experimental density, but if 0.84 g/cm^3 is an accurate value then some modifications to the structural model (already mentioned) may need to be performed so that this density can be reliably achieved.

Remarks on SAMSEN

From inspection of the measurement of collisions and the mismatches, it would appear that the simulation was very much dominated by high mismatches and required a great effort to maintain them. This may suggest that, for this polymer at least, the dihedral forces may need to be represented within the structural model. This could be achieved with a new set of templates that span the bond and second neighbour bonds between the IDT and BT units within and between repeat units. The maximum displacement away from these new templates would need to be determined by picking an energy cut-off and, as before, calculating the maximum displacement that corresponds to a change in dihedral angle. A similar correction/minimisation procedure to the other used in the fitting correction could be applied. A new structural parameter, χ_d, may need to be chosen if the value of χ_f is inappropriate.

The length of the side chains also causes some problems for the dynamical model. Calculation of the vibrational modes in the current implementation uses a sparse matrix representation which is then expanded to a dense representation for the solve. This behaviour is suboptimal as it begins to limit the number of rigid sections that can be included in the elastic network, as discussed in chapter 3. It may be beneficial therefore, to find a new method or a new implementation of the diagonalisation with less memory requirements for calculating the eigenmodes as has been suggested in previous chapters. This will also help extend the length of polymers and increase the system sizes that can be simulated under this method.

Remarks on Charge Transport Modelling

In chapter 8 it was suggested that a characteristic rate constant for charge transport processes within the IDT-BT polymer chains may be approximately 0.83 ps^{-1}, with a larger rate constant (1.4 ps^{-1}) if energetic disorder is assumed to have an Arrhenius effect. However, assuming a localised hopping model within a polymer (and such a low-disorder polymer as IDT-BT at that) may be a poor way of representing transport and charge delocalisation within chains. As we saw with P3HT, this approach also is a poor representation of delocalisation between chains. We also saw in chapter 7 that Förster resonance transfer, even assuming fully aligned dipoles in a crystal, was unable to produce the large exciton diffusion lengths inferred from experimental work and, in chapter 5, it was found that the electron diffusion coefficient was insensitive to changes in PCBM morphologies in both non-equilibirium and low energy states. Throughout this work it would therefore appear that not considering delocalisation and employing the commonly used hopping transport processes under a Gaussian disorder model was an insufficient approach to modelling charge transport.

Instead one should consider using a more detailed approach for charge/exciton transport throughout organic semiconductor materials. At the very least, a Marcus type approach with transfer integrals and site energies calculated from quantum-chemistry calculations should be used where delocalisation is small. In instances where this is still inappropriate, other methods such as the flexible surface hopping are designed to capture the delocalisation of charges within polymers and subsequent transport [9]. Similarly the work of Fornari et al. [11, 87] attempts to capture the effect of the torsional angles on the localisation and movement of charges within chains. It may also be useful to consider the effect of nuclear motion on the charge transport dynamics as was performed by Troisi et al. [26] and in nonadiabatic methods [92].

References

[1] Henning Sirringhaus. Organic field-effect transistors: The path beyond amorphous silicon. *Adv. Mater.*, 26(9):1319–1335, 2014.

[2] Rudolph A. Marcus. Electron transfer reactions in chemistry. theory and experiment. *Rev. Mod. Phys.*, 65:599–610, Jul 1993.

[3] Nobel Prize in Chemistry, 2000.

[4] Ying Diao, Leo Shaw, Zhenan Bao, and Stefan C. B. Mannsfeld. Morphology control strategies for solution-processed organic semiconductor thin films. *Energy Environ. Sci.*, 7:2145–2159, 2014.

[5] Peter K. Watkins, Alison B. Walker, and Geraldine L. B. Verschoor. Dynamical Monte Carlo modelling of organic solar cells: the dependence of internal quantum efficiency on morphology. *Nano Letters*, 5(9):1814–1818, 2005. PMID: 16159229.

[6] J. J.M. Van Der Holst, M. A. Uijttewaal, R. Balasubramanian, R. Coehoorn, P. A. Bobbert, G. A. De Wijs, and R. A. De Groot. Modeling and analysis of the three-dimensional current density in sandwich-type single-carrier devices of disordered organic semiconductors. *Phys. Rev. B - Condensed Matter and Materials Physics*, 79(8):1–11, 2009.

[7] J. Zhou, Y. C. Zhou, J. M. Zhao, C. Q. Wu, X. M. Ding, and X. Y. Hou. Carrier density dependence of mobility in organic solids: A Monte Carlo simulation. *Phys. Rev. B - Condensed Matter and Materials Physics*, 75(15):1–4, 2007.

[8] Chris Groves, Robin G. E. Kimber, and Alison B. Walker. Simulation of loss mechanisms in organic solar cells: A description of the mesoscopic monte carlo technique and an evaluation of the first reaction method. *The J. Chem. Phys.*, 133(14):144110, 2010.

[9] Linjun Wang and David Beljonne. Flexible surface hopping approach to model the crossover from hopping to band-like transport in organic crystals. *The Journal of Physical Chemistry Letters*, 4(11):1888–1894, 2013. PMID: 26283125.

[10] Samuele Giannini, Antoine Carof, and Jochen Blumberger. Crossover from hopping to band-like charge transport in an organic semiconductor model: Atomistic nonadiabatic molecular dynamics simulation. *The Journal of Physical Chemistry Letters*, 9(11):3116–3123, 2018. PMID: 29787275.

[11] Rocco P. Fornari and Alessandro Troisi. Theory of charge hopping along a disordered polymer chain. *Phys. Chem. Chem. Phys.*, 16:9997–10007, 2014.

[12] Simil Thomas, Jack Ly, Lei Zhang, Alejandro L. Briseno, and Jean-Luc Bredas. Improving the stability of organic semiconductors: Distortion energy versus aromaticity in substituted bistetracene. *Chemistry of Materials*, 28(23):8504–8512, 2016.

[13] Naveen Kumar Elumalai and Ashraf Uddin. Open circuit voltage of organic solar cells: an in-depth review. *Energy Environ. Sci.*, 9:391–410, 2016.

[14] Dominik A. Gollmer, Christopher Lorch, Frank Schreiber, Dieter P. Kern, and Monika Fleischer. Enhancing light absorption in organic semiconductor thin films by one-dimensional gold nanowire gratings. *Phys. Rev. Materials*, 1:054602, 2017.

[15] Eva Bundgaard and Frederik C. Krebs. Low band gap polymers for organic photovoltaics. *Solar Energy Materials and Solar Cells*, 91(11):954 – 985, 2007. Low Band Gap Polymer Materials for Organic Solar Cells.

[16] Stephen R. Forrest. Excitons and the lifetime of organic semiconductor devices. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2044):20140320, 2015.

[17] Joseph Kalowekamo and Erin Baker. Estimating the manufacturing cost of purely organic solar cells. *Solar Energy*, 83(8):1224 – 1231, 2009.

[18] Natalie P. Holmes, Melissa Marks, James M. Cave, Krishna Feron, Matthew G. Barr, Adam Fahy, Anirudh Sharma, Xun Pan, David A.L. Kilcoyne, Xiaojing Zhou, David A. Lewis, Mats R. Andersson, Jan Van Stam, Alison B. Walker, Ellen Moons, Warwick J. Belcher, and Paul C. Dastoor. Engineering two-phase and three-phase microstructures from water-based dispersions of nanoparticles for eco-friendly polymer solar cell applications. *Chemistry of Materials*, 30(18):6521–6531, 2018.

[19] P.K. Watkins. *Modelling Charge and Exciton Dynamics in Organic Solar Cells.* (Thesis). University of Bath, 2006.

[20] Robin G. E. Kimber, Alison B. Walker, Gerd E. Schrder-Turk, and Douglas J. Cleaver. Bicontinuous minimal surface nanostructures for polymer blend solar cells. *Phys. Chem. Chem. Phys.*, 12:844–851, 2010.

[21] Mauro Furno, Hans Kleemann, Gregor Schwartz, and Jan Blochwitz-Nimoth. 39.4l: Late-news paper: Vertical organic transistors (V-OFETs) for truly flexible amoled displays. *SID Symposium Digest of Technical Papers*, 46(1):597–600, 2015.

[22] Tamer Dogan, Roy Verbeek, Auke J. Kronemeijer, Peter A. Bobbert, Gerwin H. Gelinck, and Wilfred G. van der Wiel. Short-channel vertical organic field-effect transistors with high on/off ratios. *Advanced Electronic Materials*, 5(5):1900041, 2019.

