

Investigation of the electronic conduction of large molecules via semi-empirical electronic structure techniques

Gareth Jones

A Thesis submitted for the degree of Doctor of Philosophy

School of Physics and Astronomy

Cardiff University

August 2012

Abstract

In this thesis a new computer code is developed to perform non-equilibrium Green's function based calculations of electronic transmission, using a Hamiltonian computed from self consistent extended Hückle theory as input. Individual elements of this code are tested to ensure correctness. To evaluate its usefulness, the code is tested on porphyrin based systems against the more traditional density functional theory methods of generating the required Hamiltonian. It is then used on more complex porphyrin systems, and comments are made on the use of porphyrin in molecular electronics. Finally it is used on DNA based systems too large to be dealt with efficiently via density functional theory to provide predictions of the effects of DNA structure on its conductance.

Declaration

This work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

Signed Date

Statement 1

The thesis is being submitted in partial fulfilment of the requirements for the degree of PhD.

Signed Date

Statement 2

This thesis is the result of my own independent work/investigation, except where otherwise stated. All sources are acknowledged by explicit references. The views expressed are my own.

Signed Date

Statement 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed Date

Acknowledgements

This work was supported by a studentship grant from EPSRC. I also acknowledge Advanced Research Computing @ Cardiff (ARCCA) for providing run time on the Merlin HPC cluster. I would like to thank my supervisors Dr C C Matthai and Dr M Elliott for all of their help and discussion.

Contents

1	Introduction	1
2	Theory	7
2.1	Introduction	8
2.2	Models of conduction	8
2.2.1	A simple one level model	8
2.2.2	A realistic model: The Non-equilibrium Green's function method	11
2.3	Density functional theory	16
2.3.1	Introduction	16
2.3.2	The Born-Oppenheimer approximation	17
2.3.3	The Hohenberg-Kohn theorem	18
2.3.4	The Kohn-Sham equations	21
2.3.5	Pseudopotentials	23
2.3.6	Geometry optimisation	24
2.3.7	Extensions	26
2.4	Extended Hückel Theory	27
2.5	Periodic systems	30
3	Description of Code	33
3.1	Software	34
3.2	EHTransport	34
3.2.1	The input/output module	35
3.2.2	The EHT module	35

3.2.3	The NEGF module	51
4	Choice of Parameters and Testing of EHTransport	57
4.1	Introduction	58
4.2	Choice of SIESTA Parameters	58
4.3	Testing of EHTransport's EHT module	61
4.3.1	Overlaps of Slater orbitals	61
4.3.2	EHT tests of small molecules	62
4.3.3	Mulliken analysis tests - Tetrafluoromethane	64
4.3.4	Tests of EHTransport's Monkhorst-Pack implementation	65
4.3.5	Poisson solver	65
4.3.6	Band structure and orbital PDOS	67
4.4	Tests of the Green's functions code	68
4.4.1	Model Hamiltonian input	68
4.4.2	SIESTA Hamiltonian input	70
4.5	Test calculations	71
4.5.1	Gold capacitor	71
4.5.2	Monatomic gold wire	73
5	Conduction in Porphyrin Molecules	75
5.1	Introduction	76
5.2	Modelling the contacts	77
5.3	The porphyrin-gold interface	79
5.4	Effects of rotation on system energy	82
5.4.1	Chain contacts	83
5.4.2	Bulk contacts	84
5.5	Effects of rotation on molecular conductance	85
5.5.1	Chain contacts	85
5.5.2	Bulk contacts	87
5.6	Computation time	89
5.7	Conclusion	90

6	EHT calculations of large Porphyrin Systems	93
6.1	Introduction	94
6.2	Effects of wire length on molecular conductance	94
6.3	Conductance in the Iron porphyrin system	97
6.4	Effects of water molecules of molecular conductance	99
6.5	Conductance in the cytochrome b562 protein	101
6.6	Conclusion	104
7	EHT calculations of DNA based systems	105
7.1	Introduction	106
7.2	Structure	107
7.3	Effects of DNA sequence on molecular conductance	110
7.4	Effects of DNA structural failings on molecular conductance	112
7.5	Conclusion	113
8	Conclusion	115
8.1	Summary	116
8.2	Further work	117
A	EHTransport manual	119
A.1	Introduction	120
A.2	Performing an Extended Hückel calculation	122
A.2.1	Setup, input and output	122
A.2.2	Valid keywords	123
A.3	Performing a NEGF calculation	132
A.3.1	Setup, input and output	132
A.3.2	Valid keywords	132
A.4	Other available job types	134
A.4.1	UnitTest	134
A.4.2	SiestaTransport	134
A.4.3	PrecalcLeads	134
A.5	Code layout	135

A.5.1	File structure	135
A.5.2	Adding additional atomic species	137

List of Figures

1.1	A schematic of a typical molecular conduction experiment, with an additional gate contact included.	2
1.2	A free base porphyrin molecule, in which there the central metal ion is replaced by a pair of hydrogen atoms. The colour scheme shows the electron density in the plane of the molecule, computed by DFT.	4
1.3	A DNA strand ten base pairs in length, containing adenine and thymine base pairs. . .	5
2.1	An energy level schematic of the one level model.	8
2.2	A plot of current and the differential conductance against voltage at different levels of complexity of the one level model, assuming that ϵ lies at the mid point between μ_L and μ_R	11
2.3	An example system of a monatomic gold wire between bulk contacts, with the scattering region marked.	12
2.4	An example system showing the extent of the lead and scattering regions. In this example, each large circle represents one principal layer of the lead, and the two smaller circles represent the molecule of interest.	12
2.5	Comparison of the $n = 2$ single particle wavefunctions computed using all-electron DFT and using a pseudopotential. Beyond a cutoff distance the two overlap exactly.	24
2.6	Comparison of the Slater radial wave-function with the hydrogenic one for $n = 3$, $l = 0$. The angular part is identical in both cases.	29
3.1	A plots of orbital energy against excess charge for the valence orbitals of a single isolated atom. From left to right and top to bottom, the atomic species are gold, carbon, hydrogen and iron. Over the range shown, the curves are very close to quadratic.	51

3.2	The contour of integration followed when computing the density matrix of the scattering region via the Green's function method.	54
4.1	Convergence of system energy while varying the real space mesh energy cutoff (left) and k-point mesh density (right) of a bulk sodium system in SIESTA. For k-point calculations, the Monkhorst-Pack matrix had diagonal elements n and off-diagonal elements 0. The x axis values correspond to n	59
4.2	Convergence of energy while varying the real space mesh energy cutoff (left) and k-point mesh density (right) of a bulk gold system in SIESTA. For k-point calculations, the Monkhorst-Pack matrix had diagonal elements n and off-diagonal elements 0. The x axis values correspond to n	60
4.3	The HOMO and LUMO energies of a hydrogen molecule as the interatomic spacing is varied.	63
4.4	The energy levels of a benzene molecule computed with EHTransport and SIESTA. The two diverge significantly above the Fermi level.	64
4.5	The potential gradient of an isolated cell containing a single hydrogen atom (left), and a pair of hydrogen atoms arranged along the z axis (right) across which a 0.1V bias is applied, taken as a cross-section through the y axis. More rapidly varying colours designate more rapidly changing potential. Bands of the same colour which are not connected do not correspond to the same potential.	67
4.6	The potential gradient of a periodic cell containing a pair of hydrogen atoms (top) and a chain of hydrogen atoms (bottom), across which a 0.1V bias is applied, taken as a cross-section through the atoms. More rapidly varying colours designate more rapidly changing potential. Bands of the same colour which are not connected do not correspond to the same potential.	67
4.7	The bulk gold band structure (left) plotted along lines connecting the high symmetry points of the FCC Brillouin zone and the density of states projected onto orbitals of each l value (right) computed using EHTransport and the Cerda parameter set.	68
4.8	The structure of the one dimensional potential barrier.	69
4.9	The transmission (left) and DOS (right) of one dimensional potential barrier computed using EHTransport. The transmission is also compared to the analytic result.	69

4.10	The transmission (left) and density of states (right) of a bulk hydrogen system at the Γ point.	71
4.11	The transmission (left) and density of states (right) of a bulk hydrogen system using a Monkhorst-Pack matrix with diagonal elements $12 \times 12 \times 1$	71
4.12	A gold (111) capacitor shows exponential scaling of transmission at the Fermi level with plate separation.	72
4.13	The structure of the monatomic gold wire system, represented orthographically. The leads are 3 atoms deep.	73
4.14	Electronic transmission across a monatomic gold wire between two gold (111) electrodes, for different numbers of repetitions of the ABC wire unit.	73
5.1	The structure of a free-base porphyrin ring, without a metal ion at the center.	76
5.2	The structure of the system under study.	77
5.3	The DOS of a monatomic gold chain computed with the different available parameter sets in EHTransport, and compared to SIESTA.	78
5.4	Two ABC layers of bulk gold electrode. The direction of conduction is along the z axis, along the [111] direction, from the left to the right of the image.	79
5.5	The structure of the two porphyrin dimers.	80
5.6	The four possible positions on the gold (111) surface at which possible binding was considered. One set of ABC layers are shown, with the C layer (green) being on the surface, followed by the B layer (red) and A layer (yellow).	81
5.7	Binding energy of the thiol group to various positions on the gold surface. The FCC hollow site was found to be the preferred binding point.	81
5.8	Variation of the energy with the distance between the sulphur atom of the porphyrin dimer and the gold surface. The lowest points of the curves were been chosen to be at zero energy, and are at 2.27Å for the gold chain and 2.10Å for the 111 hollow site. . . .	82
5.9	The energy difference against dihedral angle relative to the energy minimum for a zinc-porphyrin dimer with a 2C backbone (left) and a 4C backbone (right) attached to gold chain leads.	83

5.10	The energy difference against dihedral angle relative to the energy minimum for a zinc-porphyrin dimer with a 2C backbone (left) and a 4C backbone (right) attached to bulk gold (111) surfaces.	85
5.11	The energy difference against dihedral angle relative to the energy minimum for a zinc-porphyrin dimer with a 4C backbone attached to gold 111 surfaces, performing a geometry optimization at each step.	85
5.12	The DOS and transmission spectra of a porphyrin dimer between gold chains comparing EHTransport with self consistency enabled and disabled.	86
5.13	Transmission through a Zn-porphyrin dimer between gold chains as a function of dihedral angle, computed using Transiesta (Left) and EHTransport. (Right)	87
5.14	Transmission through a Zn-porphyrin dimer between gold chains at several angles, computed using Transiesta (Left) and EHTransport. (Right)	87
5.15	Bond currents through the porphyrin molecule, computed using Transiesta and Inelastica (Top) and EHTransport (Bottom). Inelastica does not provide currents through contact atoms or directionality information, and labels atoms by index as opposed to the atomic number labels used by EHTransport.	88
5.16	The DOS (left) and transmission spectra (right) of a Zn-porphyrin dimer between gold (111) surfaces comparing EHTransport with self consistency enabled and disabled.	88
5.17	Transmission through a Zn-porphyrin dimer between gold (111) surfaces as a function of dihedral angle, computed using Transiesta (Left) and EHTransport (Right). The Fermi energy has been offset in the EHTransport case in order to line up with the position predicted by Transiesta.	89
5.18	DOS and Transmission through a Zn-porphyrin dimer between gold (111) surfaces at several angles, computed using Transiesta (Top) and EHTransport (Bottom).	90
6.1	The DOS (left) and transmission spectra (right) of a zinc-porphyrin wire of various lengths between gold (111) surfaces.	95
6.2	The conductance as a function of wire length of a zinc-porphyrin wire between gold contacts, at several energies. The pink lines are exponential fits to the first 3 data points at each energy.	96

6.3	The energy gap between the homo and lumo of a zinc-porphyrin wire as a function of chain length, with no gold contacts or thiol groups present.	96
6.4	The DOS (left) and transmission spectra (right) of an iron-porphyrin dimer between gold chains.	97
6.5	The DOS (left) and transmission spectra (right) of an iron-porphyrin dimer between (111) surfaces.	98
6.6	Colour coded representation of the orbitals involved in conduction through a zinc-porphyrin dimer. Red shows s character, green p and blue d.	98
6.7	Colour coded representation of the orbitals involved in conduction through an iron-porphyrin dimer. Red shows s character, green p and blue d.	99
6.8	The structure of the water porphyrin dimer surrounded by water.	99
6.9	The DOS (left) and Transmission spectra (right) of a system consisting of a zinc porphyrin dimer between bulk gold contacts with and without water present, as well as pure water with no porphyrin present.	100
6.10	A charge difference plot showing the difference in real space compared to that of the porphyrin system without water present for the contacts and dimer, and an isolated water molecule for all water molecules in the system. Red denotes excess electrons, and blue a deficit.	101
6.11	Bond-currents plot at the Fermi level of a zinc-porphyrin dimer between bulk gold contacts, and surrounded by polarised water.	101
6.12	Bond-currents plot 2 eV below the Fermi level of a zinc-porphyrin dimer between bulk gold contacts, and surrounded by polarised water.	102
6.13	Cytochrome b_{562} , showing the position of the heme group relative to the protein backbone.	102
6.14	The DOS (left) and transmission spectra (right) of Cytochrome b_{562} between gold chains.	103
7.1	A DNA strand ten base pairs in length, of sequence ATATATATAT.	106
7.2	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads using the original W3DNA generated structure and a version optimized using UFF forcefields.	108
7.3	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads using various possible positioning of hydrogen atoms.	109

7.4	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads using several lead positioning variations.	109
7.5	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads, exploring a narrow energy range around the Fermi level.	110
7.6	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads showing the effect of adding replacing the central CG sequence with an AT in a CGCGCG strand.	111
7.7	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads showing the difference between A-DNA and B-DNA structures. . . .	112
7.8	The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads of a CGCGCGCGCG sequence, in which the 5th base pair in the chain has either been replaced with a C-C combination, or removed completely.	112

List of Tables

3.1	The parameter set used for all EHT calculations in this thesis, excluding the self consistent energy corrections.	50
3.2	The computed parameters α_i and β_i from equation 3.27 for all considered atomic species, using a quadratic best fit.	52
4.1	A selection of overlaps of double-zeta Slater orbitals separated by 1 Bohr in the Y and Z directions. The sign correction has not been included.	62
4.2	A sample set of k-points generated for a simple cubic unit cell of side length one Bohr, a Monkhorst-Pack matrix with diagonal elements (2, 2, 2), zero off-diagonal elements and a zero displacement vector.	65
4.3	A sample set of k-points generated for a FCC unit cell of side length one Bohr, a Monkhorst-Pack matrix with diagonal elements (3, 3, 3), zero off-diagonal elements and a displacement vector of (0.5, 0.5, 0.5). SIESTA is able to uncover additional symmetries to identify equivalent k-points, and hence merge four pairs of k-points together.	66
4.4	The 'bond current' matrix elements corresponding to the first 3 atoms in the bulk hydrogen system, sampled at the gamma point.	72

Chapter 1

Introduction

Manufacturers of semiconductor devices already offer products built on a 22nm process node. Even the 45nm process from a few generations ago could have a gate length of as little of 35nm, with only 1nm of insulation separating the gate from the channel. In pure silicon with its lattice constant of 0.543nm, this is equivalent to features measuring less than 100 atoms across, with only a few atoms of insulation. On this scale quantum effects are already important, with leakage of current from the gate into the channel being a particular problem. As feature sizes shrink further these effects will become progressively more important, to the point where using existing materials is not a viable option; a 0.5nm thickness of oxide insulation, which was predicted to be obtained at the 22nm node, is expected to be the smallest possible [1] preventing further miniaturisation at least in that aspect without resorting to alternative materials which give a larger equivalent physical oxide thickness. To ensure their continued ability to shrink components and provide increased transistor densities, manufactures are looking for novel materials and techniques. Replacing the traditional doped semiconductor channel with an individual molecule is one possibility. Another possibility is identifying molecules which can exist in multiple conformations, each of which is stable at room temperature, allowing the production of molecular memory devices.