[23] Ryo Sugano, Tomoya Tashiro, Tomohito Sekine, Hiroyuki Matsui, Daisuke Kumaki, Fabrice Domingues Dos Santos, Atsushi Miyabo, and Shizuo Tokito. Switching time in ferroelectric organic field-effect transistors. *physica status solidi (a)*, 215(11):1701059, 2018.

[24] David L Cheung and Alessandro Troisi. Theoretical study of the organic photovoltaic electron acceptor PCBM: Morphology, electronic structure and charge localisation. *J. Chem. Phys.*, 114:20479–20488, 2010.

[25] Alessandro Troisi, Giorgio Orlandi, and John E. Anthony. Electronic interactions and thermal disorder in molecular crystals containing cofacial pentacene units. *Chemistry of Materials*, 17(20):5024–5031, 2005.

[26] Alessandro Troisi. Charge transport in high mobility molecular semiconductors: classical models and new theories. *Chem. Soc. Rev.*, 40:2347–2358, 2011.

[27] Thomas F. Harrelson, Varuni Dantanarayana, Xiaoyu Xie, Correy Koshnick, Dingqi Nai, Ryan Fair, Sean A. Nuez, Alan K. Thomas, Tucker L. Murrey, Michael A. Hickner, John K. Grey, John E. Anthony, Enrique D. Gomez, Alessandro Troisi, Roland Faller, and Adam J. Moul. Direct probe of the nuclear modes limiting charge mobility in molecular semiconductors. *Mater. Horiz.*, 6:182–191, 2019.

[28] M. Anderson, C. Ramanan, C. Fontanesi, A. Frick, S. Surana, D. Cheyns, M. Furno, T. Keller, S. Allard, U. Scherf, D. Beljonne, G. D'Avino, E. von Hauff,

and E. Da Como. Displacement of polarons by vibrational modes in doped conjugated polymers. *Phys. Rev. Materials*, 1:055604, 2017.

[29] Jing Li, Gabriele D'Avino, Anton Pershin, Denis Jacquemin, Ivan Duchemin, David Beljonne, and Xavier Blase. Correlated electron-hole mechanism for molecular doping in organic semiconductors. *Phys. Rev. Materials*, 1:025602, 2017.

[30] Huipeng Ma, Na Liu, and Jin-Dou Huang. A DFT study on the electronic structures and conducting properties of rubrene and its derivatives in organic field-effect transistors. *Scientific Reports*, 7(1):331, 2017.

[31] J. Henderson, M. Masino, L. E. Hatcher, G. Kociok-Khn, T. Salzillo, A. Brillante, P. R. Raithby, A. Girlando, and E. Da Como. New polymorphs of perylene:tetracyanoquinodimethane charge transfer cocrystals. *Crystal Growth & Design*, 18(4):2003–2009, 2018.

[32] Thomas F. Harrelson, Adam J. Moul, and Roland Faller. Modeling organic electronic materials: bridging length and time scales. *Molecular Simulation*, 43(10-11):730–742, 2017.

[33] Loup Verlet. Computer "experiments" on classical fluids. i. thermodynamical properties of Lennard-Jones molecules. *Phys. Rev.*, 159:98–103, 1967.

[34] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

[35] Matthew L. Jones and Eric Jankowski. Computationally connecting organic photovoltaic performance to atomistic arrangements and bulk morphology. *Molecular Simulation*, 43(10-11):756–773, 2017.

[36] Manuele Lamarra, Luca Muccioli, Silvia Orlandi, and Claudio Zannoni. Temperature dependence of charge mobility in model discotic liquid crystals. *Phys. Chem. Chem. Phys.*, 14(16):5368–5375, 2012.

[37] Riccardo Alessandri, Jaakko J. Uusitalo, Alex H. de Vries, Remco W. A. Havenith, and Siewert J. Marrink. Bulk heterojunction morphologies with atomistic resolution from coarse-grain solvent evaporation simulations. *Journal of the American Chemical Society*, 139(10):3697–3705, 2017. PMID: 28209056.

[38] M. Mesta, C. Schaefer, J. de Groot, J. Cottaar, R. Coehoorn, and P. A. Bobbert. Charge-carrier relaxation in disordered organic semiconductors studied by dark injection: Experiment and modeling. *Phys. Rev. B*, 88:174204, 2013.

[39] M. Mesta, J. Cottaar, R. Coehoorn, and P. A. Bobbert. Study of charge-carrier relaxation in a disordered organic semiconductor by simulating impedance spectroscopy. *Applied Physics Letters*, 104(21):213301, 2014.

[40] Victor Rhle, Alexander Lukyanov, Falk May, Manuel Schrader, Thorsten Vehoff, James Kirkpatrick, Bjrn Baumeier, and Denis Andrienko. Microscopic simulations of charge transport in disordered organic semiconductors. *Journal of Chemical Theory and Computation*, 7(10):3335–3345, 2011. PMID: 22076120.

[41] Feng Gao, Jianpu Wang, James C. Blakesley, Inchan Hwang, Zhe Li, and Neil C. Greenham. Quantifying loss mechanisms in polymer:fullerene photovoltaic devices. *Advanced Energy Materials*, 2(8):956–961, 2012.

[42] Z. G. Yu, D. L. Smith, A. Saxena, R. L. Martin, and A. R. Bishop. Molecular geometry fluctuations and field-dependent mobility in conjugated polymers. *Phys. Rev. B*, 63:085202, 2001.

[43] Pascal Kordt, Thomas Speck, and Denis Andrienko. Finite-size scaling of charge carrier mobility in disordered organic semiconductors. *Phys. Rev. B*, 94:014208, 2016.

[44] R. Coehoorn and P. A. Bobbert. Effects of gaussian disorder on charge carrier transport and recombination in organic semiconductors. *physica status solidi (a)*, 209(12):2354–2377, 2012.

[45] A. Massé, R. Coehoorn, and P. A. Bobbert. Universal size-dependent conductance fluctuations in disordered organic semiconductors. *Phys. Rev. Lett.*, 113:116604, 2014.

[46] Alexander R. Smith, Ian R. Thompson, and Alison B. Walker. Simulating morphologies of organic semiconductors by exploiting low-frequency vibrational modes. *Journal of Chemical Physics*, 150(16):164115, 2019.

[47] Gordon H. Purser. Lewis structures are models for predicting molecular structure, not electronic structure. *Journal of Chemical Education*, 76(7):1013, 1999.

[48] Linus Pauling and G. W. Wheland. The nature of the chemical bond. v. the quantum-mechanical calculation of the resonance energy of benzene and naph-

thalene and the hydrocarbon free radicals. *The Journal of Chemical Physics*, 1(6):362–374, 1933.

[49] Antoine Kahn. Fermi level, work function and vacuum level. *Mater. Horiz.*, 3:7–10, 2016.

[50] Jean-Luc Bredas. Mind the gap! *Mater. Horiz.*, 1:17–19, 2014.

[51] J. Nelson. *The Physics of Solar Cells*. Series on Properties of Semiconductor Materials. Imperial College Press, 2003.

[52] Alessandro Troisi and Giorgio Orlandi. Charge-transport regime of crystalline organic semiconductors: Diffusion limited by thermal off-diagonal electronic disorder. *Phys. Rev. Lett.*, 96:086601, 2006.

[53] I. Yavuz. Dichotomy between the band and hopping transport in organic crystals: insights from experiments. *Phys. Chem. Chem. Phys.*, 19:25819–25828, 2017.

[54] Gregory D. Scholes and Garry Rumbles. Excitons in nanoscale systems. *Nature Materials*, 5(9):683–696, 2006.

[55] Gregory D. Scholes and Cathal Smyth. Perspective: Detecting and measuring exciton delocalization in photosynthetic light harvesting. *The Journal of Chemical Physics*, 140(11):110901, 2014.

[56] William Barford, David G. Lidzey, Dmitry V. Makhov, and Anthony J. H. Meijer. Exciton localization in disordered poly(3-hexylthiophene). *The Journal of Chemical Physics*, 133(4):044504, 2010.

[57] T. T. To and S. Adams. Modelling of P3HT:PCBM interface using coarse-grained forcefield derived from accurate atomistic forcefield. *Phys. Chem. Chem. Phys.*, 16(10):4653–4663, 2014.

[58] Arne C. Morteani, Paiboon Sreearunothai, Laura M. Herz, Richard H. Friend, and Carlos Silva. Exciton regeneration at polymeric semiconductor heterojunctions. *Phys. Rev. Lett.*, 92:247402, 2004.