The engineering challenges in constructing such a device are formidable, however. Until the latter part of the twentieth century, no one had even been able to attach contact pads to very small objects such as these individual molecules in order to investigate their behaviour, and therefore measurements of their current/voltage characteristics could not be made. These measurements have now been performed for very small wires and some individual molecules [2,3] using tools such as scanning tunnelling microscopes (STMs), which work by passing a small current through a sample, and where the STM tip can be used as a contact. More recently measurements involving an additional gate contact to investigate transistor-like behaviour have also been made [4].




Figure 1.1: A schematic of a typical molecular conduction experiment, with an additional gate contact included.

These molecules are too small to have a meaningful voltage-independent conductance defined, and cannot be described by the normal bulk conductivity equation $\vec{J} = \sigma \vec{E}$. Hence more complex models which account for the individual possible electron energy levels are required. These can lead to results that are not intuitive, such as a resistance that is independent of molecule length [5], (but is still non-zero, and attributed solely to the molecule-wire contacts [6], in contrast to true superconductors,) and non-linear relationships between voltage and current, including negative differential resistance [7]. In the case of carbon nanotubes, even simple logical circuits have been constructed from molecular transistors [8], although we are still some way from building complete microprocessors or other devices from molecular building blocks. One of the most important results is that conductance is quantised: In a conducting channel, conductance cannot exceed a fixed, finite value (G_0) per each available energy level in that channel. The value of G_0 depends only on the charge carried by an electron q and Plank’s constant h (Equation 1.1). This effect can even be measured at room temperature with suitably small wires [9].

$$G_0 = \frac{q^2}{h} = 25.8k\Omega^{-1} \dots \quad (1.1)$$

Several procedures exist for theoretical modelling of electrical conduction through molecules. A simple scheme is Simmons’ model [10, 11], which treats a lead-molecule-lead system as a one dimensional potential barrier problem, with flat lead potentials, separated by the barrier. A more detailed description of electronic transmission is captured by the Landauer-Büttiker formalism [12, 13], which is valid for ballistic conductance, where the separation between the contacts is less than the mean free path of electrons in the scattering region. A more complex description is the non-equilibrium Green’s function [14–16] (NEGF) approach that computes the electronic transmission spectra of a system described by a given Hamiltonian at some energy by computing the Green’s function of that Hamiltonian.

If the system under study is described by a set of atomic positions, which is the likely case for study of a real molecule, the Hamiltonian must first be obtained from this set of atomic positions by some other method. One such method is Density Functional Theory [17] (DFT), which is most often employed for this purpose, but is computationally quite expensive. Other methods can potentially also be used in its place, subject to the ability to retrieve the Hamiltonian matrix elements in an appropriate form. Hartree-Fock [18] and its later derivatives collectively

known as post Hartree-Fock methods [19] are one such group, the more advanced of which can be considerably more accurate than DFT, at the expense of computational scaling that places even small molecules out of the reach of current computational resources. At the other end of the spectrum are empirical and semi-empirical techniques, which could be far less computationally expensive, but only at the expense of reduced accuracy, or smaller domains of applicability.




Figure 1.2: A free base porphyrin molecule, in which the central metal ion is replaced by a pair of hydrogen atoms. The colour scheme shows the electron density in the plane of the molecule, computed by DFT.

One particular family of molecules being investigated as possible candidates for molecular conductors are the porphyrins [20–22], characterised by a planar conjugated ring containing 20 carbon atoms and 4 nitrogen atoms, usually with a metal ion in the centre (Figure 1.2). This family and its derivatives include many common molecules in biology, including electrically active molecules such as the chlorophylls¹, (which undergo a complex redox reaction using energy from an absorbed photon,) which would be expected to show interesting transport behaviours. It may also be possible to construct molecular wires using porphyrins as the building blocks [20,21]. Experiments have shown metal-porphyrins with multiple separate conduction states [23], believed to be due to conformational changes. Porphyrins have also been studied experimentally, and also theoretically by the DFT/NEGF approach [24]. This combination of interesting electrical properties, prior results and wide variety of possible chemical and conformational changes make the porphyrins a useful tool for benchmarking new systems. Although small enough to run with DFT individually, the wide variety of these possible structures make exploring all possible variations a daunting task.

¹Chlorophylls are a derivative, rather than a true porphyrin; the ring contains one unsaturated bond, making it a chlorin. Chlorophylls all also contain a magnesium ion at their centre. The variation between chlorophylls is in the structure of the side chains of the ring.




Figure 1.3: A DNA strand ten base pairs in length, containing adenine and thymine base pairs.

Another molecule which has shown interesting electrical transport properties is DNA [25], possibly the best known of all complex organic molecules. It consists of a twisted backbone pair, connected together by a ladder of base pairs (Figure 1.3). However, even a fairly short DNA sequence of eight base pairs contains hundreds of atoms, making theoretical modelling using these DFT/NEGF techniques difficult with currently available computational resources. It is the DFT step, an operation that scales as $\mathcal{O}(N^3)$ with the number of electrons in a system, that normally consumes the greatest amount of time in this calculation. This raises the question of whether the computationally complex ab-initio DFT stage could be replaced by a less complex empirical or semi-empirical step while still obtaining meaningful and useful results. Further more, a smaller basis set is often used in such simpler calculations, which results in a reduced problem size at the NEGF step, (which also scales as $\mathcal{O}(N^3)$, and would become the limiting factor in a calculation were the DFT stage significantly improved,) speeding that up as a by-product.

One such semi-empirical method which may be appropriate is the Extended Hückel Theory [26] (EHT). Despite the popularity of more complex ab-initio methods, EHT remains in common use, particularly on systems of large size [27–29]. By utilising a set of empirical parameters for each atomic species in a system, EHT is able to efficiently produce a Hamiltonian suitable for use in the NEGF equations for systems of large size. Later enhancements to EHT, such as the self-consistent Extended Hückel Theory [30, 31] (SC-EHT) improve the transferability of these empirical parameters, improving their domain of applicability.

In this thesis, I report on the development, functioning and usefulness of a new SC-EHT/NEGF code which I have named EHTransport. Results from EHTransport have been compared to the existing DFT/NEGF code SIESTA to answer the question of whether the SC-EHT/NEGF com-

bination is capable of producing useful results, and then it is used on its own for the investigation of several system types. In chapter two I have described the underpinning theory of DFT, EHT and NEGF, and then in chapter three the specific algorithmic implementations of these within EHTransport, and develop the empirical parameter set required where existing parameters are not available. In chapter four I examine the effects of some adjustable parameters within SIESTA and present tests of each code unit within EHTransport. In chapter five I have applied the EHTransport code to a selection of porphyrin based systems, which are also investigated with SIESTA for comparisons of speed and results. In chapter six I have investigated further porphyrin systems, including a porphyrin-based protein, using only EHT. In chapter seven I have used the code to investigate DNA based structures, which are too large to study via DFT/NEGF techniques using currently available computational resources, and have compared the results to recent experiment. The final chapter lists the conclusions of this thesis, and additional directions in which EHTransport could be taken.

Chapter 2

Theory

2.1 Introduction

In this chapter, I have built up the NEGF description of electronic transport in molecules, starting from a simple single level model and expanding to the general case of multiple interacting levels. I have also described the EHT and DFT procedures, and how to use them to construct a Hamiltonian suitable for performing NEGF calculations.

2.2 Models of conduction

2.2.1 A simple one level model

To build a simple one level model of molecular conduction we consider a three component system, made up of left and right leads, and a central molecule [32]. Both leads are assumed to contain a completely flat continuum of states, which extends to all energies, and which is filled up to some Fermi level μ_L and μ_R , which may be different for each lead in order to allow for a potential difference across the molecule. The molecule is assumed to contain only a single possible state, at some energy level ϵ . Both leads are coupled to the molecule with coupling constants γ_L and γ_R , but not to each other. Any electron moving from one lead to another must therefore exist within the molecular energy state at some instant.




Figure 2.1: An energy level schematic of the one level model.

For one of the leads in contact with the molecule, and assuming that the presence of the molecule does not have any effect on the lead, electrons will flow between the lead and the molecule to give an occupancy according to the Fermi function

$$f(E) = \frac{1}{1 + e^{(E-\mu)/k_B T}} . \quad (2.1)$$

At low temperatures, this simply implies that the level ϵ will be occupied if it is positioned below μ and unoccupied if above it. The rate of flow of charge between the lead and the molecule can be written as

$$I = -eC(2f(\epsilon) - N) , \quad (2.2)$$

where C is a rate constant of dimension time^{-1} and N is the number of electrons already on the molecule. This rate constant is typically written as γ/\hbar , where γ is the previously mentioned coupling constant, which can now be seen to have units of energy. This seems reasonable: Connected to one lead, the occupation of the molecule will change over time until $f(\epsilon) = N$, with larger coupling resulting in faster transfer. An analogous equation can also be written for the second lead, to give a second current. In a steady state the current between the molecule and one lead must be the negative of that through the other. (i.e. N is constant)

$$-e\frac{\gamma_L}{\hbar}(f_L(\epsilon) - N) = I_L = -I_R = e\frac{\gamma_R}{\hbar}(f_R(\epsilon) - N) . \quad (2.3)$$

It can be seen from inspection of equation 2.2 that the system will always tend towards this steady state, as if the initial occupation of the molecule is not such that $I_L = I_R$ then the difference between I_L and I_R will be in the direction towards the preferred N . It is possible to factor out this preferred N from equation 2.3, giving an expression for N (and hence the current) in terms of the coupling constants.

$$\begin{aligned} N &= \frac{\gamma_L f_L(\epsilon) + \gamma_R f_R(\epsilon)}{\gamma_L + \gamma_R} \\ I &= \frac{e}{\hbar} \frac{\gamma_L \gamma_R}{\gamma_L + \gamma_R} (f_L(\epsilon) - f_R(\epsilon)) . \end{aligned} \quad (2.4)$$

At this point, it is obvious that something is wrong. As already mentioned, previous experimental results have shown that there exists a maximum possible conductance through a single energy level of $G_0 \equiv e^2/h$, whereas this result suggests that conductance can be increased indefinitely by strengthening the coupling between molecule and contact. What has been missed is a broadening effect caused by the coupling to the leads. Although it may be a reasonable assumption to make that the leads are unchanged by the presence of the molecule, (being semi-infinite in size, any finite perturbation can be ignored,) the molecule can be greatly affected by

the presence of the leads. If the original single energy level is broadened, then some portion of it will be above (or below) both μ_L and μ_R , and hence will not contribute to the conduction. If we replace the single level with a Lorentzian function centred at ϵ , integrated area e and with broadening such that $G \rightarrow G_0$ as $\gamma_L, \gamma_R \rightarrow \infty$ the broadening turns out to be $\gamma_L + \gamma_R$, and the density of states is given by

$$D_\epsilon(E) = \frac{(\gamma_L + \gamma_R)/2\pi}{(E - \epsilon)^2 + ((\gamma_L + \gamma_R)/2)^2} . \quad (2.5)$$

The current is now found by an integration over energy.

$$I = \frac{e}{\hbar} \frac{\gamma_L \gamma_R}{\gamma_L + \gamma_R} \int dE D_\epsilon(E) (f_L(E) - f_R(E)) . \quad (2.6)$$

A further effect which has yet to be accounted for is that adding or removing electrons modifies the potential in the molecule. In this simple one level model, the result will be that ϵ will float up and down around its original value ϵ_0 . We can introduce another parameter U which measures the shift in the energy level for each extra electron added to the channel.

$$\epsilon = \epsilon_0 + U(N - N_0) . \quad (2.7)$$

This sets up a circular dependency between the position of ϵ and the number of electrons in the level, forcing the equations to be solved self consistently. Thinking about it as a classical Coulomb repulsion suggests that it will be reduced the more smeared out the molecular energy level is in real space.

One final point to address in this 'one' level model is that an electron feels no interaction due to itself, so why should the level float around in this way when the only change in the electron distribution is in itself? The answer of course is that this level in fact contains two electrons (or holes,) with opposite spins. Up until now it has been implicitly assumed that the occupation number of these two spins has been equal, but in reality there is no such restriction.

$$\begin{aligned} \epsilon_\uparrow &= \epsilon_{\uparrow 0} + U(N_\downarrow - N_{\downarrow 0}) \\ \epsilon_\downarrow &= \epsilon_{\downarrow 0} + U(N_\uparrow - N_{\uparrow 0}) . \end{aligned} \quad (2.8)$$

In the limit of small U this reduces back to the restricted equation. For large U however, there may now be a case where the occupation of one spin is large, which shifts the other spin to the point where it is no longer accessible to either lead, and no longer plays any role in the conduction until the potential difference is great enough to overcome the effect. This effect is known as a Coulomb blockade. The results of these progress effects on the I/V characteristics of this simple system are displayed in figure 2.2.

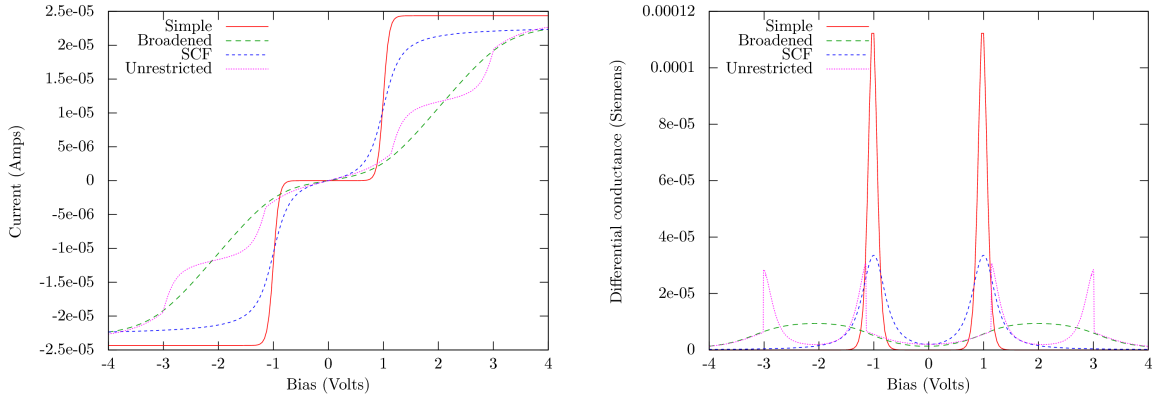


Figure 2.2: A plot of current and the differential conductance against voltage at different levels of complexity of the one level model, assuming that ϵ lies at the mid point between μ_L and μ_R .

The final point that needs to be addressed for this simple model is what exactly it means to apply a voltage across the molecule. It is clear that $\mu_L - \mu_R = V$, but this says nothing about the relationship of μ_L and μ_R with ϵ . In reality it will depend on how the system is set up. For example a third gate contact could be added to the system, fixing the molecular energy level at some ground voltage. (And indeed, transistor action would be observed if the voltage was varied to shift ϵ towards and away from the lead potentials.) In the completely isolated case with symmetrical contacts, ϵ will remain fixed while $\mu_L = \mu_{L0} + V/2$ and $\mu_R = \mu_{R0} - V/2$, but alternative behaviours can lead to interesting effects, such as asymmetries in the current-voltage curves.

2.2.2 A realistic model: The Non-equilibrium Green's function method

The simple one level model may have a use in reality for some tiny molecule at low temperature and bias which has only one energy level within a few $k_B T$ of the lead potentials, but in general a multi-level model will be needed. Such a model could be trivially built from the one level model




Figure 2.3: An example system of a monatomic gold wire between bulk contacts, with the scattering region marked.

by making the assumption that different energy levels within the molecule do not interact, leading to a simple summation of one level model calculations with differing ϵ , but this is unlikely to prove successful as energy levels within the molecule will typically interact with each other more strongly than with the leads. Generalising the logic of the unrestricted model suggests that we will need an $N \times N$ matrix where N is the number of available levels in the molecule to describe the potential shifts of each level due to the occupation of each other level, with the energy becoming a vector of length N , (or $2N$, if spins are still to be treated separately,) along with the coupling energies. This is already a much harder to solve self-consistency problem than the two variable case from the unrestricted one level model, but it nevertheless fails to capture other interactions beyond the classical Coulomb repulsion, such as the effects of electrons hopping from one molecular level to another, rather than from one lead, to a single level, and directly to another lead. Furthermore, we would still have the issue that the leads are described by a flat continuum of states, rather than a realistic surface which will have some structure and possibly even gaps.




Figure 2.4: An example system showing the extent of the lead and scattering regions. In this example, each large circle represents one principal layer of the lead, and the two smaller circles represent the molecule of interest.

First a description of the layout of a more realistic system is needed. Instead of talking only about energy levels, it will be described in terms of orbitals localised in real space, in order to

more easily see the effects of the leads. The system consists of some molecule, connected on two sides to a much larger lead. The leads are taken to be semi-infinite, having some periodic layered structure that extends away from the molecule. Rather than a vector of energy levels, the full Hamiltonian matrix is required. Including the contribution from the leads, this will be infinite in size, but it can be built up from an infinitely repeating series of finite and calculable components. The system is described with a localised basis set such that the range of \mathbf{H} is limited, and only off-diagonal elements between basis functions in some finitely separated layers of the lead are non-zero. Furthermore, there are no non-zero elements coupling functions in the left hand lead to the right hand lead. For each lead a principal layer is chosen, which is at minimum the smallest integer number of the periodic layers required so that no coupling exists between next nearest principal layers. Each principal layer now has a finite Hamiltonian \mathbf{H}_l and \mathbf{H}_r , and another matrix which describes the coupling between each principal layer \mathbf{H}_{lc} and \mathbf{H}_{rc} . The molecule also has a finite Hamiltonian \mathbf{H}_s and coupling matrices between itself and each lead \mathbf{H}_{slc} and \mathbf{H}_{src} . The infinite Hamiltonian describing the system can then be written as

$$\mathbf{H} = \begin{bmatrix} \ddots & \mathbf{H}_{lc} & 0 & 0 & 0 & 0 & 0 \\ \mathbf{H}_{lc}^\dagger & \mathbf{H}_l & \mathbf{H}_{lc} & 0 & 0 & 0 & 0 \\ 0 & \mathbf{H}_{lc}^\dagger & \mathbf{H}_l & \mathbf{H}_{slc} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{H}_{slc}^\dagger & \mathbf{H}_s & \mathbf{H}_{src} & 0 & 0 \\ 0 & 0 & 0 & \mathbf{H}_{src}^\dagger & \mathbf{H}_r & \mathbf{H}_{rc} & 0 \\ 0 & 0 & 0 & 0 & \mathbf{H}_{rc}^\dagger & \mathbf{H}_r & \mathbf{H}_{rc} \\ 0 & 0 & 0 & 0 & 0 & \mathbf{H}_{rc}^\dagger & \ddots \end{bmatrix}. \quad (2.9)$$

This Hamiltonian can now be constructed by analysing three separate finite sized systems. First are two principal layers of the left hand lead with periodic boundary conditions, giving a matrix

$$\mathbf{H}_1 = \begin{bmatrix} \mathbf{H}_l & \mathbf{H}_{lc} \\ \mathbf{H}_{lc}^\dagger & \mathbf{H}_l \end{bmatrix}. \quad (2.10)$$

The same calculation can be performed for the right hand lead. The final section, the extended molecule, or the scattering region, contains not only the molecule under study, but also enough layers of the lead on each side that the coupling between one principal layer and the next has

returned to the bulk value.¹ The matrices \mathbf{H}_{slc} and \mathbf{H}_{src} are then equal to \mathbf{H}_{lc} and \mathbf{H}_{rc} , and the extended molecule can be analysed for the final required element \mathbf{H}_s .

To convert this infinite matrix into something tractable, the non-equilibrium Green's function (NEGF) formalism is used. The Green's function equation for the full (infinite) system is given by

$$(E^+ \mathbf{S} - \mathbf{H}) \mathbf{G}(E) = \mathbf{I} , \quad (2.11)$$

where \mathbf{G} is the Green's function and \mathbf{S} is the overlap matrix:

$$\mathbf{S}_{ij} = \langle \psi_i | \psi_j \rangle , \quad (2.12)$$

where ψ are the functions of the basis set.

\mathbf{S} is the identity matrix for normalised orthogonal basis functions. E^+ is the energy at which the Green's function is being calculated with a positive infinitesimal imaginary component added, which is only a mathematical device to avoid a singularity in the calculation rather than having physical significance.² In terms of the Hamiltonian we have previously constructed, this equation reads

$$\left(E^+ \begin{bmatrix} \ddots & \mathbf{S}_{lc} & 0 & 0 & 0 \\ \mathbf{S}_{lc}^\dagger & \mathbf{S}_l & \mathbf{S}_{lc} & 0 & 0 \\ 0 & \mathbf{S}_{lc}^\dagger & \mathbf{S}_s & \mathbf{S}_{rc} & 0 \\ 0 & 0 & \mathbf{S}_{rc}^\dagger & \mathbf{S}_r & \mathbf{S}_{rc} \\ 0 & 0 & 0 & \mathbf{S}_{rc}^\dagger & \ddots \end{bmatrix} - \begin{bmatrix} \ddots & \mathbf{H}_{lc} & 0 & 0 & 0 \\ \mathbf{H}_{lc}^\dagger & \mathbf{H}_l & \mathbf{H}_{lc} & 0 & 0 \\ 0 & \mathbf{H}_{lc}^\dagger & \mathbf{H}_s & \mathbf{H}_{rc} & 0 \\ 0 & 0 & \mathbf{H}_{rc}^\dagger & \mathbf{H}_r & \mathbf{H}_{rc} \\ 0 & 0 & 0 & \mathbf{H}_{rc}^\dagger & \ddots \end{bmatrix} \right) \begin{bmatrix} \ddots & \vdots & \ddots \\ \dots & \mathbf{G}_s(E) & \dots \\ \ddots & \vdots & \ddots \end{bmatrix} = \mathbf{I} . \quad (2.13)$$

The central block of the Green's function is what is required for the calculation of current. We could write a reduced version of the Green's function equation using only the region of the Hamiltonian corresponding to the basis functions present in the scattering region of the system

¹Boundary conditions on this scattering region are more complex; periodic conditions are only acceptable if the left and right leads have identical principal layers. If not, there is no periodicity used, and additional principal layers are added to each lead to hide any surface effects. In this case the Hamiltonian output from the calculation needs to be modified to strip out any elements due to these extra principal layers before use.

²There are actually advanced and retarded Green's functions, dependent on if E^+ or E^- is in use. In this work only the advanced Green's function is used.

$$(E^+ \mathbf{S}_s - \mathbf{H}_s - \mathbf{\Sigma}_l - \mathbf{\Sigma}_r) \mathbf{G}_s(E) = \mathbf{I} \rightarrow \mathbf{G}_s = [E^+ \mathbf{S}_s - \mathbf{H}_s - \mathbf{\Sigma}_l - \mathbf{\Sigma}_r]^{-1} , \quad (2.14)$$

where $\mathbf{\Sigma}$ is an additional term such that \mathbf{G}_s becomes equal to the result from equation 2.13. This term is known as a self-energy matrix, and can be calculated from the sections of the Hamiltonian related to the leads. For the left lead it has the form

$$\mathbf{\Sigma}_l = (E^+ \mathbf{S}_{lc}^\dagger - \mathbf{H}_{lc}^\dagger) \mathbf{G}_l (E^+ \mathbf{S}_{lc} - \mathbf{H}_{lc}) \quad (2.15)$$

where \mathbf{G}_l is the surface Green's function of the lead, which is the bottom right block of the full semi-infinite Green's function representing the lead. With the scattering region Green's function known, the current is given by an integration over energy.

$$I = \frac{e}{h} \int dE \text{Tr} \left[\mathbf{\Gamma}_l \mathbf{G}_s^\dagger \mathbf{\Gamma}_r \mathbf{G}_s \right] (f(E - \mu_l) - f(E - \mu_r)) \quad (2.16)$$

where $\mathbf{\Gamma}$ is just the anti-hermitian part of the self-energy matrix, and is the equivalent of the broadening term γ from the one level model

$$\mathbf{\Gamma} = i \left[\mathbf{\Sigma} - \mathbf{\Sigma}^\dagger \right] . \quad (2.17)$$

The trace term $\text{Tr} \left[\mathbf{\Gamma}_l \mathbf{G}_s^\dagger \mathbf{\Gamma}_r \mathbf{G}_s \right]$ gives the electronic transmission through the scattering region at energy E . We now have an expression for the current which is equivalent to equation 2.6 in the one level model, with the original single broadened energy level replaced by the Green's function. The final step is calculating the surface Green's functions of each lead. Since calculating the whole semi-infinite Green's function is not possible, another method is required exploiting the periodicity of the leads. Several such techniques have been developed [33–35], and the algorithm used in EHTransport will be described in chapter 3.

This procedure is now equivalent to the one level model at the point broadening was added, but is still missing charging effects. A density matrix ρ can be calculated simply by integrating the Green's function over energy, and plays the role of the level occupancy N from the one level model.

$$\rho = \frac{1}{2\pi} \int dE \mathbf{G}(E) . \quad (2.18)$$

As well as the density matrix, an equivalent of the charging energy U from the single level model is also required, to relate the density matrix to an energy shift. This is a more complex task, and its method of calculation depends on the method of obtaining the Hamiltonian. Once obtained however, its effects are identical to the one level model; we simply add the potential matrix to the Hamiltonian, so that the Green's function becomes

$$\mathbf{G}_s = [E^+ \mathbf{S}_s - \mathbf{H}_s - \mathbf{U} - \mathbf{\Sigma}_l - \mathbf{\Sigma}_r]^{-1} . \quad (2.19)$$

The Green's function now needs to be calculated self-consistently, as was the case with the one level model.

Although this provides a method of calculating the current, a method of constructing the input is still needed. Specifically the five repeated elements that make up the infinite Hamiltonian matrix, and their corresponding overlap matrices are required. This requires solving the Schrödinger equation for the lead and scattering regions, a many-body quantum problem. For this we will initially use density functional theory.

2.3 Density functional theory

2.3.1 Introduction

In physics there are many problems which do not have analytic solutions, and must be solved by numerical methods. One of the best known is the classical many body problem, where even three objects moving non-relativistically with a simple gravitational attraction between each object cannot be solved analytically in the general case. How much harder then must be an electronic structure problem, which even for small molecules may contain hundreds of interacting quantum particles, interacting in far more complex ways than simple Newtonian gravity. Here, even an attempt at a numerical solution to the Schrödinger equation will fail due to the sheer scale of the problem, and approximate theories are required. Density Functional Theory (DFT) [17, 36] is one of the most popular; although other methods such as Hartree-Fock [37] and quantum

Monte-Carlo [38] also exist, DFT tends to offer a good balance between calculation quality and the required computational effort [39]. DFT is a method that is interested in the total electronic density at each point in space, rather than attempting to obtain the many-particle wave-function directly.

2.3.2 The Born-Oppenheimer approximation

In an example molecular system, we have one or more nuclei, surrounded by an electron cloud, with no external interactions. This system has the Schrödinger equation

$$(\hat{T}_e + \hat{T}_n + \hat{U}_{ee} + \hat{U}_{nn} + \hat{U}_{ne})\Psi = E\Psi , \quad (2.20)$$

which includes the kinetic energy of the electrons and nuclei, an electron-electron interaction term, a nuclei-nuclei interaction term and a nuclei-electron interaction term. This is a function of $3(N_e + N_n)$ coordinates, and is analytically insoluble for any system more complex than a single hydrogen atom. Approximations must be sought to make any headway.

A first item of note is that the mass of the electrons is several orders of magnitude less than that of the nuclei, (even the case of a hydrogen atom, the lightest possible nuclei, the mass ratio between an electron and the nucleus is already near 2000 to 1,) and at a given temperature will have greater velocity. If we only desire to calculate electronic properties, a useful assumption we can make in this case is that there is a separation of the time and energy scales between these fast moving electrons and the much slower nuclei, and that the electrons instantly adapt to any movement of the nuclei. This is the Born-Oppenheimer approximation [40], and is equivalent to saying that the electron energies are determined by a static external potential due to the nuclei, which are now considered fixed classical particles for the purposes of calculating the electronic structure, removing three of the terms from our Schrödinger equation; \hat{T}_n , \hat{U}_{nn} and \hat{U}_{ne} . In exchange, we gain one additional term \hat{V} , describing the effect on the electrons of the static external potential. This approximation removes our ability to see effects due to quantum interactions between electrons and nuclei, such as electron-phonon scattering, but it reduces the total number of particles in the system, as well as ensuring that only a single particle type (the electron) exists in the system.