[59] Stavros Athanasopoulos, Emmanuelle Hennebicq, David Beljonne, and Alison B. Walker. Trap limited exciton transport in conjugated polymers. *The Journal of Physical Chemistry C*, 112(30):11532–11538, 2008.

[60] Anna Khler and Heinz Bssler. What controls triplet exciton transfer in organic semiconductors? *J. Mater. Chem.*, 21:4003–4011, 2011.

[61] Florian Schindler, Josemon Jacob, Andrew C. Grimsdale, Ullrich Scherf, Klaus Mllen, John M. Lupton, and Jochen Feldmann. Counting chromophores in conjugated polymers. *Angewandte Chemie International Edition*, 44(10):1520–1525, 2005.

[62] Shuji Abe, J. Yu, and W. P. Su. Singlet and triplet excitons in conjugated polymers. *Phys. Rev. B*, 45:8264–8271, 1992.

[63] M Wohlgenannt and Z V Vardeny. Spin-dependent exciton formation rates in -conjugated materials. *Journal of Physics: Condensed Matter*, 15(3):R83–R107, 2003.

[64] M. Wohlgenannt, Kunj Tandon, S. Mazumdar, S. Ramasesha, and Z. V. Vardeny. Formation cross-sections of singlet and triplet excitons in -conjugated polymers. *Nature*, 409(6819):494–497, 2001.

[65] Z. Shuai, D. Beljonne, R. J. Silbey, and J. L. Brédas. Singlet and triplet exciton formation rates in conjugated polymer light-emitting diodes. *Phys. Rev. Lett.*, 84:131–134, 2000.

[66] Yong Cao, Ian D. Parker, Gang Yu, Chi Zhang, and Alan J. Heeger. Improved quantum efficiency for electroluminescence in semiconducting polymers. *Nature*, 397(6718):414–417, 1999.

[67] D. L. Dexter. A theory of sensitized luminescence in solids. *The Journal of Chemical Physics*, 21(5):836–850, 1953.

[68] J. J. M. Halls, J. Cornil, D. A. dos Santos, R. Silbey, D.-H. Hwang, A. B. Holmes, J. L. Brédas, and R. H. Friend. Charge- and energy-transfer processes at polymer/polymer interfaces: A joint experimental and theoretical study. *Phys. Rev. B*, 60:5721–5727, 1999.

[69] R. A. Marcus. Chemical and electrochemical electron-transfer theory. *Annual Review of Physical Chemistry*, 15(1):155–196, 1964.

[70] Jenny Nelson, James Kirkpatrick, and P. Ravirajan. Factors limiting the efficiency of molecular photovoltaic devices. *Phys. Rev. B*, 69:035337, 2004.

[71] Sven Stafstrm. Electron localization and the transition from adiabatic to nonadiabatic charge transport in organic conductors. *Chem. Soc. Rev.*, 39:2484–2499, 2010.

[72] Allen Miller and Elihu Abrahams. Impurity conduction at low concentrations. *Phys. Rev.*, 120:745–755, Nov 1960.

[73] H. van Eersel. *Device Physics of organic light emitting diodes: interplay between charges and excitons.* (Thesis). TUE, Eindhoven University of Technology, 2015.

[74] Stavros Athanasopoulos, Steffen Tscheuschner, Heinz Bssler, and Anna Khler. Efficient charge separation of cold charge-transfer states in organic solar cells through incoherent hopping. *The Journal of Physical Chemistry Letters*, 8(9):2093–2098, 2017. PMID: 28436660.

[75] W. F. Pasveer, J. Cottaar, C. Tanase, R. Coehoorn, P. A. Bobbert, P. W. M. Blom, D. M. de Leeuw, and M. A. J. Michels. Unified description of charge-carrier mobilities in disordered semiconducting polymers. *Phys. Rev. Lett.*, 94:206601, 2005.

[76] J. Cottaar, L. J. A. Koster, R. Coehoorn, and P. A. Bobbert. Scaling theory for percolative charge transport in disordered molecular semiconductors. *Phys. Rev. Lett.*, 107:136601, 2011.

[77] Robin G. E. Kimber, Edward N. Wright, Simon E. J. O'Kane, Alison B. Walker, and James C. Blakesley. Mesoscopic kinetic monte carlo modeling of organic photovoltaic device characteristics. *Phys. Rev. B*, 86:235206, Dec 2012.

[78] Th Forster. Energiewanderung und Fluoreszenz. *Die Naturwissenschaften*, 33(6):166–175, 1948.

[79] Oleksandr V. Mikhnenko, Paul W. M. Blom, and Thuc-Quyen Nguyen. Exciton diffusion in organic semiconductors. *Energy Environ. Sci.*, 8:1867–1888, 2015.

[80] Krishna Feron, James M. Cave, Mahir N. Thameel, Connor OSullivan, Renee Kroon, Mats R. Andersson, Xiaojing Zhou, Christopher J. Fell, Warwick J. Belcher, Alison B. Walker, and Paul C. Dastoor. Utilizing energy transfer in binary and ternary bulk heterojunction organic solar cells. *ACS Applied Materials & Interfaces*, 8(32):20928–20937, 2016. PMID: 27456294.

[81] Steven S. Vogel, Tuan A. Nguyen, B. Wieb van der Meer, and Paul S. Blank. The impact of heterogeneity and dark acceptor states on fret: Implications for using fluorescent protein donors and acceptors. *PLOS ONE*, 7(11):1–14, 2012.

[82] Maria Khrenova, Igor Topol, Jack Collins, and Alexander Nemukhin. Estimating orientation factors in the fret theory of fluorescent proteins: The tagrfp-kfp pair and beyond. *Biophysical Journal*, 108(1):126 – 132, 2015.

[83] J. Knoester and J. E. Van Himbergen. Effect of molecular reorientation on excitation decay due to incoherent energy transfer. *The Journal of Chemical Physics*, 81(10):4380–4388, 1984.

[84] L. M. Herz, C. Silva, R. H. Friend, R. T. Phillips, S. Setayesh, S. Becker, D. Marsitsky, and K. Müllen. Effects of aggregation on the excitation transfer in perylene-end-capped polyindenofluorene studied by time-resolved photoluminescence spectroscopy. *Phys. Rev. B*, 64:195203, 2001.

[85] Hannah J. Eggimann, Florian Le Roux, and Laura M. Herz. How -phase content moderates chain conjugation and energy transfer in polyfluorene films. *The Journal of Physical Chemistry Letters*, 10(8):1729–1736, 2019. PMID: 30900449.

[86] E.N. Wright. *Modelling Charge and Energy Transport in Organic Devices*. (Thesis) University of Bath, 2014.

[87] Rocco P. Fornari and Alessandro Troisi. Narrower bands with better charge transport: The counterintuitive behavior of semiconducting copolymers. *Advanced Materials*, 26(45):7627–7631, 2014.

[88] Haoyuan Li and Jean-Luc Bredas. Modeling of actual-size organic electronic devices from efficient molecular-scale simulations. *Advanced Functional Materials*, 28(29):1801460, 2018.

[89] Victor Rühle, Christoph Junghans, Alexander Lukyanov, Kurt Kremer, and Denis Andrienko. Versatile Object-Oriented toolkit for Coarse-Graining applications. *Journal of Chemical Theory and Computation*, 5(12):3211–3223, 2009.

[90] G. te Velde, F. M. Bickelhaupt, E. J. Baerends, C. Fonseca Guerra, S. J. A. van Gisbergen, J. G. Snijders, and T. Ziegler. Chemistry with adf. *J. Comput. Chem.*, 22(9):931–967, 2001.

[91] Veaceslav Coropceanu, Jrme Cornil, Demetrio A. da Silva Filho, Yoann Olivier, Robert Silbey, and Jean-Luc Brdas. Charge transport in organic semiconductors. *Chemical Reviews*, 107(4):926–952, 2007. PMID: 17378615.

[92] Timothy Clark. Simulating charge transport in flexible systems. *Perspectives in Science*, 6:58 – 65, 2015. Proceedings of the Beilstein Bozen Symposium 2014 Chemistry and Time.

[93] Reinder Coehoorn, Harm van Eersel, Peter Bobbert, and Ren Janssen. Kinetic monte carlo study of the sensitivity of oled efficiency and lifetime to materials parameters. *Advanced Functional Materials*, 25(13):2024–2037, 2015.

[94] K. Feron, X. Zhou, W. J. Belcher, and P. C. Dastoor. Exciton transport in organic semiconductors: Frster resonance energy transfer compared with a simple random walk. *Journal of Applied Physics*, 111(4):044510, 2012.