This separation of energies is easy to see in molecules, where electron excitation energies are

typically around two orders of magnitude larger than the vibrational modes of the nuclei, which are in turn another couple of orders of magnitude greater than rotational energies. In an insulator, the picture is muddled by the broadening of electron energy levels into bands within which there exist a continuum of states with arbitrarily small energy differences, but due to Pauli's exclusion principle preventing electrons moving into already occupied states, and without any temperature effects the only electronic excitations allowed will be across a band gap which is still large in comparison with nuclei energies, and so the separation still holds. In a conductor however there is no band gap, and arbitrarily small electronic excitations are allowed. In this case it is not obvious that the Born-Oppenheimer approximation should be valid, although in practice it is still generally used³, and provides reasonable results because at normal temperatures only a small number of electrons near the the Fermi surface are affected, and do not have a great deal of influence on the properties of interest of the system.

2.3.3 The Hohenberg-Kohn theorem

Applying the Born-Oppenheimer approximation, the nuclear related terms in the Schrödinger equation are replaced by a new external potential term V , leading to

$$(\hat{T} + \hat{U} + \hat{V})\Psi = E\Psi , \quad (2.21)$$

where \hat{T} is the operator for kinetic energy, \hat{U} the operator for the interaction between electrons, \hat{V} the operator for external potential and Ψ is a many-particle wave-function that describes all electrons in the system but no other particles. It is obvious from inspection of the Schrödinger equation that a given value of the external potential will lead to a wave-function which is a solution of the equation, which will in turn lead to some total electronic density. The basic premise of DFT is that the reverse is also possible; i.e. given a total electronic density there exists a functional with which one can calculate the external potential, (and hence the full Hamiltonian,) of a system. The Hohenberg-Kohn theorem [17], essentially the proof of the existence of this functional, is simple for non-degenerate and non-relativistic systems. Firstly, to

³Indeed, until recently its use was required for DFT, independent of any computational restraints. Moving beyond the Born-Oppenheimer approximation results in more than one particle type in the system, and requires a multi-component formulation of DFT, which was not well developed until the last ten years [41, 42] and is not discussed here.

prove that knowing the wave-function leads to a unique potential we assume that the opposite is true, that there exists some wave-function which is a solution to two different potentials:

$$\begin{aligned}(\hat{T} + \hat{U} + \hat{V})\Psi &= E\Psi \\(\hat{T} + \hat{U} + \hat{V}')\Psi &= E'\Psi .\end{aligned}\tag{2.22}$$

Subtracting these two equations gives:

$$(\hat{V} - \hat{V}')\Psi = (E - E')\Psi .\tag{2.23}$$

Since E and E' are simple numbers, the difference in the external potential must be the same except for a constant.⁴ Likewise, assuming that there exist two wave-functions differing by more than a phase term which each give rise to the same density η leads to

$$\langle \Psi | \hat{H} | \Psi \rangle = E < \langle \Psi' | \hat{H} | \Psi' \rangle = \langle \Psi' | \hat{H}' + \hat{V} - \hat{V}' | \Psi' \rangle = E' + \int d^3r \eta(\mathbf{r})(v(\mathbf{r}) - v'(\mathbf{r})) ,\tag{2.24}$$

where we've used the result from the first stage of the proof that if the wave-functions differ then so must the external potentials. Substituting the roles of the two wave-functions gives the similar expression

$$E' < E + \int d^3r \eta(\mathbf{r})(v'(\mathbf{r}) - v(\mathbf{r}))\tag{2.25}$$

and finally, adding these leads to the contradiction

$$E + E' < E' + E\tag{2.26}$$

Since the electronic density uniquely determines the Hamiltonian, it also determines the wave-function and all other ground states properties of the system such as energy. There must therefore exist some functional of the electron density $F[\eta]$ that maps the density to the ground

⁴Without further information on the nature of the operator \hat{V} , it is possible to conceive a system where Ψ is an eigenfunction of each of \hat{V} , \hat{V}' and the combination $\hat{V} - \hat{V}'$. The external potential operator is multiplicative however, so with the assumption that there's at least one particle in the system (and hence $\Psi \neq 0$) Ψ cancels and the proof holds.

state energy. This functional is currently unknown, and is almost certainly extremely complex since it encapsulates the whole of the interacting many-body problem. It can however be written more simply in terms of another functional $G[\eta]$ which maps the density to its wave-functions.

$$F[\eta] = \langle G[\eta] | \hat{T} + \hat{U} + \hat{V} | G[\eta] \rangle = \langle G[\eta] | \hat{T} + \hat{U} | G[\eta] \rangle + \int d^3r n(\mathbf{r})v(\mathbf{r}) . \quad (2.27)$$

This in itself is not useful since the functional $G[\eta]$ is also not known. It does however prove that the form of the functional $F[\eta]$ is independent of the potential, (since \hat{T} and \hat{U} are independent,) and also provides the stage for the proof of the second of the Hohenberg-Kohn theorems, which states that for a given function $v(\mathbf{r})$ (i.e. a known potential,) then the functional $F[\eta]$ is minimised by the 'correct' electronic density, as long as the number of particles in the system is kept constant.

$$F[\eta'] > F[\eta] \text{ for } \eta' \neq \eta . \quad (2.28)$$

This follows trivially from the Rayleigh-Ritz variational principle, since \hat{T} and \hat{U} are not dependent on the density, and so fixing the external potential fixes the whole Hamiltonian.

$$\langle \Psi | H | \Psi \rangle < \langle \Psi' | H | \Psi' \rangle . \quad (2.29)$$

These proofs can also be extended to finite temperatures [43], degenerate [44] or relativistic [45] systems, although the maths becomes correspondingly more complex. It's also a point of note that the exact forms of \hat{T} and \hat{U} do not matter in the proof, as long as their forms are not dependent on the density. At this point we could start with a guess for the total electronic density, perhaps just a sum of the densities of the atomic orbitals of isolated atoms, or from some simple tight binding model, and perform a minimisation to find the ground state density. Unfortunately, although the Hohenberg-Kohn theorem guarantees the existence of some density functional which can be minimised to find the ground state properties, it does not give any clues as to the form of the functional, or how to construct approximations to it.

2.3.4 The Kohn-Sham equations

Since the proof of the existence of the density functional does not depend on the exact forms of \hat{T} and \hat{U} , it is reasonable to imagine the existence of a hypothetical system that lacks inter-particle interactions, (and hence \hat{U} is always zero,) but where the density remains unchanged from the original problem. Since there is now no interaction between particles, the many-body problem vanishes and the system can be treated as a set of independent particles. The kinetic energy functional can be written exactly in terms of the single particle orbitals ψ_i ⁵, which are the first N solutions of a one particle Schrödinger equation.

$$T[\eta] = \frac{-\hbar^2}{2m} \sum_{i=1}^N \langle \psi_i | \nabla^2 | \psi_i \rangle . \quad (2.30)$$

The Schrödinger equation for this system contains only the single particle kinetic energy and the external potential, and no electron-electron interaction term.

$$\left(\frac{-\hbar^2}{2m} \nabla^2 + v[\eta] \right) \psi = E\psi . \quad (2.31)$$

And finally the density is just the sum of the squared magnitudes of the occupied single particle wave-functions.

$$\eta(\mathbf{r}) = e \sum_{i=1}^N |\psi_i|^2 , \quad (2.32)$$

where N is the number of electrons in the system.

It is then possible to construct a new effective external potential $v[\eta]$ such that the ground state energy and density matches that of the original interacting system. This is the Kohn-Sham scheme [46], which has proved successful at simplifying DFT calculations. In order for this to be useful as a method of solving the interacting electron problem, some way of creating this new effective potential is needed. We can start by writing down the system energy in its functional form

$$E[\eta] = T[\eta] + \int d^3r v_{ext}(\mathbf{r})\eta(\mathbf{r}) + \frac{1}{2} \int \int d^3r d^3r' \frac{\eta(\mathbf{r})\eta(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + E_{XC}[\eta] , \quad (2.33)$$

⁵Although this is exact for the non-interacting system, it is missing terms compared to the energy from the interacting system. The missing fraction is added back via the exchange-correlation functional.

where the first term is the kinetic energy, the second is the external potential from the original interacting system, the third is the classical Coulomb energy⁶, and the final term is the exchange-correlation functional containing everything else that is unknown about the system. The effective potential is required to result in this energy, along with the original density, and is obtained by applying the variational principle to the energy functional.

$$v[\eta] = v_{ext}(\mathbf{r}) + \frac{1}{2} \int d^3r' \frac{\eta(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta E_{XC}[\eta]}{\delta \eta} . \quad (2.34)$$

As this expression is both dependent on the density and required for the calculation of it, the equations must be solved self-consistently.

Since the exchange-correlation term E_{XC} is unknown, an approximation must be used. Fortunately, it is easier to construct these than for the original problem of interacting particles, and there are many well tested functionals in common use. The most basic is the local density approximation (LDA), which depends only on the electronic density at each point in space, and not on its derivatives or any other terms. It is most commonly derived from a homogeneous electron gas, resulting in an expression which is exact in the limit of zero electronic density gradient, and then applying that expression at each point in space.

$$E_{XC}[\eta] = \int d^3r \eta(\mathbf{r}) e_{xc}^{hom}(\eta(\mathbf{r})) , \quad (2.35)$$

where $e_{xc}^{hom}(\eta(\mathbf{r}))$ is the exchange-correlation energy density of a homogeneous electron gas of density $\eta(\mathbf{r})$.

This works surprisingly well, even in cases where the true electron density is far from homogeneous. It can be shown that the LDA is the first term in an expansion of the exchange-correlation energy in terms of the gradient of the density [46, 47], but later attempts to add in additional terms turned out to produce worse results than the LDA. There are various sum rules that the true exchange-correlation functional obeys, such as the fact that the exchange hole should displace exactly one electron, and on closer examination it turned out that although the LDA obeyed these rules adding in additional gradient terms was breaching them [48]. More advanced func-

⁶This term is no longer an electron-electron interaction, but is a functional of density. As such it contains a spurious self interaction between single electrons and themselves. It can be partially fixed by the exchange-correlation functionals, but due to its non-local nature to remove it completely requires a non-local functional, something which is generally avoided for efficiency reasons.

tionals such as PBE [49] were created to include terms dependent on the gradient of the density while still preserving the most important of these sum rules, the so-called generalized gradient approximation (GGA). There also exist even more complex functionals such as metaGGA class which includes terms dependent on the kinetic energy density, an exact-exchange functional, and hybrid functionals [50] such as B3LYP which add in part of the Hartree-Fock energy, but these require much greater computational resources.

While these Kohn-Sham states have no direct link to any single particle orbitals in the interacting problem, (which, as has been mentioned, are not physical and do not actually exist,) it is common to think of them as such. Similarly the energies of these orbitals are thought of as single particle energies in the interacting problem, and are commonly used for calculating excitation energies or plotting band structures. In fact, it turns out that these energies are indeed a reasonable approximation, and can give good band plots for the occupied states. Agreement with unoccupied states is typically not as good, however.

2.3.5 Pseudopotentials

There is one more important approximation that can be made before the scheme is used. As two distant atoms are brought closer together, it is their outermost electrons which interact with each other most strongly. As long as the atoms are not brought in too close, the core electrons remain relatively unperturbed from the isolated atom case. This suggests that instead of treating all electrons in the problem explicitly, we could 'freeze' the core electrons, merging them in with the orbital and making them a part of the static potential. The combined potential due to the combined nucleus and core electrons becomes the pseudopotential. This greatly reduces the size of the problem, as with all but the first two rows of the periodic table of elements more than half of the electrons are removed from the system. By constructing the pseudopotentials in such a way that they are slowly varying in space it also greatly improves efficiency when representing the Kohn-Sham orbitals in a plane wave basis set, as the tightly bound core electron wave-functions are normally quickly varying in space and require large numbers of plane waves to represent accurately (Figure 2.5). Furthermore, the orthogonality condition of valence electrons would force them to also curve sharply near the core electrons.

One downside from using pseudopotentials is that some information is lost from the final




Figure 2.5: Comparison of the $n = 2$ single particle wavefunctions computed using all-electron DFT and using a pseudopotential. Beyond a cutoff distance the two overlap exactly.

results. One important example is that it is no longer known what the total system energy is, although comparisons can still be made between a systems containing the same set of atoms in different configurations. Also it assumes that core electrons are completely unperturbed from the isolated atom case, which may not be a good approximation for systems under very high pressures, or which are strongly ionic. All the work here is done with norm-conserving⁷ pseudopotentials supplied with SIESTA.

2.3.6 Geometry optimisation

We now have a method of calculating the electronic properties of a system, but even though we are working within the Born-Oppenheimer approximation we may still want to allow the nuclei to move, particularly for finding the ground state configuration of atomic positions for problems such as discovering how a molecule is likely to bind to the leads. It would be possible to simply perform the same electronic structure calculation repeatedly with the nuclei in different positions, but there is a large many-dimensional space to explore and so the process would be inefficient. As an alternative, we can apply the Hellmann-Feynman theorem [51], which proves that we can calculate forces on the nuclei by performing the electronic structure calculation at a single point

$$\frac{\partial E}{\partial \lambda} = \int d^3r \Psi^*(\lambda) \frac{\partial \hat{H}_\lambda}{\partial \lambda} \Psi(\lambda) . \quad (2.36)$$

In this case the parameter λ that we are varying relates to the positions of the nuclei. The

⁷i.e. the integral of the wavefunction is kept the same as the all electron wavefunction.

full theorem is not actually required here; the only terms in the Kohn-Sham Hamiltonian with a dependence on the nuclear positions are the nuclei-nuclei interactions, and the interaction between the density and the external potential, so a simple derivative of the energy functional is sufficient. Explicitly including the nuclei-nuclei interaction term and dropping the insignificant terms gives

$$E[\eta] = \int d^3r \left\{ \eta(\mathbf{r}) \sum_i \frac{Z_i e}{|\mathbf{r} - \mathbf{R}_i|} + \frac{1}{2} \sum_{i \neq j} \frac{Z_i Z_j e^2}{|\mathbf{R}_i - \mathbf{R}_j|} + \dots \right\}. \quad (2.37)$$

This can be differentiated with respect to \mathbf{R}_i to give the force on nucleus i ⁸

$$\mathbf{F}_i = \frac{\partial E[\eta]}{\partial \mathbf{R}_i} = \int d^3r \left\{ \eta(\mathbf{r}) \frac{Z_i e}{(\mathbf{r} - \mathbf{R}_i)^2} + \frac{1}{2} \sum_j \frac{Z_j Z_i e^2}{(\mathbf{R}_j - \mathbf{R}_i)^2} \right\}. \quad (2.38)$$

While this form looks simple, as usual practical issues introduce further complications. Extra Pulay forces [52] are introduced if the movement of the nuclei affects the overlap of the basis functions, as is the case with any non-complete set of localised orbitals centred around the nuclei. There are also further terms introduced if the density used isn't the exact solution of the system, and we will generally be working with pseudopotentials, and hence moving the whole atomic core rather than the bare nuclei. Nevertheless, with appropriate corrections this does provide a method to locate a minimum in energy by moving the atoms in the direction of the forces until the forces become zero, or by defining a timestep, keeping track of the velocities of the nuclei and applying accelerations due to the forces, first principles Born Oppenheimer molecular dynamics (BOMD) simulations can be carried out.