[95] H. Bssler. Charge transport in disordered organic photoconductors a monte carlo simulation study. *physica status solidi (b)*, 175(1):15–56, 1993.

[96] Tudor H. Thomas, David J. Harkin, Alexander J. Gillett, Vincent Lemaur, Mark Nikolka, Aditya Sadhanala, Johannes M. Richter, John Armitage, Hu Chen, Iain McCulloch, S. Matthew Menke, Yoann Olivier, David Beljonne, and Henning Sirringhaus. Short contacts between chains enhancing luminescence quantum yields and carrier mobilities in conjugated copolymers. *Nature Communications*, 10(1):2614, 2019.

[97] J. J.M. Van Der Holst, F. W.A. Van Oost, R. Coehoorn, and P. A. Bobbert. Electron-hole recombination in disordered organic semiconductors: Validity of the Langevin formula. *Phys. Rev. B - Condensed Matter and Materials Physics*, 80(23):1–8, 2009.

[98] Yu. N. Gartstein and E. M. Conwell. High-field hopping mobility in molecular systems with spatially correlated energetic disorder. *Chemical Physics Letters*, 245(4):351–358, 1995.

[99] M. Bouhassoune, S. L. M. van Mensfoort, P. A. Bobbert, and R. Coehoorn. Carrier-density and field-dependent charge-carrier mobility in organic semiconductors with correlated gaussian disorder. *Organic Electronics*, 10(3):437–445, 2009.

[100] Yifan Zheng, Jaemin Kong, Di Huang, Wei Shi, Lyndsey McMillon-Brown, Howard E. Katz, Junsheng Yu, and Andr D. Taylor. Spray coating of the PCBM electron transport layer significantly improves the efficiency of p-i-n planar perovskite solar cells. *Nanoscale*, 10:11342–11348, 2018.

[101] Tracey M. Clarke, Amy M. Ballantyne, Jenny Nelson, Donal D. C. Bradley, and James R. Durrant. Free energy control of charge photogeneration in polythiophene/fullerene solar cells: The influence of thermal annealing on P3HT/PCBM blends. *Advanced Functional Materials*, 18(24):4029–4035, 2008.

[102] Tom J. Savenije, Jessica E. Kroeze, Xiaoniu Yang, and Joachim Loos. The formation of crystalline P3HT fibrils upon annealing of a PCBM:P3HT bulk heterojunction. *Thin Solid Films*, 511-512:2 – 6, 2006. EMSR 2005 - Proceedings of Symposium F on Thin Film and Nanostructured Materials for Photovoltaics.

[103] C. A. Otlora, A. F. Loaiza, and G. Gordillo. Influence of solvent on the molecular ordering of thin films of P3HT:PCBM blends and precursor solution. In *2014 IEEE 40th Photovoltaic Specialist Conference (PVSC)*, pages 1754–1757, 2014.

[104] Y.-M. Lu, C.-H. Chiang, and S. Lien-Chung Hsu. The performance of polymer solar cells based on P3HT:PCBM after post-annealing and adding titanium dioxide nanoparticles. *Materials Research Innovations*, 18(sup3):S3–102–S3–105, 2014.

[105] David E. Motaung, Gerald F. Malgas, Steven S. Nkosi, Gugu H. Mhlongo, Bonex W. Mwakikunga, Thomas Malwela, Christopher J. Arendse, Theophillus F. G. Muller, and Franscious R. Cummings. Comparative study: the effect of annealing conditions on the properties of P3HT:PCBM blends. *Journal of Materials Science*, 48(4):1763–1778, 2013.

[106] Deepak Venkateshvaran, Mark Nikolka, Aditya Sadhanala, Vincent Lemaur, Mateusz Zelazny, Michal Kepa, Michael Hurhangee, Auke Jisk Kronemeijer, Vincenzo Pecunia, Iyad Nasrallah, Igor Romanov, Katharina Broch, Iain McCulloch, David Emin, Yoann Olivier, Jerome Cornil, David Beljonne, and Henning Sirringhaus. Approaching disorder-free transport in high-mobility conjugated polymers. *Nature*, 515:384–388, 2014.

[107] M. R. Stalker, J. Grant, C. W. Yong, L. A. Ohene-Yeboah, T. J. Mays, and S. C. Parker. Molecular simulation of hydrogen storage and transport in cellulose. *Molecular Simulation*, 2019.

[108] Daniel J. Cosgrove. Growth of the plant cell wall. *Nature Reviews Molecular Cell Biology*, 6(11):850–861, 2005.

[109] David M. Huang, Roland Faller, Khanh Do, and Adam J. Moule. Coarse-grained computer simulations of polymer/fullerene bulk heterojunctions for organic photovoltaic applications. *J. Chem. Theory Comput.*, 6(2):526–537, 2010.

[110] Rodrigo Noriega, Jonathan Rivnay, Koen Vandewal, Felix P V Koch, Natalie Stingelin, Paul Smith, Michael F. Toney, and Alberto Salleo. A general relationship between disorder, aggregation and charge transport in conjugated polymers. *Nat. Mater.*, 12(11):1038–1044, 2013.

[111] Ilhan Yavuz, Blanton N. Martin, Jiyong Park, and K. N. Houk. Theoretical study of the molecular ordering, paracrystallinity, and charge mobilities of oligomers in different crystalline phases. *Journal of the American Chemical Society*, 137(8):2856–2866, 2015. PMID: 25658235.

[112] A M Hindeleh and R Hosemann. Paracrystals representing the physical state of matter. *Journal of Physics C: Solid State Physics*, 21(23):4155–4170, 1988.

[113] M. L. Jones, D. M. Huang, B. Chakrabarti, and Chris Groves. Relating molecular morphology to charge mobility in semicrystalline conjugated polymers. *J. Phys. Chem. C*, 120(8):4240–4250, 2016.

[114] Samuel E. Root, Nicholas E. Jackson, Suchol Savagatrup, Gaurav Arya, and Darren J. Lipomi. Modelling the morphology and thermomechanical behaviour of low-bandgap conjugated polymers and bulk heterojunction films. *Energy Environ. Sci.*, 10(2):558–569, 2017.

[115] Stavros Athanasopoulos, Evguenia V. Emelianova, Alison B. Walker, and David Beljonne. Exciton diffusion in energetically disordered organic materials. *Phys. Rev. B*, 80:195209, 2009.

[116] Matthew L. Jones. *Examining the Links Between Organic Photovoltaic Operation and Complex Morphological Structures*. Durham University, 2015.

[117] Ian R. Thompson, Mary K. Coe, Alison B. Walker, Matteo Ricci, Otello M. Roscioni, and Claudio Zannoni. Microscopic origins of charge transport in triphenylene systems. *Phys. Rev. Materials*, 2:064601, 2018.

[118] Anne A.Y. Guilbert, Mohamed Zbiri, Alan D.F. Dunbar, and Jenny Nelson. Quantitative Analysis of the Molecular Dynamics of P3HT:PCBM Bulk Heterojunction. *J. Phys. Chem. B*, 121(38):9073–9080, 2017.

[119] Pascal Friederich, Vadim Rodin, Florian Von Wrochem, and Wolfgang Wenzel. Built-in potentials induced by molecular order in amorphous organic thin films. *ACS Appl. Mater. Interfaces*, 10(2):1881–1887, 2018.

[120] Tobias Neumann, Denis Danilov, Christian Lennartz, and Wolfgang Wenzel. Modeling disordered morphologies in organic semiconductors. *J. Comput. Chem.*, 34(31):2716–2725, 2013.

[121] Junmei Wang, Romain M. Wolf, James W. Caldwell, Peter A. Kollman, and David A. Case. Development and testing of a general amber force field. *Journal of Computational Chemistry*, 25(9):1157–1174, 2004.

[122] Naga Rajesh Tummala, Christopher Bruner, Chad Risko, Jean Luc Brédas, and Reinhold H. Dauskardt. Molecular-scale understanding of cohesion and fracture in P3HT:Fullerene blends. *ACS Appl. Mater. Interfaces*, 7(18):9957–9964, 2015.

[123] Monika Williams, Naga Rajesh Tummala, Saadullah G. Aziz, Chad Risko, and Jean Luc Brédas. Influence of molecular shape on solid-state packing in disordered PC61BM and PC71BM fullerenes. *J. Phys. Chem. Letters*, 5(19):3427–3433, 2014.

[124] Roland Faller. Automatic coarse graining of polymers. *Polymer*, 45(11):3869 – 3876, 2004.

[125] A.R. Leach. *Molecular Modelling: Principles and Applications*. Prentice Hall, 2001.

[126] Robert E. Rudd and Jeremy Q. Broughton. Coarse-grained molecular dynamics and the atomic limit of finite elements. *Phys. Rev. B*, 58:R5893–R5896, 1998.

[127] John Tinsley Oden, Kathryn Farrell, and Danial Faghihi. Estimation of error in observables of coarse-grained models of atomic systems. *Advanced Modeling and Simulation in Engineering Sciences*, 2(1):5, 2015.

[128] Frank H. Stillinger and Thomas A. Weber. Inherent structure theory of liquids in the hardsphere limit. *The Journal of Chemical Physics*, 83(9):4767–4775, 1985.

[129] D. Frenkel and B. Smit. *Understanding Molecular Simulation: From Algorithms to Applications*. Academic Press, 1996.

[130] Pablo G. Debenedetti and Frank H. Stillinger. Supercooled liquids and the glass transition. *Nature*, 410(6825):259–267, 2001.

[131] Kurt Binder and Walter Kob. *Glassy Materials and Disordered Solids*. World Scientific, revised edition, 2011.

[132] C. S. O'Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel. Jamming at zero temperature and zero applied stress: the epitome of disorder. *Phys. Rev. E*, 68:1–19, 2003.

[133] Jun Zhao, Ann Swinnen, Guy Van Assche, Jean Manca, Dirk Vanderzande, and Bruno Van Mele. Phase diagram of P3HT:PCBM blends and its implication for the stability of morphology. *J. Phys. Chem. A*, 113:1587–1591, 2009.

[134] Trinh Tung Ngo, Duc Nghia Nguyen, and Van Tuyen Nguyen. Glass transition of PCBM, P3HT and their blends in quenched state. *Adv. Nat. Sci: Nanosci. Nanotechnol.*, 3(4):045001, 2012.

[135] Christian Mller. On the glass transition of polymer semiconductors and its impact on polymer solar cell stability. *Chemistry of Materials*, 27(8):2740–2754, 2015.

[136] Diane M. Walters, Lucas Antony, Juan J. de Pablo, and M. D. Ediger. Influence of molecular shape on the thermal stability and molecular orientation of vapor-deposited organic semiconductors. *The Journal of Physical Chemistry Letters*, 8(14):3380–3386, 2017. PMID: 28677392.

[137] Shakeel S. Dalal, Diane M. Walters, Ivan Lyubimov, Juan J. de Pablo, and M. D. Ediger. Tunable molecular orientation and elevated thermal stability of vapor-deposited organic semiconductors. *Proceedings of the National Academy of Sciences*, 112(14):4227–4232, 2015.

[138] Frank H. Stillinger and Thomas A. Weber. Packing structures and transitions in liquids and solids. *Science*, 225(4666):983–989, 1984.

[139] Frank H. Stillinger and Thomas A. Weber. Hidden structure in liquids. *Phys. Rev. A*, 25(2):978–989, 1982.

[140] Frank H. Stillinger. A topographic view of supercooled liquids and glass formation. *Science*, 267(5206):1935–1939, 1995.

[141] Sai Manoj Gali, Gabriele DAvino, Philippe Aurel, Guangchao Han, Yuanping Yi, Theodoros A. Papadopoulos, Veaceslav Coropceanu, Jean-Luc Brdas, Georges Hadziioannou, Claudio Zannoni, and Luca Muccioli. Energetic fluctuations in amorphous semiconducting polymers: Impact on charge-carrier mobility. *The Journal of Chemical Physics*, 147(13):134904, 2017.

[142] Stefan Goedecker. Minima hopping: An efficient search method for the global minimum of the potential energy surface of complex molecular systems. *The Journal of Chemical Physics*, 120(21):9911–9917, 2004.

[143] Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 99(20):12562–12566, 2002.

[144] Andrea Ninarello, Ludovic Berthier, and Daniele Coslovich. Models and algorithms for the next generation of glass transition studies. *Phys. Rev. X*, 7:021039, 2017.

[145] Tomás S. Grigera and Giorgio Parisi. Fast monte carlo algorithm for supercooled soft spheres. *Phys. Rev. E*, 63:045102, 2001.

[146] B Brooks and M Karplus. Harmonic dynamics of proteins: normal modes and fluctuations in bovine pancreatic trypsin inhibitor. *Proc. Natl. Acad. Sci. U.S.A.*, 80(21):6571–6575, 1983.

[147] A. R. Atilgan, S. R. Durell, R. L. Jernigan, M. C. Demirel, O. Keskin, and I. Bahar. Anisotropy of fluctuation dynamics of proteins with an elastic network model. *Biophys. J.*, 80(1):505–515, 2001.

[148] Stephen Wells, Scott Menor, Brandon Hespenheide, and M. F. Thorpe. Constrained geometric simulation of diffusive motion in proteins. *Phys. Biol.*, 2(4):S127–S136, 2005.

[149] C.I. Brändén and J. Tooze. *Introduction to Protein Structure*. Garland Pub., 1999.

[150] A.R. Rees. *The Antibody Molecule: From Antitoxins to Therapeutic Antibodies*. Oxford medical histories. Oxford University Press, 2015.

[151] Ivet Bahar, Ali Rana Atilgan, Melik C. Demirel, and Burak Erman. Vibrational dynamics of folded proteins: Significance of slow and fast motions in relation to function and stability. *Phys. Rev. Lett.*, 80(12):2733–2736, 1998.

[152] Hao Hu, Michael W. Clarkson, Jan Hermans, and Andrew L. Lee. Increased rigidity of eglin c at acidic ph: evidence from NMR spin relaxation and MD simulations. *Biochemistry*, 42(47):13856–13868, 2003. PMID: 14636053.

[153] Emma Langella, Roberto Improta, and Vincenzo Barone. Checking the ph-induced conformational transition of prion protein by molecular dynamics simula-

tions: Effect of protonation of histidine residues. *Biophysical Journal*, 87(6):3623 – 3632, 2004.

[154] Lennart Nilsson and Andrey Karshikoff. Multiple pH regime molecular dynamics simulation for pK calculations. *PLOS ONE*, 6(5):1–9, 2011.

[155] AC Martin, Janet C Cheetham, and Anthony R Rees. Modeling antibody hypervariable loops: a combined algorithm. *Proceedings of the National Academy of Sciences*, 86(23):9268–9272, 1989.

[156] Nicholas R.J. Whitelegg and Anthony R. Rees. WAM: an improved algorithm for modelling antibodies on the WEB. *Protein Engineering, Design and Selection*, 13(12):819–824, 2000.

[157] Hiroki Shirai, Nobuyuki Nakajima, Junichi Higo, Akinori Kidera, and Haruki Nakamura. Conformational sampling of CDR-H3 in antibodies by multicanonical molecular dynamics simulation. *Journal of Molecular Biology*, 278(2):481 – 496, 1998.

[158] Daisuke Kuroda, Hiroki Shirai, Masato Kobori, and Haruki Nakamura. Structural classification of cdr-h3 revisited: A lesson in antibody modeling. *Proteins: Structure, Function, and Bioinformatics*, 73(3):608–620, 2008.

[159] M.F. Thorpe. Continuous deformations in random networks. *Journal of Non-Crystalline Solids*, 57(3):355 – 370, 1983.

[160] Brandon M. Hespenheide, A.J. Rader, M.F. Thorpe, and Leslie A. Kuhn. Identifying protein folding cores from the evolution of flexible regions during unfolding. *Journal of Molecular Graphics and Modelling*, 21(3):195 – 207, 2002.

[161] Donald J. Jacobs, A. J. Rader, Leslie A. Khun, and M. F. Thorpe. Protein flexibility predictions using graph theory. *Proteins*, 44:150–160, 2001.

[162] A J Rader. Thermostability in rubredoxin and its relationship to mechanical rigidity. *Physical Biology*, 7(1):016002, 2009.

[163] Stephen A Wells, Martin T Dove, and Matthew G Tucker. Finding best-fit polyhedral rotations with geometric algebra. *Journal of Physics: Condensed Matter*, 14(17):4567, 2002.

[164] Stephen Wells, Martin Dove, and Matthew Tucker. Reverse Monte Carlo with geometric analysis - RMC+GA. *Journal of Applied Crystallography*, 37(4):536–544, 2004.

[165] Stephen A. Wells and Asel Sartbaeva. GASP: Software for geometric simulations of flexibility in polyhedral and molecular framework structures. *Mol. Simul.*, 41(16-17):1409–1421, 2015.

[166] J E Jimenez-Roldan, R B Freedman, R A Römer, and S A Wells. Rapid simulation of protein motion: merging flexibility, rigidity and normal mode analyses. *Phys. Biol.*, 9(1):016008, 2012.