A major issue with using this method to locate a ground state geometry is that without some idea of the ground state positions to start with, starting from a random atomic configuration is far more likely to reach a local minimum than the real ground state of a system. In the case of our systems, it would mean that although a geometry optimisation could be used to give you the height a molecule prefers to sit above a site, (e.g. hcp, fcc, bridge, etc.) the molecule would typically not move between sites. In that sort of calculation it is common to start with many randomised states, or to use finite temperature molecular dynamics together with a method to remove energy from the system.

⁸This needs to be multiplied by a sign term to account for the ordering of the nuclei.

2.3.7 Extensions

Since the Hohenberg-Kohn theorem ensures that *all* ground state properties are functionals of the electronic density, there should also exist functionals to extract individual spin densities or other order parameters. In practice these would be complicated, as would be exchange-correlation functionals encapsulating the effects of such symmetry breaking, so extensions to DFT in the form of additional densities are used instead. These range from the commonly used spin-DFT, which includes a $(\eta_{\uparrow} - \eta_{\downarrow})$ density in addition to the $(\eta_{\uparrow} + \eta_{\downarrow})$ term from standard DFT and is useful for magnetic or open shell systems, through to multi-body DFT which allows the inclusion of additional particle types and which is required for going beyond the Born-Oppenheimer approximation, to even an ab-initio description of superconductivity including order parameters as densities [53].

Implementations acting directly on the density rather than introducing some sort of virtual non-interacting particle are also available [54],⁹ (which has become known as orbital-free DFT to differentiate it from Kohn-Sham based schemes,) but since only fairly crude approximations to the real energy functional are known, (in particular, and in addition to exchange-correlation problems, the kinetic energy has no obvious expression in terms of the density,) it tends to be less accurate than Kohn-Sham based DFT. It is still used because, by its nature, it tends to scale linearly with atom count without further approximations, making it very efficient for large systems. Since there are now well developed linear scaling Kohn-Sham based implementations this advantage has been partially lost, however. The majority of DFT calculations performed today use the Kohn-Sham formalism to specify an equivalent problem of N non-interacting particles with orbital-free DFT mainly reserved for large scale metallic systems, which tend to be difficult to treat using linear scaling Kohn-Sham DFT, which often fails in the absence of a large band gap between occupied and unoccupied electronic states.

In the case of modelling electrical properties of systems, one extension that is likely to be important is time dependence. Traditional DFT asks the question of what the electronic density looks like in the presence of a static external potential, but does not answer what happens if that

⁹In fact, the first attempt at constructing this functional was made by Thomas and Fermi in 1927, before the Hohenberg-Kohn theorem had even put DFT onto a sound theoretical footing. This took the kinetic energy, exchange and correlation terms from the homogeneous electron gas in a manner equivalent to the Kohn-Sham version of the LDA, but proved too inaccurate for practical use. It could not reproduce atomic shells, gave divergent densities at the nuclei sites, and gave incorrect decaying behaviour away from the system.

potential is varying with time. IBM have reported transistors operating of speeds of 155GHz, which is equivalent an on-off-on transition time of 6.5 pico-seconds. A transistor operating at the upper limit of its operational frequency will never achieve a steady state, and Time Dependent Density Functional Theory (TDDFT) would need to be invoked in order to describe what happens as the transistor switches between states.¹⁰ There is a Runge-Gross theorem [55] which, in an analogous way to the Hohenberg-Kohn theorem, proves that the time dependent wavefunction of a system has a one to one mapping to the time dependent electron density of a system, given a time dependent external potential. Some work on combining TDDFT with NEGF to calculate time dependent transport behaviour has already been carried out [56].

2.4 Extended Hückel Theory

Extended Hückel Theory [26] (EHT) is a semi-empirical linear combination of atomic orbitals (LCAO) method for studying the electronic configuration of a structure. As the name suggests, EHT is an extension of an earlier (non-extended) Hückel theory [57–60], which was a simple LCAO method very restricted in its applicability, developed to study conjugated hydrocarbons. The 'extension' however had more to do with the range of systems of applicability than the actual mathematical technique, which is markedly different between the two. Hückel theory uses no explicit real-space orbitals and instead assumes that the overlap matrix is the identity matrix, and then the Hamiltonian is parametrised based on atomic type and position. By comparison, the defining feature of EHT is the assumption that the each element of the Hamiltonian matrix is directly proportional to the corresponding element of the (non-diagonal) overlap matrix.

$$\mathbf{H}_{ij} = \langle \psi_i | \hat{H} | \psi_j \rangle \approx c_{ij} \langle \psi_i | \psi_j \rangle = c_{ij} \mathbf{S}_{ij} , \quad (2.39)$$

where ψ are the functions of the basis set, which in comparison to the original Hückel theory must now have some definition that allows the calculation of these overlaps. There are a few variations on the form of the values c_{ij} , but they are always parameter based and do not depend on the properties of ψ , removing the requirement for any complex three or higher centre integrals, and

¹⁰This frequency of transistor operation is still substantially slower than typical molecular vibration frequencies, calling into question the usefulness of applying TDDFT to this case while still under the Born-Oppenheimer approximation. In using this, the assumption is made that effects of molecular vibration on the properties of the system are insignificant compared to the effects of, e.g., the widely swinging electric fields on a gate electrode.

limiting the two centre integrals to simple overlaps between basis functions, which can normally be efficiently calculated by analytical methods.

In a standard EHT calculation, the basis set chosen is a set of Slater orbitals positioned at the site of each atom, representing each of the valence states of the atom. The outermost (i.e. highest n) filled or partially filled orbitals of each angular momentum are used.

$$\phi_{nlm}(\vec{r} - \vec{R}_0) = (2\zeta_{nl})^n \sqrt{\frac{2\zeta_{nl}}{(2n)!}} |\vec{r} - \vec{R}_0|^{n-1} e^{-\zeta_{nl}|\vec{r} - \vec{R}_0|} Y_l^m(\vec{r} - \vec{R}_0), \quad (2.40)$$

where \vec{R}_0 is the location of the orbital and n , l and m are the usual quantum numbers. ζ_{nl} is the Slater exponent; an empirically chosen parameter. Like the orbitals which result from the use of pseudopotentials in DFT, these Slater orbitals match the form of the hydrogenic orbitals at high r , but are smoother at lower r , without nodes (Figure 2.6).

A second empirically chosen parameter I is assigned to each n and l combination of each atomic species, and can loosely be regarded as an ionisation energy from that orbital. Since normally only valence orbitals are considered, there will not be two orbitals for one atomic species with the same l but different n , leaving $2\max(l)$ empirical parameters per atomic species.¹¹ The diagonal elements of \mathbf{H} will have $C_{ii} = I_i$. The c_{ij} of the off-diagonal elements will have an average of I_i and I_j , which may be either arithmetic or geometric, multiplied by a further scaling factor k , which is known as the Wolfsberg-Helmholtz constant, and is usually chosen to have a fixed value of 1.75.

$$\mathbf{H}_{ii} = I_i \quad (2.41)$$

$$\mathbf{H}_{ij} = k\mathbf{S}_{ij} \frac{I_i + I_j}{2} \quad (2.42)$$

or $\mathbf{H}_{ij} = k\mathbf{S}_{ij} \sqrt{I_i I_j}$.

The parameters are chosen to match some behaviour, such as band structure, energy levels of some molecule, etc.

The transferability of the EHT parameters can be an issue, for example a carbon atom in the backbone of an octane molecule is likely to behave differently to a carbon atom in a slice

¹¹From inspection of equation 2.40 it can be seen that the overlap between two orbitals positioned at the same site with different n but equal l will not be zero, so in any case where additional core or empty levels are included special care needs to be taken.




Figure 2.6: Comparison of the Slater radial wave-function with the hydrogenic one for $n = 3$, $l = 0$. The angular part is identical in both cases.

of graphene. One major source of these errors is the charge transfer that occurs between atoms of different electronegativities in a molecule. To some extent, this can be taken care of by the parameter choice; since fluorine atoms typically become negatively charged, their ionisation energy can be adjusted appropriately. This has negative effects on transferability though; e.g. a carbon bonded to the fluorine will likely have a more positive charge than if it were bonded to a hydrogen, and the same ionisation potential could not be used in both cases. One solution to this problem is to employ an iterative scheme, in which corrections to the Hamiltonian are calculated at each step, and the new Hamiltonian used to compute new corrections, until self-consistency is reached. In one simple self-consistent Extended Hückel Theory scheme [30,31], the on-site energy of each orbital is assumed to be a linear function of the charge Δq on the atom upon which the orbital is centred, adding one additional empirical parameter α per atomic species (The s, p and higher orbitals are assumed to share the same value of α) .

$$I_i = I_i^0 + \Delta q_i \alpha_i . \quad (2.43)$$

Many schemes are available for estimating the atomic charges, from the simple Mulliken population analysis [61] to natural charges [62] to integrating the electronic charge density in real space. Other schemes for calculating the corrections to the ionisation potential I_i also exist, ranging from the simple addition of additional parameters multiplied by higher powers of Δq_i to more complex Poisson solvers [63] or NEGF methods [64].

2.5 Periodic systems

The DFT and EHT equations described so far have been in the general terms of a system of unspecified composition and size. In the NEGF formalism described, the leads at least are *required* to have an infinite extent in the direction of conduction, and both the leads and scattering regions have no bounds on their size in directions orthogonal to conduction. Obviously a real calculation on an infinite system cannot be conducted directly, so an alternative is required. In this work, all systems studied that have infinite extent along one or more axis are assumed to be completely periodic along those axis; i.e. that we can define a single finite unit cell, and then build up the full infinite system by repeating it periodically¹². We can therefore apply the concept of Bloch waves [65] to obtain an approximation to the infinite Hamiltonian.

Bloch's theorem states that the wave-function of an infinite periodic system can be written as the product of two functions, one with periodicity equal to the system under study, and the other a plane wave. First we define a set of Hamiltonians \mathbf{H}^n where n is a label for the periodic images of the unit cell, and \mathbf{H}^0 is equivalent to the non-periodic Hamiltonian.

$$\begin{aligned}\mathbf{S}_{ij}^n &= \langle \psi_i^0 | \psi_j^n \rangle \\ \mathbf{H}_{ij}^n &= \langle \psi_i^0 | \hat{H} | \psi_j^n \rangle .\end{aligned}\tag{2.44}$$

A second set of Hamiltonians \mathbf{H}^k can then be defined, where \vec{R}_n is the vector pointing from the original unit cell to unit cell n , and $\phi_{\vec{k}}$ is a (possibly complex) phase factor.

$$\begin{aligned}\mathbf{S}'^{\vec{k}} &= \sum_n \phi_{\vec{k}}(\vec{R}_n) \mathbf{S}^n \\ \mathbf{H}'^{\vec{k}} &= \sum_n \phi_{\vec{k}}(\vec{R}_n) \mathbf{H}^n .\end{aligned}\tag{2.45}$$

The \vec{k} here is a k-point, describing a point in the reciprocal lattice vector. The k-point defines the plane wave part of the wave-function, so the more k-points used the more Bloch states are included in the calculation. This is a controlled approximation; as the range of n and as the

¹²We have in fact already made this assumption implicitly for the leads in the direction of conduction by the existence of the lead 'principal layer' in the NEGF formalism.

density of k-points tends to infinity, results from the set of Hamiltonians tends towards the results which would be obtained from the original infinite Hamiltonian. Since it is not computationally feasible to sample an infinite number of k-points, a subset needs to be chosen which provides a good approximation for the number of points. Algorithms such as Monkhorst-Pack [66] exist for picking this set of k-points. It is also obvious from inspection that the Hamiltonian at the Γ point (i.e. at $\vec{k} = 0$) of a system in which interactions between the periodic images are negligible reduces back to the original form.

The matrices at each k-point have a number of elements equal to what would exist in the Hamiltonian of a single isolated periodic cell. Since the calculation time scales linearly with the number of k-points, but as $\mathcal{O}(N^3)$ with the number of orbitals involved in the system, for a given calculation quality it is vastly faster to obtain e.g. a density of states from a minimal unit cell and many k-points than it is to obtain the same thing by performing a Γ point calculation with a much larger unit cell.

Chapter 3

Description of Code

3.1 Software

Two software packages capable of computing transmission spectra were used; SIESTA [67–70], (Spanish Initiative for Electronic Simulations with Thousands of Atoms,) and EHTransport. SIESTA is an MPI-based DFT code which uses a numerical localised basis set. This set has a strictly limited range; i.e. rather than the exponentially decaying tail that may be expected for a localised atomic orbital, the basis set is constructed to be exactly zero beyond a certain cutoff distance [71,72]. This is advantageous for the NEGF calculations, as the computed Hamiltonian is already in a localised form which makes the calculations straightforward, and the strict cutoff ensures that the matrix is sparse, ensuring the existence of a principal layer that doesn’t require neglecting any matrix elements, and improving the efficiency of the matrix operations involved in the Green’s function calculations. SIESTA also includes the subcomponents Transiesta and TBTrans, which perform the required Green’s function calculations [64,73].

EHTransport is my own custom written OpenMP based code which uses the semi-empirical Extended Hückel Theory (EHT) to compute the Hamiltonian of a given system in a localised basis set of Slater orbitals, and then uses the NEGF method to compute transmission spectra. It is written in C++, and currently totals approximately 6000 lines of code. It is described in detail in the following sections.

3.2 EHTransport

EHTransport is divided up into 3 main modules, along with a number of specialised utility modules. The first is the input/output module, responsible for reading user input and using it to build the parameters for other modules, calling them, then processing the raw results obtained into useful output files. The second is the EHT module, which takes a list of atomic species and positions as input and delivers as its primary output a Hamiltonian and overlap matrix, along with the estimated Fermi level of the system. The third module is responsible for NEGF calculations. It accepts as input a set of three Hamiltonian and overlap matrices, representing the leads and scattering region, along with the Fermi levels of the leads and the potential difference between them. It provides as output a transmission spectra and density of states in the scattering region. Additional cross links are available between the EHT and NEGF

module to allow charge consistent EHT calculations for the scattering region, but this code is not currently active. Utility modules provide representations of collections of atoms and the means to perform manipulations such as rescaling or the construction of the super-cells. This strict subdivision allows modification and testing of one part of the code without risking breaking another. The aim of the application is to provide results which, given an appropriately chosen parameter set, are in reasonable agreement with Transiesta. This might be for investigation of systems too large to be accessible to Transiesta, or for large numbers of smaller systems, the most interesting of which may warrant further study with far more computationally intensive Transiesta.