[167] Rudolf A. Römer, Stephen A. Wells, J. Emilio Jimenez-Roldan, Moitrayee Bhattacharyya, Saraswathi Vishweshwara, and Robert B. Freedman. The flexibility and dynamics of protein disulfide isomerase. *Proteins: Structure, Function, and Bioinformatics*, 84(12):1776–1785, 2016.

[168] W. G. Krebs, Vadim Alexandrov, Cyrus A. Wilson, Nathaniel Echols, Haiyuan Yu, and Mark Gerstein. Normal mode analysis of macromolecular motions in a database framework: Developing mode concentration as a useful classifying statistic. *Proteins*, 48(4):682–695, 2002.

[169] Monique M. Tirion. Large amplitude elastic motions in proteins from a single-parameter, atomic analysis. *Phys. Rev. Lett.*, 77(9):1905–1908, 1996.

[170] Holger Gohlke and M. F. Thorpe. A natural coarse graining for simulating large biomolecular motion. *Biophysical Journal*, 91(6):2115–2120, 2006.

[171] Ivet Bahar, Ali Rana Atilgan, and Burak Erman. Direct evaluation of thermal fluctuations in proteins using a single-parameter harmonic potential. *Folding and Design*, 2(3):173 – 181, 1997.

[172] Chakra Chennubhotla, A J Rader, Lee-Wei Yang, and Ivet Bahar. Elastic network models for understanding biomolecular machinery: from enzymes to supramolecular assemblies. *Physical Biology*, 2(4):S173–S180, 2005.

[173] Konrad Hinsen. Analysis of domain motions by approximate normal mode calculations. *Proteins: Structure, Function, and Bioinformatics*, 33(3):417–429, 1998.

[174] Eran Eyal, Lee-Wei Yang, and Ivet Bahar. Anisotropic network model: systematic evaluation and a new web interface. *Bioinformatics*, 22(21):2619–2627, 2006.

[175] N Go, T Noguti, and T Nishikawa. Dynamics of a small globular protein in terms of low-frequency vibrational modes. *Proc Natl Acad Sci U S A*, 80(12):3696–3700, Jun 1983.

[176] Matthieu Wyart. On the rigidity of amorphous solids. *Ann. Phys. Fr.*, 30(3):1–96, 2005.

[177] M. Wyart, S. R. Nagel, and T. A. Witten. Geometric origin of excess low-frequency vibrational modes in weakly connected amorphous solids. *Europhysics Letters*, 72(3):486–492, 2005.

[178] Matthieu Wyart, Leonardo E. Silbert, Sidney R. Nagel, and Thomas A. Witten. Effects of compression on the vibrational modes of marginally jammed solids. *Phys. Rev. E - Statistical, Nonlinear, and Soft Matter Physics*, 72(5):1–11, 2005.

[179] Carolina Brito and Matthieu Wyart. Heterogeneous dynamics, marginal stability and soft modes in hard sphere glasses. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):L08003, 2007.

[180] Asaph Widmer-Cooper, Peter Harrowell, and H. Fynewever. How reproducible are dynamic heterogeneities in a supercooled liquid? *Phys. Rev. Lett.*, 93(13):1–4, 2004.

[181] Asaph Widmer-Cooper and Peter Harrowell. Predicting the long-time dynamic heterogeneity in a supercooled liquid on the basis of short-time heterogeneities. *Phys. Rev. Lett.*, 96(18):2–5, 2006.

[182] Asaph Widmer-Cooper, Heidi Perry, Peter Harrowell, and David R. Reichman. Irreversible reorganization in a supercooled liquid originates from localized soft modes. *Nature Physics*, 4(9):711–715, 2008.

[183] Asaph Widmer-Cooper, Heidi Perry, Peter Harrowell, and David R. Reichman. Localized soft modes and the supercooled liquid's irreversible passage through its configuration space. *J. Chem. Phys.*, 131(19), 2009.

[184] M. Marinov and N. Zotov. Model investigation of the Raman spectra of amorphous silicon. *Phys. Rev. B - Condensed Matter and Materials Physics*, 55(5):2938–2944, 1997.

[185] Mitch Mailman, Carl F. Schreck, Corey S. O'Hern, and Bulbul Chakraborty. Jamming in systems composed of frictionless ellipse-shaped particles. *Phys. Rev. Lett.*, 102(25):1–4, 2009.

[186] Z. Zeravcic, N. Xu, A. J. Liu, S. R. Nagel, and W. Van Saarloos. Excitations of ellipsoid packings near jamming. *Epl*, 87(2), 2009.

[187] Elsen Tjhung and Takeshi Kawasaki. Excitation of vibrational soft modes in disordered systems using active oscillation. *Soft Matter*, 13(1):111–118, 2017.

[188] M. L. Manning and A. J. Liu. Vibrational modes identify soft spots in a sheared disordered packing. *Phys. Rev. Lett.*, 107(10):2–5, 2011.

[189] Matthias Merkel and M. Lisa Manning. A geometrically controlled rigidity transition in a model for confluent 3D tissues. *New Journal of Physics*, 20(2), 2018.

[190] Ludovic Berthier, Angel J. Moreno, and Grzegorz Szamel. Increasing the density melts ultrasoft colloidal glasses. *Phys. Rev. E - Statistical, Nonlinear, and Soft Matter Physics*, 82(6):1–4, 2010.

[191] Nicoletta Gnan and Emanuela Zaccarelli. The microscopic role of deformation in the dynamics of soft colloids. *Nature Physics*, 15(7):683–688, 2019.

[192] Leonardo E. Silbert, Andrea J. Liu, and Sidney R. Nagel. Normal modes in model jammed systems in three dimensions. *Phys. Rev. E - Statistical, Nonlinear, and Soft Matter Physics*, 79(2):1–7, 2009.

[193] V. K. de Souza and P. Harrowell. Rigidity percolation and the spatial heterogeneity of soft modes in disordered materials. *Proceedings of the National Academy of Sciences*, 106(36):15136–15141, 2009.

[194] M. Wyart, S. R. Nagel, and T. A. Witten. Geometric origin of excess low-frequency vibrational modes in weakly connected amorphous solids. *EPL (Europhysics Letters)*, 72(3):486, 2005.

[195] J. Clerk Maxwell. On the calculation of the equilibrium and stiffness of frames. *Philos. Mag.*, 27(182):294–299, 1864.

[196] H. Schober and C. Oligschleger. Low-frequency vibrations in a model glass. *Phys. Rev. B - Condensed Matter and Materials Physics*, 53(17):11469–11480, 1996.

[197] Philip B. Allen, Joseph L. Feldman, Jaroslav Fabian, and Frederick Wooten. Diffusons, locons and propagons: Character of atomie yibrations in amorphous Si. *Philosophical Magazine B: Physics of Condensed Matter; Statistical Mechanics, Electronic, Optical and Magnetic Properties*, 79(11-12):1715–1731, 1999.

[198] Peter K. Morse and Eric I. Corwin. Echoes of the glass transition in athermal soft spheres. *Phys. Rev. Lett.*, 119:118003, 2017.

[199] Richard M. Stratt. The Instantaneous Normal Modes of Liquids. *Accounts of Chemical Research*, 28(5):201–207, 1995.

[200] N.W. Ashcroft and N.D. Mermin. *Solid State Physics*. HRW international editions. Holt, Rinehart and Winston, 1976.

[201] Hamid Reza Seyf, Wei Lv, Andrew Rohskopf, and Asegun Henry. The importance of phonons with negative phase quotient in disordered solids. *Scientific Reports*, 8(1):2627, 2018.

[202] Antina Ghosh, Vijayakumar Chikkadi, Peter Schall, and Daniel Bonn. Connecting structural relaxation with the low frequency modes in a hard-sphere colloidal glass. *Phys. Rev. Lett.*, 107:188303, 2011.

[203] Matthieu Wyart and Michael E. Cates. Does a growing static length scale control the glass transition? *Phys. Rev. Lett.*, 119:195501, 2017.

[204] J.C. Phillips. Topology of covalent non-crystalline solids i: Short-range order in chalcogenide alloys. *Journal of Non-Crystalline Solids*, 34(2):153–181, 1979.

[205] R. J. Boyd. The relative sizes of atoms. *J. Phys. B*, 10(12):2283–2291, 1977.

[206] Salvatore Alfano and Meredith L. Greer. Determining If Two Solid Ellipsoids Intersect. *J. Guid. Control Dyn.*, 26(1):106–110, 2003.

[207] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed 2017-03-19].

[208] Boris D. Lubachevsky and Frank H. Stillinger. Geometric properties of random disk packings. *Journal of Statistical Physics*, 60(5):561–583, 1990.

[209] M Galassi et al. *GNU Scientific Library Reference Manual - Third Edition*. Network Theory Ltd., 3rd edition, 2009.

[210] F.A.M. Gomes and D. Sorensen. *ARPACK++, a C++ implementation of ARPACK eigenvalue package*. Center for Research on Parallel Computation, Rice University, 1998.

[211] Anton Melnyk. *Relating chemical structure and molecular packing to charge transport in conjugated polymers*. Johannes Gutenberg-Universität Mainz, 2016.

[212] H. Caswell. *Matrix Population Models: Construction, Analysis, and Interpretation*. Sinauer Associates, 2001.

[213] Rachael L. Fleurence and Christopher S. Hollenbeak. Rates and probabilities in economic modelling. *PharmacoEconomics*, 25(1):3–6, 2007.

[214] L. J. A. Koster. Charge carrier mobility in disordered organic blends for photovoltaics. *Phys. Rev. B*, 81:205318, 2010.

[215] Paul E. Shaw, Arvydas Ruseckas, and Ifor D. W. Samuel. Exciton diffusion measurements in poly(3-hexylthiophene). *Advanced Materials*, 20(18):3516–3520, 2008.

[216] Anita Zeidler and Philip S. Salmon. Pressure-driven transformation of the ordering in amorphous network-forming materials. *Phys. Rev. B*, 93(21), 2016.

[217] Leonardo E. Silbert, Deniz Erta, Gary S. Grest, Thomas C. Halsey, and Dov Levine. Geometry of frictionless and frictional sphere packings. *Phys. Rev. E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 65(3):1–6, 2002.

[218] Giorgio Parisi and Francesco Zamponi. The ideal glass transition of hard spheres. *J. Chem. Phys.*, 123(14), 2005.

[219] Jeppe C Dyre. Simple liquids' quasiuniversality and the hard-sphere paradigm. *Journal of Physics: Condensed Matter*, 28(32):323001, 2016.

[220] J.-P. Hansen and I.R. McDonald. *The Theory of Simple Liquids*. Academic Press, 3rd edition, 2008.

[221] Eran Bouchbinder and Edan Lerner. Universal disorder-induced broadening of phonon bands: from disordered lattices to glasses. *New Journal of Physics*, 20(7):073022, 2018.

[222] M van Hecke. Jamming of soft particles: geometry, mechanics, scaling and isostaticity. *Journal of Physics: Condensed Matter*, 22(3):033101, 2009.

[223] Jürgen Horbach and Walter Kob. Relaxation dynamics of a viscous silica melt: The intermediate scattering functions. *Phys. Rev. E*, 64(4):14, 2001.

[224] Harry E. Pence and Antony Williams. Chemspider: An online chemical information resource. *J. Chem. Educ.*, 87(11):1123–1124, 2010.

[225] Marcus D. Hanwell, Donald E. Curtis, David C. Lonie, Tim Vandermeersch, Eva Zurek, and Geoffrey R. Hutchison. Avogadro: an advanced semantic chemical

editor, visualization, and analysis platform. *Journal of Cheminformatics*, 4(1):17, 2012.

[226] S. Plimpton. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.*, 117:1–19, 1995.

[227] 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. *J. Am. Chem. Soc.*, 118(45):11225–11236, 1996.

[228] Pascal Friederich, Franz Symalla, Velimir Meded, Tobias Neumann, and Wolfgang Wenzel. Ab initio treatment of disorder effects in amorphous organic materials: Toward parameter free materials simulation. *Journal of Chemical Theory and Computation*, 10(9):3720–3725, 2014.

[229] Furitsu Suzuki, Shosei Kubo, Tatsuya Fukushima, and Hironori Kaji. Effects of structural and energetic disorders on charge transports in crystal and amorphous organic layers. *Sci. Rep.*, 8(1):1–9, 2018.

[230] Azadeh Farahzadi, Maryam Beigmohamadi, Phenwisa Niyamakom, Stephan Kremers, Nico Meyer, Michael Heuken, and Matthias Wuttig. Characterization of amorphous organic thin films, determination of precise model for spectroscopic ellipsometry measurements. *Appl. Surf. Sci.*, 256(22):6612–6617, 2010.

[231] Maki Shibata, Yoshiya Sakai, and Daisuke Yokoyama. Advantages and disadvantages of vacuum-deposited and spin-coated amorphous organic semiconductor films for organic light-emitting diodes. *J. Mater. Chem. C*, 3(42):11178–11191, 2015.

[232] M Ricci, OM Roscioni, L Querciagrossa, and C Zannoni. Molc. a reversible coarse grained approach using anisotropic beads for the modelling of organic functional materials. *Physical Chemistry Chemical Physics (to be published)*, 2019.

[233] Francesco Frigerio, Mos Casalegno, Chiara Carbonera, Tommaso Nicolini, Stefano Valdo Meille, and Guido Raos. Molecular dynamics simulations of the solvent- and thermal history-dependent structure of the PCBM fullerene derivative. *J. Mater. Chem.*, 22:5434–5443, 2012.

[234] Shuichi Nosé. A unified formulation of the constant temperature molecular dynamics methods. *J. Chem. Phys.*, 81(1):511–519, 1984.

[235] Stephan Büchner and Andreas Heuer. Potential energy landscape of a model glass former: Thermodynamics, anharmonicities, and finite size effects. *Phys. Rev. E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 60(6):6507–6518, 1999.

[236] Andreas Heuer and Stephan Büchner. Why is the density of inherent structures of a Lennard-Jones-type system Gaussian? *Journal of Physics Condensed Matter*, 12(29):6535–6541, 2000.

[237] F. Sciortino, W. Kob, and P. Tartaglia. Inherent structure entropy of supercooled liquids. *Phys. Rev. Lett.*, 83(16):3214–3217, 1999.

[238] Pascal Friederich, Artem Fediai, Simon Kaiser, Manuel Konrad, Nicole Jung, and Wolfgang Wenzel. Toward design of novel materials for organic electronics. *Advanced Materials*, 31(26):1808256, 2019.

[239] R.G.E. Kimber. *Charge and Energy Transport in Organic Semiconductors*. (Thesis). University of Bath, 2011.

[240] C. Groves, R. A. Marsh, and N. C. Greenham. Monte carlo modeling of geminate recombination in polymer-polymer photovoltaic devices. *The J. Chem. Phys.*, 129(11):114903, 2008.

[241] Cheng Kuang Lee and Chun Wei Pao. Nanomorphology evolution of p3ht/pcbm blends during solution-processing from coarse-grained molecular simulations. *J. Phys. Chem. C*, 118(21):11224–11233, 2014.

[242] Cheng Kuang Lee, Chun Wei Pao, and Chih Wei Chu. Multiscale molecular simulations of the nanoscale morphologies of P3HT:PCBM blends for bulk heterojunction organic photovoltaic cells. *Energy and Environmental Science*, 4(10):4124–4132, 2011.

[243] Samuel E. Root, Suchol Savagatrup, Christopher J. Pais, Gaurav Arya, and Darren J. Lipomi. Predicting the Mechanical Properties of Organic Semiconductors Using Coarse-Grained Molecular Dynamics Simulations. *Macromolecules*, 49(7):2886–2894, 2016.

[244] Joachim Loos. Volume morphology of printable solar cells. *Materials Today*, 13(10):14 – 20, 2010.

[245] Thusitha Etampawala, Dilru Ratnaweera, Brian Morgan, Souleymane Diallo, Eugene Mamontov, and Mark Dadmun. Monitoring the dynamics of miscible

P3HT:PCBM blends: A quasi elastic neutron scattering study of organic photovoltaic active layers. *Polymer*, 61:155 – 162, 2015.

[246] Anne A. Y. Guilbert, Malte Schmidt, Annalisa Bruno, Jizhong Yao, Simon King, Sachetan M. Tuladhar, Thomas Kirchartz, M. Isabel Alonso, Alejandro R. Goi, Natalie Stingelin, Saif A. Haque, Mariano Campoy-Quiles, and Jenny Nelson. Spectroscopic evaluation of mixing and crystallinity of fullerenes in bulk heterojunctions. *Advanced Functional Materials*, 24(44):6972–6980, 2014.

[247] Brian A. Collins, John R. Tumbleston, and Harald Ade. Miscibility, crystallinity, and phase development in P3HT/PCBM solar cells: Toward an enlightened understanding of device morphology and stability. *The Journal of Physical Chemistry Letters*, 2(24):3135–3145, 2011.