3.2.1 The input/output module

EHTransport’s input/output module is responsible for reading an input file from disc, reading a job type and parameters from that file, and then setting up and running the appropriate calculation. Once the calculation is complete, it also takes the raw data obtained and processes it to create useful output files. The input files use the same FDF format as SIESTA, including using many of the same keywords. In most cases, this overlap, combined with the flexibility of the FDF format that ignores unrecognised keywords, allows the same input file to be processed with both SIESTA and EHTransport, which is a useful feature for performing comparisons. FDF is a plain text, human readable format, allowing for visual inspection of scripts. Full details of the format are available in appendix 1, or in the SIESTA documentation.

3.2.2 The EHT module

The EHT module takes as its primary input a list of atomic species and positions, and the size of the cell containing the calculation. It also requires a number of parameters: the k-point grid on which to perform the calculation, the size of the super-cell to use for periodicity, a cutoff range beyond which orbital overlaps are assumed to be zero, the type of self-consistency to use and the criteria for deciding when convergence has been reached. It also requires a set of EHT parameters for each atomic species involved in the calculation. Its main task is to provide overlap and Hamiltonian matrices to represent the system. (Or more than one, in the case of a periodic calculation with multiple k-points.) It can also provide the system Fermi level, projected and

total density of states, self-consistent atomic charges, band structures, and a representation of charge movement and potential in real space.

Computing the overlap matrix

The first step is the computation of an overlap matrix \mathbf{S}_{ij}

$$\mathbf{S}_{ij} = \langle \phi_i | \phi_j \rangle . \quad (3.1)$$

EHTransport exclusively uses Slater orbitals [74] centred around each atom in the system as its basis set. Tesseral harmonics are used for the angular part, rather than the spherical harmonics.

$$\phi_{nlm}(\vec{r} - \vec{R}_0) = (2\alpha)^n \sqrt{\frac{2\zeta}{(2n)!}} |\vec{r} - \vec{R}_0|^{n-1} e^{-\alpha|\vec{r}-\vec{R}_0|} Z_{lm}(\vec{r} - \vec{R}_0) , \quad (3.2)$$

where \vec{R}_0 is the location of the atom, n , l and m are the usual quantum numbers, Z_{lm} are the tesseral harmonics and α is the Slater exponent; an empirically chosen parameter. As the basis set is normalised, the diagonals are taken to be 1. Slater orbitals with the same centre and differing l or m are orthogonal, so overlaps of different orbitals centred around the same atom are taken to be 0. A distance cutoff may also be applied, with overlaps between orbitals centred around atoms separated by more than a specific distance also taken as 0. If none of these rules apply, the overlap can be calculated for arbitrary atomic numbers either via a fast analytic method, [75] or, for testing purposes, code also exists for calculating them via much slower volume integrals for arbitrary n and $l \leq 2$. Since EHTransport uses tesseral harmonics for the angular part of the orbitals, the overlap matrix is real and symmetric. The overlap matrix is also positive definite, and so care needs to be taken with applying a distance cutoff as it can destroy this property.

The original expression for the analytic overlaps is given by the system of equations from 3.3 to 3.12, where n_i , l_i and m_i are the usual quantum numbers, α_i is the Slater exponent of the orbital, and \vec{R} the vector joining the first orbital centre to the second. Y_l^m are the spherical harmonics. u and v are the transformed coordinates $u = (\alpha_2 + \alpha_1)|\vec{R}|/2$ and $v = (\alpha_2 - \alpha_1)|\vec{R}|/2$. The quantities in square brackets in equations 3.3 are the 3-j coefficients and the bracketed quantities in equation 3.10 represent the binomial coefficients. In equation 3.6 any non-integer

values resulting from the divisions by 2 are rounded towards zero before performing the factorial.

$$I(n_1, l_1, m_1, \alpha_1, n_2, l_2, m_2, \alpha_2, \vec{R}) = (-1)^{m_1} \sum_{\Lambda} A_{\Lambda}(\dots) \begin{bmatrix} l_1 & l_2 & \Lambda \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} l_1 & l_2 & \Lambda \\ m_1 & -m_2 & m_2 - m_1 \end{bmatrix} Y_{\Lambda}^{m_1 - m_2}(\vec{R}) , \quad (3.3)$$

$$A_{\Lambda}(n_1, l_1, m_1, \alpha_1, n_2, l_2, m_2, \alpha_2, \vec{R}) = \sum_{\lambda, L} F_{\Lambda\lambda L}(\dots) Q_{\lambda L}(\dots) , \quad (3.4)$$

$$F_{\Lambda\lambda L} = (-1)^{l+\lambda} (2L+1) \sqrt{4\pi(2\Lambda+1)(2l_1+1)(2l_2+1)} \frac{\Phi_{l_2 L \gamma} \Phi_{L \Lambda(l_1-\lambda)}}{\Phi_{\Lambda l_2 l_1}} , \quad (3.5)$$

$$\Phi_{ijk} = \frac{((i+j+k)/2)!(i+j-k)!}{((i+j-k)/2)!((j+k-i)/2)!((k+i-j)/2)!(i+j+k+1)!} , \quad (3.6)$$

$$Q_{\lambda L}() = \sum_{k=0}^{L/2} \sum_{p=0}^{L-2k} D_{Lkp} |\vec{R}|^{l_1-\lambda+L-2k-2p-1} J(n_2 + \lambda + 2k - L, n_1 - l_1, p) , \quad (3.7)$$

$$D_{Lkp} = \frac{(-1)^k (2L-2k)!}{2^{2L-2k} k! (L-k)! p! (L-2k-p)!} , \quad (3.8)$$

$$J(M, N, p) = |\vec{R}|^{M+N+2p+2} \sum_{\mu=0}^{M+N} B_{MN\mu} F_{\mu+p}(u) G_{M+N-\mu+p}(v) , \quad (3.9)$$

$$B_{MN\mu} = \frac{1}{2^{M+N+1}} \sum_b (-1)^{N-\mu+b} \binom{M}{b} \binom{N}{\mu-b} , \quad (3.10)$$

$$F_s(u) = s! \frac{e^{-u}}{u^{s+1}} \sum_{i=0}^s \frac{u^i}{i!} , \quad (3.11)$$

$$G_s(v) = \frac{s!}{2v^{s+1}} \left[e^v \sum_{i=0}^s \frac{(-v)^i}{i!} - e^{-v} \sum_{i=0}^s \frac{v^i}{i!} \right] . \quad (3.12)$$

All summations without explicit limits are bounded by the range of possible values, e.g. by

the range of values for which the 3-j coefficients are defined in the case of Λ . These are given by equations 3.13 to 3.16. In the case of equation 3.12, the summation of many terms of opposite signs can lead to numerical instability. There are 3 possible outcomes: For reasonably large v , the solution of equation 3.12 can be computed without difficulty. For smaller v , as shown in [75] the summation can be recast into the open form given in equation 3.17. Although theoretically an infinite series, for $|v| \ll 1$ convergence will be reached very quickly. (Using a standard EHT parameter set the limits of 64-bit floating point accuracy are typically reached within 20 terms for any combination of orbitals.) The final special case is for $v = 0$, for which both expressions as written reduce to $0/0$. This case is very common, occurring in any system with more than one atom of the same atomic species. (Hence giving rise to overlaps for which $\alpha_1 = \alpha_2$.) For this special case, equation 3.18 is used. No numerical difficulties arise for the case of small but non-zero v , for which the sum in equation 3.17 reduces to a single term. One final source of numerical divergence is the case of \vec{R} approaching zero, which which this system of equations fail. Two atoms could not approach arbitrarily closely in a real system, but there is still the case of multiple orbitals centred on the same atom to consider. The solution to this case is trivial and given in equation 3.19,¹ but these cases need to be treated as a special case in any computation.

$$\max(|l_1 - l_2|, |m_1 - m_2|) \leq \Lambda \leq l_1 + l_2 , \quad (3.13)$$

$$0 \leq \lambda \leq l_1 , \quad (3.14)$$

$$\max(|l_2 - \lambda|, |\Lambda + \lambda - l_1|) \leq L \leq \min(|l_2 + \lambda, \Lambda + l_1 - \lambda|) , \quad (3.15)$$

$$\max(0, \mu - N) \leq b \leq \min(\mu, M) , \quad (3.16)$$

$$G_s(v) = \frac{s!}{2v^{s+1}} \left[e^{-v} \sum_{i=s+1}^{\infty} \frac{v^i}{i!} - e^v \sum_{i=s+1}^{\infty} \frac{(-v)^i}{i!} \right] , \quad (3.17)$$

$$G_s(0) = \delta_{s \bmod 2, 1} \frac{1}{s+1} , \quad (3.18)$$

¹In actuality, unlike standard atomic orbitals, the overlap between multiple Slater orbitals centred on the same site, of different n but equal l and m is not 0. Usually, EHT calculations are not performed in which a single atomic species has multiple orbitals of the same l and different n , so this case isn't important. In a case where they are used, however, the overlap should still be set artificially to 0.

$$I(n_1, l_1, m_1, \alpha_1, n_2, l_2, m_2, \alpha_2, 0) = \delta_{n_1 n_2} \delta_{l_1 l_2} \delta_{m_1 m_2} . \quad (3.19)$$

Although it was stated that EHTransport utilises tesseral harmonics, the equations as written are in terms of the spherical harmonics. It is preferable to recast the results into a real form, in order to obtain real and symmetric overlap and Hamiltonian matrices at the Γ point. Using spherical harmonics would result in complex Hermitian matrices, and the overlap matrix is still positive definite, but the computational cost of further stages of the calculation would be increased by as much as a factor of four. It is also desirable to know the exact real space representations of the orbitals in use, to allow for visualisation of any real-space quantities following a calculation. For this it is necessary to carefully define the spherical polar coordinate transformations and other manipulations to remove any ambiguity. In our case, $\theta = \arccos(\vec{R}_z/|\vec{R}|)$ and $\phi = \arctan(y/x)$.² The transformation from spherical to tesseral harmonics, and hence from a complex to a real basis set is then given by equation 3.20.

²In programming terms, this is implemented as `atan2(y, x)`, to avoid uncertainties over quadrants.

$$I_{\text{real}}(m_1, m_2, \dots) = \text{sign}(m_2)(-1)^{l_1+l_2+m_1+m_2} I_c(m_1, m_2, \dots) \text{ where}$$

$$I_c(m_1, m_2) = \begin{cases} m_1 = 0, m_2 = 0 \rightarrow I(m_1, m_2) \\ m_1 > 0, m_2 = 0 \rightarrow (I(m_1, m_2) + (-1)^{m_1} I(-m_1, m_2))/\sqrt{(2)} \\ m_1 < 0, m_2 = 0 \rightarrow (I(-m_1, m_2) - (-1)^{m_1} I(m_1, m_2))/(\sqrt{(2)}i) \\ m_1 = 0, m_2 > 0 \rightarrow (I(m_1, m_2) + (-1)^{m_2} I(m_1, -m_2))/\sqrt{(2)} \\ m_1 > 0, m_2 > 0 \rightarrow (I(m_1, m_2) + (-1)^{m_1} I(-m_1, m_2) + (-1)^{m_2} I(m_1, -m_2) + (-1)^{m_1+m_2} I(-m_1, -m_2))/2 \\ m_1 < 0, m_2 > 0 \rightarrow (I(-m_1, m_2) - (-1)^{m_1} I(m_1, m_2) + (-1)^{m_2} I(m_1, -m_2) - (-1)^{m_1+m_2} I(m_1, -m_2))/2i \\ m_1 = 0, m_2 < 0 \rightarrow (I(m_1, -m_2) + (-1)^{m_2} I(m_1, m_2))/(\sqrt{(2)}i) \\ m_1 > 0, m_2 < 0 \rightarrow (I(m_1, -m_2) + (-1)^{m_1} I(-m_1, -m_2) - (-1)^{m_2} I(m_1, m_2) - (-1)^{m_1+m_2} I(-m_1, m_2))/2i \\ m_1 < 0, m_2 < 0 \rightarrow -(I(-m_1, -m_2) - (-1)^{m_1} I(m_1, -m_2) - (-1)^{m_2} I(-m_1, m_2) + (-1)^{m_1+m_2} I(m_1, m_2))/2 \end{cases} \quad (3.20)$$

One final point is that standard EHT parameter sets commonly use double-zeta orbitals, especially for $l = 2$ orbitals in the transition metals. Mathematically, this is equivalent to a linear combination of two standard orbitals, each with different α but otherwise identical (Equation 3.21). The coefficients are usually, but not always, normalised such that the overlap of the double zeta orbital with itself is 1 (Equation 3.22, where the one of the two results in which $0 \leq c_b \leq 1$ is taken.). In the case that they're not, the overlap is required to be less than 1, and is equivalent to a normalised triple-zeta orbital where $\alpha_c \rightarrow \infty$. (In fact, in many cases a single zeta orbital with a value of $c_a < 1$ is used, equivalent to a double-zeta orbital where $\alpha_b \rightarrow \infty$.) There is no need to perform the whole calculation twice, as the only α dependence is in equations 3.11 and 3.12, which only appear within the summation in equation 3.9. The result of multiple-zeta

orbitals is a modification of B into equation 3.23.

$$I_{2\zeta}(\dots, \alpha_{1a}, \alpha_{1b}, c_a, c_b) = I(\dots, \alpha_{1a})c_a + I(\dots, \alpha_{1b})c_b , \quad (3.21)$$

$$c_b = \frac{1}{2} - b \pm \sqrt{b^2 + 4 - 4c_a^2} \text{ where } b = c_a \frac{2^{2n+2}(\alpha_a \alpha_b)^{n+0.5}}{(\alpha_a + \alpha_b)^{2n+1}} , \quad (3.22)$$

$$J(M, N, p) = |\vec{R}|^{M+N+2p+2} \sum_{\mu=0}^{M+N} B_{MN\mu} \sum_{i_1, i_2} F_{\mu+p}(u_{i_1, i_2}) G_{M+N-\mu+p}(v_{i_1, i_2}) c_{1i} c_{2i} . \quad (3.23)$$

As a first observation, although this method is valid for arbitrary l , most EHT calculations will have a small maximum value of n and l , hence also limiting the range of m and many of the summations. With these limits, the number of 3-j coefficients and results of equations 3.8, 3.10 and 3.11 will be heavily restricted. Since the possible values are known, they can all be calculated beforehand and stored. This immediately removes one level of the nested summations and removes some computational complexity, at the expense of a very small amount of extra storage. For the case of $l \leq 2$ and double precision numbers, caching the 3-j coefficients requires 9000 bytes of storage, D requires 600 bytes, F requires 5400 bytes, and B requires 6552 bytes. Compared to the storage requirements for the Hamiltonian matrix of any non-trivial system, this required storage is negligible.³

To improve the performance of these overlap calculations further, a few other modifications can be made. The 3-j coefficients are always used in fixed pairs, so when generating the cache we can save the product of each pair rather than the true values at no cost in additional storage. In a real calculation, we always want to use all possible m values for each l , which has several implications. Firstly it reduces the cost of the transformation to a real basis set: Whereas for the calculation of a single real overlap it would be necessary to compute up to four complex overlaps, these same four complex overlaps are used for the calculation of three other real overlaps. Thus for efficiency, the overlaps are calculated in the complex basis set for all possible m_1 and m_2

³EHTransport actually over allocates space here to ease performing the lookups. e.g. no use is made of the fact that in a 3-j coefficient $\Lambda \geq |l_1 - l_2|$. Allocating a larger than required block ensures that accesses can be made by a simple chain of multiply-adds, rather than requiring a more complex calculation to locate the cached value.

combinations and stored, and then transformed into the real basis together. We can also note that equation 3.4 does not depend on m , and since in a real calculation all orbitals of the same atomic species, n and l will have the same Slater coefficient, the results of equation 3.4 can be computed once for all of $0 \leq \Lambda \leq l_1 + l_2$, and then used for all m combinations.