[248] Joshua N. Milstein and Jens-Christian Meiners. *Worm-Like Chain (WLC) Model*, pages 2757–2760. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[249] Loes M. J. Kroon-Batenburg, Peter H. Kruiskamp, Johannes F. G. Vliegenthart, and Jan Kroon. Estimation of the persistence length of polymers by MD simulations on small fragments in solution. application to cellulose. *The Journal of Physical Chemistry B*, 101(42):8454–8459, 1997.

[250] Albert P. Philipse. The random contact equation and its implications for (colloidal) rods in packings, suspensions, and anisotropic powders. *Langmuir*, 12(5):1127–1133, 1996.

[251] Yuri Yu. Tarasevich, Andrei S. Burmistrov, Taisiya S. Shinyaeva, Valeri V. Laptev, Nikolai V. Vygornitskii, and Nikolai I. Lebovka. Percolation and jamming of linear k-mers on a square lattice with defects: Effect of anisotropy. *Phys. Rev. E*, 92:062142, 2015.

[252] Yu Yu Tarasevich, D O Dubinin, V V Laptev, and N I Lebovka. Impact of defects on electrical connectivity of monolayer of ideally aligned rods. *Journal of Physics: Conference Series*, 681:012038, 2016.

[253] Syahrul Ulum, Natalie Holmes, Darmawati Darwis, Kerry Burke, A.L. David Kilcoyne, Xiaojing Zhou, Warwick Belcher, and Paul Dastoor. Determining the structural motif of P3HT:PCBM nanoparticulate organic photovoltaic devices. *Solar Energy Materials and Solar Cells*, 110:43 – 48, 2013.

[254] Natalie P. Holmes, Syahrul Ulum, Prakash Sista, Kerry B. Burke, Mitchell G. Wilson, Mihaela C. Stefan, Xiaojing Zhou, Paul C. Dastoor, and Warwick J.

Belcher. The effect of polymer molecular weight on P3HT:PCBM nanoparticulate organic photovoltaic device performance. *Solar Energy Materials and Solar Cells*, 128:369 – 377, 2014.

[255] Dian Chen, Feng Liu, Cheng Wang, Atsuhiro Nakahara, and Thomas P. Russell. Bulk heterojunction photovoltaic active layers via bilayer interdiffusion. *Nano Letters*, 11(5):2071–2078, 2011. PMID: 21476579.

[256] Patrick R. Amestoy, Alfredo Buttari, Jean-Yves LExcellent, and Theo Mary. Performance and scalability of the block low-rank multifrontal factorization on multicore architectures. *ACM Trans. Math. Softw.*, 45(1), February 2019.

[257] William Humphrey, Andrew Dalke, and Klaus Schulten. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics*, 14:33–38, 1996.

[258] Kaname Kanai, Takahiro Miyazaki, Hiroyuki Suzuki, Mina Inaba, Yukio Ouchi, and Kazuhiko Seki. Effect of annealing on the electronic structure of poly(3-hexylthiophene) thin film. *Phys. Chem. Chem. Phys.*, 12:273–282, 2010.

[259] Nico Seidler, Giovanni Mattia Lazzerini, Giovanni Li Destri, Giovanni Marletta, and Franco Cacialli. Enhanced crystallinity and film retention of P3HT thin-films for efficient organic solar cells by use of preformed nanofibers in solution. *J. Mater. Chem. C*, 1:7748–7757, 2013.

[260] P. Aruna, K. Suresh, and C.M. Joseph. Effect of fullerene doping on the electrical properties of P3HT/PCBM layers. *Materials Science in Semiconductor Processing*, 36:7 – 12, 2015.

[261] Takeaki Sakurai, Toshihiro Yamanari, Masato Kubota, Susumu Toyoshima, Tetsuya Taima, Yuji Yoshida, and Katsuhiro Akimoto. Comparative study on structural properties of poly(3-hexylthiophene) and poly(3-hexylthiophene):6,6-phenyl-c61butyric acid methyl ester thin films using synchrotron x-ray diffraction. *Japanese Journal of Applied Physics*, 49(1):01AC01, 2010.

[262] Renato Colle, Giuseppe Grosso, Alberto Ronzani, and Claudio M. Zicovich-Wilson. Structure and x-ray spectrum of crystalline poly(3-hexylthiophene) from dft-van der waals calculations. *Phys. Status Solidi B*, 248(6):1360–1368, 2011.

[263] Christopher L. Farrow and Simon J. L. Billinge. Relationship between the atomic pair distribution function and small-angle scattering: implications for modeling of nanoparticles. *Acta Crystallographica Section A*, 65(3):232–239, 2009.

[264] E Prince, H Fuess, Th Hahn, H Wondratschek, U Mller, U Shmueli, Andre Authier, Vojtech Kopsky, D. Litvin, M Rossmann, E Arnold, S Hall, and B. McMahon. *International Tables for Crystallography, Vol. C: Mathematical, Physical and Chemical Tables*. IUCR, 2006.

[265] Wichard J. D. Beenken and Tnu Pullerits. Spectroscopic units in conjugated polymers: a quantum chemically founded concept? *The Journal of Physical Chemistry B*, 108(20):6164–6169, 2004. PMID: 18950096.

[266] Jenny Clark, Carlos Silva, Richard H. Friend, and Frank C. Spano. Role of intermolecular coupling in the photophysics of disordered organic semiconductors: Aggregate emission in regioregular polythiophene. *Phys. Rev. Lett.*, 98:206406, 2007.

[267] Xinran Zhang, Hugo Bronstein, Auke J. Kronemeijer, Jeremy Smith, Youngju Kim, R. Joseph Kline, Lee J. Richter, Thomas D. Anthopoulos, Henning Sirringhaus, Kigook Song, Martin Heeney, Weimin Zhang, Iain McCulloch, and Dean M. DeLongchamp. Molecular origin of high field-effect mobility in an indacenodithiophene–benzothiadiazole copolymer. *Nature Communications*, 4:2238 EP –, 2013.

[268] Henry Opoku, Benjamin Nketia-Yawson, Eun-Sol Shin, and Yong-Young Noh. Organic field-effect transistors processed by an environmentally friendly non-halogenated solvent blend. *J. Mater. Chem. C*, 6:661–667, 2018.

[269] C. Fonseca Guerra, J. G. Snijders, G. te Velde, and E. J. Baerends. Towards an order-n dft method. *Theoretical Chemistry Accounts*, 99(6):391–403, 1998.

[270] Evert Jan Baerends, Tom Ziegler, A. J. Atkins, Jochen Autschbach, Donald Bashford, O. Baseggio, A. Bérces, F. Matthias Bickelhaupt, C. Bo, P. M. Boerritger, Luigi Cavallo, C. Daul, D. P. Chong, D. V. Chulhai, L. Deng, R. M. Dickson, J. M. Dieterich, D. E. Ellis, M. van Faassen, A. Ghysels, A. Giammona, Stan J. A. van Gisbergen, A. Goez, A. W. Götz, S. Gusarov, F. E. Harris, P. van den Hoek, Z. Hu, Christoph R. Jacob, H. Jacobsen, L. Jensen, L. Joubert, J. W. Kaminski, G. van Kessel, C. König, F. Kootstra, A. Kovalenko, Mykhaylo Krykunov, Erik van Lenthe, D. A. McCormack, A. Michalak, M. Mitoraj, S. M. Morton, Johannes Neugebauer, V. P. Nicu, Louis Noodleman, V. P. Osinga, S. Patchkovskii, Michele Pavanello, C. A. Peeples, Pierre Herman Theodoor Philipsen, D. Post, Cory C. Pye, H. Ramanantoanina, P. Ramos, W. Ravenek, J. I. Rodríguez, P. Ros, R. Rüger, P. R. T. Schipper, D. Schlüns, H. van Schoot, G. Schreckenbach,

J. S. Seldenthuis, Mike Seth, Jaap G. Snijders, Miquel Solà, Stener M., M. Swart, D. Swerhone, G. te Velde, V. Tognetti, P. Vernooijs, L. Versluis, Lucas Visscher, O. Visser, Fan Wang, T. A. Wesolowski, E. M. van Wezenbeek, G. Wiesenekker, S. K. Wolff, T. K. Woo, and A. L. Yakovlev. ADF2019, SCM, Theoretical Chemistry, Vrije Universiteit, Amsterdam, The Netherlands, https://www.scm.com, 2019.

[271] Y. Oliver. private communication, 2019.