The procedure is multithreaded by assigning different atoms to different threads. It is difficult to divide them up fairly at the start of a calculation, as different atoms have differing numbers of orbitals, and need to be overlapped with differing numbers of other atoms. (Both because of differing numbers of neighbours, and because the symmetry of the overlap matrix means that if we calculate an overlap $\langle i|j \rangle$ we do not want to waste time by also calculating $\langle j|i \rangle$.) The collection of atoms is therefore broken up into a queue of fairly small chunks, with each thread taking a new chunk from the queue as soon as it finishes the previous. This requires a small amount of extra bookkeeping overhead, but ensures that processors are not left idle.

The calculation of the overlap matrix is naturally an $\mathcal{O}(N^2)$ operation where N is the number of orbitals in the system. This reduces to $\mathcal{O}(N)$ with the application of a cutoff beyond which overlaps are assumed to be zero, and appropriate ordering of data structures. For an isolated system, this is the final overlap matrix of the system.

Periodic structures

If the system under study is periodic, further work needs to be done. Rather than computing a single overlap matrix which holds overlaps between all atoms in the system, (where 'the system' in the periodic case is a single unit cell in an infinite periodic lattice,) additional overlap matrices of the same dimension are constructed, each one representing the overlap between atoms in the original cell with each atom in a different cell. Since a cutoff distance is applied beyond which it is assumed that overlaps between orbitals will be zero, it must also be true that for cells more than some distance away there will be no atoms within the cell that have a non-zero overlap with any atom in the original unit cell. This provides a limit on how many of the matrices are required to be constructed. In the limit of a large unit cell, (i.e. one which is greater in all dimensions than the cutoff distance,) only the nearest neighbour cells are required, giving 27 overlap matrices in total. This set of real space overlap matrices are not used directly, but rather are summed, with each multiplied by a phase factor to generate a final matrix (Equation

3.24). These phase factors are dependent on the point in k-space (the k-point, \vec{k}) at which the calculation is performed. Aside from at the Γ point, (at which the phase factor for all real space matrices in the summation is 1,) these phase factors are complex, and hence the overlap and Hamiltonian matrices are no longer real in general, but are still guaranteed to be Hermitian and, and \mathbf{S} remains positive definite.

$$\mathbf{S}_{ij} = \sum_{\alpha} \mathbf{S}_{ij}^{\alpha} e^{i(\vec{k} \cdot \vec{R}_{\alpha})} \quad (3.24)$$

For the Γ point, ($\vec{k} = 0$) and with a unit cell large compared to the distribution of atoms it contains, (i.e. there are no overlaps even between nearest neighbour unit cells, and the number of chosen unit cells is one,) this again reduces to the original non-periodic result. If only one of those two conditions are true the resulting matrices are real, but in the general case they are complex.

The k-points are input to and output from EHTransport as vectors scaled by the reciprocal unit cell lattice vectors, (k' i.e. as numbers between 0 and 1. The vector \vec{k} from equation 3.24 is calculated by permuting the x, y, z variables in

$$\vec{k}_x = k'_x \left(\frac{2\pi(\vec{c}^y \times \vec{c}^z)}{\vec{c}^x \cdot (\vec{c}^y \times \vec{c}^z)} \right) \quad (3.25)$$

where \vec{c}^{xyz} are the cell unit vectors.

Computing the Hamiltonian matrix and charge self-consistency

The second step is the calculation of a Hamiltonian matrix. In EHT, elements of the Hamiltonian matrix are taken to be proportional to the corresponding elements of the overlap matrix. The diagonal elements of the Hamiltonian are made up of an empirical energy I_i parameter assigned to each basis function added to a self-consistently calculated perturbation term. The off-diagonals are multiplied by the average of the energy parameters of basis functions i and j . The Wolfsberg-Helmholtz constant is not actually taken to be a constant in EHTransport, but instead is a property of each orbital K_i , and the average of the two are used. Since these factors are all real numbers, and since the diagonal scaling factors are symmetric, the symmetry of the overlap matrix is maintained; i.e. the Hamiltonian matrix is real and symmetric in an isolated system or

at the Γ point of a periodic system, and Hermitian otherwise. The Hamiltonian matrix elements are therefore

$$\begin{aligned}\mathbf{H}_{ii} &= I_i + \Delta I_i \\ \mathbf{H}_{ij} &= \frac{K_i + K_j}{2} \mathbf{S}_{ij} \frac{\mathbf{H}_{ii} + \mathbf{H}_{jj}}{2},\end{aligned}\tag{3.26}$$

where ΔI is a modification to the energy parameter due to charging effects. EHTransport implements three methods of finding these perturbation parameters ΔI , or alternatively can fix them to zero, in which case the calculation reduces back to standard non-self-consistent EHT. The first of the three supported methods is based on standard SC-EHT, as described previously. The other two are more specialised, and are able to include the effects of a bias across the system.

The standard SC-EHT implementation uses a Mulliken population analysis to construct the population matrix. In a departure from the original, it adds a term quadratic in the atomic charges, as well as the linear term

$$\mathbf{H}_{ii} = I_i + \alpha_i \Delta Q_i + \beta_i \Delta Q_i^2,\tag{3.27}$$

where ΔQ_i is the excess charge on the atom to which orbital i belongs, calculated by using the Mulliken population analysis to count up the number of electrons on the atom, and then subtracting off the number of electrons that would be expected if the atom was isolated and neutral. The parameters α and β are chosen per orbital, as opposed to per atom. It is the simplest and most efficient of the available self-consistent methods, but does not allow for the application of a bias across the system⁴.

The Mulliken population analysis is carried out by first solving the Hermitian-definite generalised eigen-problem to obtain the eigenvalues and eigenvectors. The available states are filled up according to the number of electrons in the system

⁴An additional potential dependent on the coordinates of each atom could be added, but this would fail to account for screening effects.

$$\begin{aligned}
Q_i &= \mathbf{P}_{ii} \\
\mathbf{P} &= \mathbf{S}\mathbf{D} \\
\mathbf{D}_{uv} &= \sum_{i=0}^n \mathbf{Z}_{ui} \mathbf{Z}_{vi}^\dagger F(E_i) \\
\mathbf{S}\vec{E} &= \mathbf{Z}\mathbf{H}\vec{E} \ ,
\end{aligned} \tag{3.28}$$

where Q_i is the occupation of orbital i and $F(E_i)$ is the occupation of a level with energy E_i : 1 if $E_i < E_f$ or 0 otherwise.

The statement that ‘available states are filled up according to the number of electrons in the system’ glosses over a fairly important problem, common in this type of self-consistent scheme. It is perfectly possible to fill up states from the lowest energy and work your way upwards until you run out of electrons as described above, but doing so is only practical for systems with a large HOMO-LUMO gap. If the gap is smaller then the charge redistribution, and resulting change in the value of ΔI , between two iterations can cause two states near the Fermi level to cross, and hence on the next iteration one or two electrons will jump from one to the other. The newly occupied state will likely increase in energy in the next iteration, causing the states to cross back, and the charge to reverse. This will continue to flip-flop forever, destroying any hopes of reaching convergence with reasonable mixing parameters. This issue is solved by filling up states around the Fermi level with non-integer numbers of electrons. In EHTransport, the usual Fermi function is used, with the Fermi level picked to fix the amount of charge in the system (usually at zero.). With this modification to $F(E_i)$, the crossing of energy of two nearby levels has little effect on the occupancy of each level, removing the sudden shifts from the system. The temperature T is somewhat artificial, and depending on the behaviour of the system under study around the Fermi level, any temperature from 0K for large gap insulators to 1000K for metals may be reasonable. The Fermi function is

$$f(E_i) = \frac{2}{e^{(E_i - E_F)/kT} + 1} \ . \tag{3.29}$$

Whereas the standard SC-EHT computes a population matrix via Mulliken population analysis, it is also possible to compute it using NEGF, which allows for bias across the system and

incorporates the effects of the leads, but has a much greater computational workload. From the population matrix a set of atomic charges is computed, and from there the procedure is the same as before. The exact NEGF procedure overlaps almost completely with that of calculating the DOS, and so is described in a later section. Although code exists for this procedure in EHTransport, it is not well tested and is currently disabled: Early testing showed that the speed penalty was too severe to mesh with EHTransport’s design goals.

In the final method, a charge on each atom is computed via Mulliken population analysis as in standard SC-EHT. Rather than assuming a polynomial relationship between these charges and ΔI , a real-space charge distribution is created by placing a function with a Gaussian radial component and a constant angular component on top of each atom in the system. The Poisson equation is solved in three dimensions, using boundary conditions that fix the potential at all points along the boundary of the unit cell. In the direction of conduction, the potential is fixed at $\pm V/2$, and in the other directions $-v/2 + z(v/zmax)$. If the system is periodic, the boundary conditions in directions orthogonal to conduction are updated at each step, setting the potential at each point on the boundary to the average of the point directly next to that boundary and the point next to the boundary on the opposite side of the cell. The solution of Poisson’s equation gives us a potential at each grid point, and the ΔI values are taken to be the potential at the atomic sites, linearly interpolated from the eight neighbouring grid points. The efficient Poisson solver in the Intel MKL library was employed for carrying out this calculation, but carries the restriction that the unit cell must be orthogonal. Since a single Gaussian basis function is used for each atom this setup does not allow for polarisation of the charge around an atom, but it does allow for charge rearrangement between atoms, and for screening effects. This method was developed in [63].

Whatever the method of calculating ΔI , using the output of one cycle directly as input to the next will rarely reach a state of convergence. Instead, it is necessary to ‘mix’ the results of each cycle. EHTransport employs a linear mixing scheme on the atomic charges, initially assuming all atoms are uncharged, thus the change in the excess charge from one step to the next is given by

$$\vec{\Delta Q}_{n+1} = \lambda \vec{\Delta Q}_{new} + (1 - \lambda) \vec{\Delta Q}_n . \quad (3.30)$$

This has the advantage over mixing the Hamiltonian that there is only one vector of charges

regardless of the system or number of k-points, and it is a single vector rather than a matrix. More advanced schemes exist than linear mixing, and can in general reach a given level of convergence with fewer steps. A parameter free version of Pulay mixing guaranteed to reach convergence if a set of assumptions hold [76] was implemented, but disabled after testing showed worse convergence behaviour than straightforward linear mixing. A more standard Pulay mixing scheme has yet to be implemented.

Like the overlap matrix, the calculation of a single Hamiltonian matrix is naturally an $\mathcal{O}(N^2)$ operation, reducing to $\mathcal{O}(N)$ if a cutoff has been applied to the overlap matrix and the matrices are stored in a sparse format. For a non-self consistent EHT calculation the constant prefactor of this operation is significantly smaller than that of the overlap matrix. A self-consistent calculation makes it difficult to gauge complexity, as it is not possible to predict beforehand the number of iterations that will be required to reach any desired convergence criteria. We can at least look at the complexity of calculating the energy perturbations for each orbital: for standard SC-EHT, the limiting step is the calculation of the occupied eigenvalues and eigenvectors of the system, which is an $\mathcal{O}(N^3)$ operation. There is also a matrix multiplication involved, which although mathematically $\mathcal{O}(N^3)$ when computed by the naive algorithm⁵ can be reduced to $\mathcal{O}(N^2)$ by noting that only the diagonal elements of \mathbf{P} are required. (Since \mathbf{Z} is not Hermitian, or even sparse, it cannot be reduced to $\mathcal{O}(N)$.) Using NEGF requires further matrix maths, again with a complexity of $\mathcal{O}(N^3)$, but which also needs to be repeated at several energy points to obtain a usable sample, greatly increasing the computational cost compared to standard EHT. For the Poisson based scheme, assuming that there is a linear relationship between the number of atoms in a system and the volume of the unit cell, construction of the real space density scales as $\mathcal{O}(N^2)$ where N is the number of atoms in the system rather than orbitals. In general, calculation of the overlap matrix is the dominating step for a non self-consistent calculation, and calculation of the Hamiltonian matrix elements dominates in the self-consistent case.

Other features

For the lead systems, an estimate of the Fermi level is required. For this the generalised eigenvalue problem is solved for the eigenvalues. A specialised method is available to perform this calculation

⁵In fact, for large matrices MKL employs a more complex algorithm with a complexity between $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$.

efficiently, taking advantage of the properties of the matrices. (i.e. that they are both either real and symmetric or Hermitian, depending on periodicity, and that the overlap matrix is positive definite.) The eigenstates are filled in ascending order up to the number of electrons in the (neutral) system, and the Fermi energy is taken to be the average of the highest occupied level and the lowest unoccupied one.

This procedure is exactly equivalent to the electron counting method of the Mulliken population analysis at zero temperature. For this reason, the same code is used for both, and the previously described finite temperature method can be applied here to give a Fermi level estimate that is not exactly at the midpoint between the HOMO and LUMO. As previously noted though, this temperature procedure is somewhat artificial, and the predictions of where the Fermi level lies would not be accurate.

Other optional outputs of the EHT module include the DOS of the system, and the PDOS projected onto a user-defined subset of the orbitals used in the calculation. The initial procedure for this is identical to the Fermi level computation; the generalised eigenvalue problem is solved, but this time both eigenvalues and eigenvectors are required. The total DOS is taken to be a summation of a set of Gaussian functions centred on each eigenvalue

$$DOS(E) = \frac{1}{b\sqrt{2\pi}} \sum_i e^{(E-E_i)^2/(2b^2)} , \quad (3.31)$$

where b is a user defined broadening factor and the summation is over each eigenvalue in the basis set used in the EHT calculation. To compute the density of states projected onto a given orbital j , an additional weighting term c_{ij} is added.

$$PDOS_j(E) = \frac{1}{b\sqrt{2\pi}} \sum_i c_{ik} e^{(E-E_i)^2/(2b^2)} . \quad (3.32)$$

The coefficients c_{ij} are computed in manner similar to the Mulliken population analysis in equation 3.28, with the exception that the summation over eigenvalues is not performed, leading to a set of occupations per basis set function and per eigenstate. The PDOS is then summed over the orbitals of interest, which may be an atom, all atoms of a given species, all atoms making up a specific part of the system under study, or so on.

EHTransport is also capable of outputting band structures. Plotting the band structure of

a bulk material is a useful method of checking the EHT parameters, especially given that many sets of EHT parameters are calculated by band structure fitting in the first place. Obviously, a band structure plot can only be made for a periodic system. The procedure involves nothing new that has not already been done. Each band is a straight line in k-space, usually between two high symmetry points. A set of points is chosen equispaced along the line, and overlap and Hamiltonian matrices are calculated at each of these points. The generalised eigenvalue equation is solved for each Hamiltonian, but only for the eigenvalues and not the eigenvectors, which improves computational efficiency. The eigenvalues are plotted out for each point. It must be noted that this procedure does not account for the crossing of bands: EHTransport only outputs the eigenvalues in ascending order, so a straightforward plot will not show any crosses. Because there may be a high number of k-points involved in making a good quality plot of the band structure, and because the selection of k-points required for making such a plot may not be a good representation of the overall system, the full self consistent procedure is not followed using the k-points in a band structure calculation. Rather, a self consistent calculation is carried out using a separate selection of k-points, and the resulting charges from this initial calculation are fixed for the band structure calculation.

Another available output is the partial charges on each atom. If the original EHT calculation was self consistent, the charges from the final iteration of the self consistent loop will already be available, and can be supplied for no additional computational cost. If the original calculation was not self consistent, EHTransport will perform a Mulliken population analysis to obtain the charges.

EHT parameters

Since EHT is a semi-empirical theory, it depends on a set of parameters that must be chosen prior to performing a calculation. Commonly used values for the Slater coefficients and uncharged energy parameters exist and include the Hoffman⁶, Müller⁷ and Cerda [82] parameter sets.

Parameters describing the energy changes due to charge transfer do not have pre-existing well-known and tabulated values, and are calculated from a run of all-electron DFT calculations

⁶Although not tabulated in a single publication, these have been used in e.g. [77,78]

⁷Also not available from a single publication, but tabulated as part of the distribution of YAEHMOP 2.0, and based on work in [79–81]

Species	Orbital	I (eV)	α_a	c_a	α_b	c_b	K
H	s	-13.600	1.300	1	NA	0	1.75
C	s	-21.400	1.625	1	NA	0	1.75
C	p	-11.400	1.625	1	NA	0	1.75
N	s	-26.000	1.950	1	NA	0	1.75
N	p	-13.400	1.950	1	NA	0	1.75
O	s	-32.300	2.275	1	NA	0	1.75
O	p	-14.900	2.275	1	NA	0	1.75
F	s	-40.000	2.425	1	NA	0	1.75
F	p	-18.100	2.425	1	NA	0	1.75
P	s	-18.600	2.122	1	NA	0	1.75
P	p	-14.000	1.827	1	NA	0	1.75
S	s	-20.000	2.122	1	NA	0	1.75
S	p	-11.000	1.827	1	NA	0	1.75
Fe	s	-9.100	1.900	1	NA	0	1.75
Fe	p	-5.320	1.000	1	NA	0	1.75
Fe	d	-12.600	5.350	0.5505	2.000	0.6260	1.75
Zn	s	-12.410	2.010	1	NA	0	1.75
Zn	p	-6.530	1.700	1	NA	0	1.75
Au (Chain)	s	-10.929	2.602	1	NA	0	1.75
Au (Chain)	p	-5.550	2.293	1	NA	0	1.75
Au (Chain)	d	-12.605	2.292	0.596	NA	0	1.75
Au (Bulk)	s	-12.134	2.316	0.603	NA	0	2.3
Au (Bulk)	p	-6.740	1.745	0.627	NA	0	2.3
Au (Bulk)	d	-14.026	2.327	0.376	5.445	0.794	2.3

Table 3.1: The parameter set used for all EHT calculations in this thesis, excluding the self consistent energy corrections.

using the SIESTA subcomponent ATOM. The procedure is to take a single isolated atom and tabulate how the Kohn-Sham energies of each orbital vary with the atomic charge. The range of charge is somewhat limited; a complete electron can be removed from the system, but only a fraction can be added, as the procedure used by ATOM fails if the eigenvalue of any occupied state becomes positive, but there is more than sufficient range to compute a best fit. From the data either a best fit second order polynomial can be calculated, or else a linear fit can be made at zero charge for systems where the charge on each atom remains close to zero. The second order polynomial provides a very close fit to the results, so further orders are not needed. There are separate values for each l quantum number; i.e. one for each of the s, p and d orbitals.

One consideration as already mentioned previously is that the original non-self consistent values have some level of charging effects 'built in'; it is recognised that some atoms are typically




Figure 3.1: A plots of orbital energy against excess charge for the valence orbitals of a single isolated atom. From left to right and top to bottom, the atomic species are gold, carbon, hydrogen and iron. Over the range shown, the curves are very close to quadratic.

charged in molecules, and to some extent the standard uncharged ionization energy parameter for non-iterative EHT reflects this. Thus the 'best' uncharged energy parameters may differ between the self-consistent and non-self-consistent cases. New values could be computed to match the NEGF density of states to the Transiesta results.

3.2.3 The NEGF module

Input to the NEGF calculation is a set of 3 overlap and Hamiltonian matrices, to represent the scattering region, and the left and right leads, along with the lead Fermi levels. If no right lead data is given, it is assumed to be identical to the left lead. The lead matrices must represent two principal layers, and must have a number of elements no greater than that of the scattering region. (The result of this constraint is that no orbital in the scattering region is coupled to both leads, which allows some simplification of the required equations.) The right hand lead overlap

Species	Orbital	$\alpha_i(\text{eV per electron})$	$\beta_i(\text{eV per electron}^2)$
H	s	-11.249	-2.454
C	s	-10.321	-1.896
C	p	-9.874	-2.024
N	s	-12.096	-2.026
N	p	-11.665	-2.140
O	s	-13.853	-2.072
O	p	-13.424	-2.186
F	s	-15.582	-2.115
F	p	-15.147	-2.229
P	s	-8.433	-0.979
P	p	-7.853	-0.963
S	s	-9.487	-0.994
S	p	-8.915	-0.963
Fe	s	-7.590	-1.221
Fe	p	-5.199	-3.229
Fe	d	-12.113	-2.197
Zn	s	-7.100	-1.142
Zn	p	-5.489	-1.412
Au	s	-6.945	-0.506
Au	p	-4.943	-0.990
Au	d	-7.807	-0.633

Table 3.2: The computed parameters α_i and β_i from equation 3.27 for all considered atomic species, using a quadratic best fit.

and Hamiltonian matrices may be omitted. If so, it is assumed that both leads are identical, and the left hand matrices are replicated. This allows for some further optimisation.

In an NEGF calculation at an energy E , the first step that is required is the calculation of the self energy matrices (and hence broadening matrices) of each lead. There are several available methods of extracting these self energy matrices. EHTransport uses an iterative block diagonalisation procedure [34], at an energy taken to be $E \pm V/2 + \delta i$ where V is the bias across the scattering region to allow for the inclusion of finite bias, and δi is the same small complex part as in the scattering region calculations. First a set of four matrices, $\mathbf{v}^{\mathbf{ab}}$, are calculated.

$$\begin{aligned}
\mathbf{v}^{00} &= \mathbf{v}^{11} = (E\mathbf{S}^{nn} - \mathbf{H}^{nn})^\dagger \\
\mathbf{v}^{10} &= E\mathbf{S}_{ij}^{nm} - \mathbf{H}_{ij}^{nm} \\
\mathbf{v}^{01} &= (E\mathbf{S}_{ij}^{nm} - \mathbf{H}_{ij}^{nm})^\dagger
\end{aligned} \tag{3.33}$$

These are then used to compute a new set of matrices \mathbf{v}'^{ab} , defined as

$$\begin{aligned}
\mathbf{v}'^{00} &= \mathbf{v}^{00} - \mathbf{v}^{01}(\mathbf{v}^{11})^{-1}\mathbf{v}^{10} \\
\mathbf{v}'^{11} &= \mathbf{v}^{11} - (\mathbf{v}^{01}(\mathbf{v}^{11})^{-1}\mathbf{v}^{10} + \mathbf{v}^{10}(\mathbf{v}^{11})^{-1}\mathbf{v}^{01}) \\
\mathbf{v}'^{10} &= (-\mathbf{v}^{10})\mathbf{v}^{11})^{-1}\mathbf{v}^{10} \\
\mathbf{v}'^{01} &= (-\mathbf{v}^{01})\mathbf{v}^{11})^{-1}\mathbf{v}^{01} .
\end{aligned} \tag{3.34}$$

Equations 3.34 are then iterated, using the new set of \mathbf{v}'^{ab} matrices as the \mathbf{v}^{ab} input matrices for the next cycle. With each iteration, the matrices \mathbf{v}'^{10} and \mathbf{v}'^{01} will reduce in magnitude. EHTransport continues the iteration until the largest element of \mathbf{v}'^{01} is less than 10^{-6} . The matrices \mathbf{v} can then be used to compute the final self energy $\mathbf{\Sigma}$ and broadening matrices $\mathbf{\Gamma}$ of the lead

$$\begin{aligned}
\mathbf{\Sigma} &= \mathbf{v}^{01}(\mathbf{v}'^{00})^{-1}\mathbf{v}^{10} \\
\mathbf{\Gamma} &= i(\mathbf{\Sigma} - \mathbf{\Sigma}^\dagger) .
\end{aligned} \tag{3.35}$$

Note that in equation 3.35 it is the original \mathbf{v}_{01} and \mathbf{v}_{10} that are used, and not the final set. At this point, the matrices \mathbf{v}'^{10} and \mathbf{v}'^{01} are, to within some small tolerance, zero.

The rest of the procedure used to compute the required Green's function matrix is as described in chapter two. First the Green's function of the scattering region is calculated, and the self energies $\mathbf{\Sigma}$ of the leads subtracted off. Finally the scattering region Green's function is multiplied by the lead broadening matrices $\mathbf{\Gamma}$. The current is calculated by computing a transmission spectrum across a range of energies from a few $k_B T$ below the lower of the two adjusted lead Fermi energies up to a few $k_B T$ above the greater of them, and then integrating over the transmission multiplied by the difference in the Fermi functions of each lead. The DOS can also be calculated by a method similar to that of the transmission.

When NEGF is being used for charge-consistency in the EHT calculation, the same input of 3 sets of matrices are required. The charge-consistency applies only to the scattering region; the leads are fixed. As the lead self energies and broadening matrices are used in the same way as in




Figure 3.2: The contour of integration followed when computing the density matrix of the scattering region via the Green's function method.

the transmission calculation, and the lead Hamiltonians are fixed, they are calculated only once in the first SCF loop and then saved. In this case there is no obvious lower limit to the integration so a point somewhere below the lowest eigenvalue of the scattering Hamiltonian is taken. Since unbroadened states that don't contribute to transport are important, the integration path follows a semicircle into the imaginary axis, and only returns to the real axis at the lower of the two lead Fermi energies, above which the Cauchy integral theorem is not valid. This procedure will pick up on unbroadened states that exist below the Fermi levels of both leads, but will miss out states that are unbroadened and that exist between the two Fermi energies of the leads. A more complicated procedure is available to account for them [64], that performs the integration between the two lead Fermi levels on a line offset from the real axis, above some number of the poles that lie on the imaginary axis, but this has yet to be implemented.

$$\mathbf{D} = \frac{1}{\pi} \int_{-\infty}^{\min(E_{fL}, E_{fR})} \left(\mathbf{G}\mathbf{\Gamma}_L\mathbf{G}^\dagger + \mathbf{G}\mathbf{\Gamma}_R\mathbf{G}^\dagger \right) dE + \frac{1}{\pi} \int_{\min(E_{fL}, E_{fR})}^{\max(E_{fL}, E_{fR})} \left(\mathbf{G}\mathbf{\Gamma}_L\mathbf{G}^\dagger F(E - E_{fL}) + \mathbf{G}\mathbf{\Gamma}_R\mathbf{G}^\dagger F(E - E_{fR}) \right) dE \quad (3.36)$$

$$p_i = \mathbf{D}_{ii} .$$

The NEGF module has no intrinsic support for periodicity; it doesn't accept a list of k-points as input, but only individual overlap and Hamiltonian matrices . Rather, it is the responsibility of the input module to call the NEGF procedure multiple times, once for each k-point, and to make an appropriate summation of the results accounting for the k-point weightings. It is important that the x and y components of the k points match up between the leads and the

scattering region. (Since the system is not periodic in the z direction, (the direction in which conduction occurs in EHTransport,) the z component of each k -point must always be zero) Two forms of the NEGF module were written; one that assumes real and symmetric input matrices, and hence that $E\mathbf{S} - H$ is complex and symmetric, and one that assumes complex Hermitian input matrices. In this latter case, $E\mathbf{S} - H$ has no special properties.

For efficiency, the self energy matrices of the leads may be cached to the hard drive. This is useful if the same leads are to be used with several different systems, since the computation of the self energy matrices can take of the order of half of the total NEGF calculation time in cases where the lead contains many orbitals. (e.g. the case of a molecule trapped between bulk gold contacts.) This can take a very large amount of disc space if a lot of energy points are desired; usually many gigabytes.

Bond currents

The NEGF formalism also provides a means of computing 'bond currents' [41, 42]. These are somewhat abstract quantities rather than a precise measure of an actual current flowing through individual molecular bonds, but still provide a method of identification of the sections of a molecule which are most important to current flow. They also give some idea of the type of conduction which is taking place. (i.e. whether through s or p orbitals.)

The bond currents for the left hand lead \mathbf{B}_l are given by

$$\mathbf{B}_{lij} = 2\text{Im}(\mathbf{H}_{ij}\mathbf{A}_{lji}) , \quad (3.37)$$

where \mathbf{A}_l is given by

$$\begin{aligned} \mathbf{A}_l &= \sqrt{\mathbf{S}}^{-1} \mathbf{G}' \mathbf{\Lambda}'_l \mathbf{G}'^\dagger \sqrt{\mathbf{S}}^{-1} \\ \mathbf{\Lambda}_l &= \sqrt{\mathbf{S}}^{-1} \Gamma_l \sqrt{\mathbf{S}}^{-1} \\ \mathbf{G}' &= \sqrt{\mathbf{S}} \mathbf{G} \sqrt{\mathbf{S}} . \end{aligned} \quad (3.38)$$

The bond current matrix is real and antisymmetric. (i.e. $\mathbf{B}_l = -\mathbf{B}_l^T$) An equivalent bond current matrix can be defined for the right contact by replacing the left hand lead broadening

matrix Γ_l with the equivalent right hand lead matrix Γ_r . For a symmetrical system $\mathbf{B}_l = \mathbf{B}_r^T$, although this is not true in general.

Chapter 4

Choice of Parameters and Testing of EHTransport

4.1 Introduction

With any new code, it's vitally important to thoroughly test all aspects of it to ensure correctness. Where a code contains configurable parameters, these also need to be carefully chosen in order to give meaningful results. In this chapter choices are made for the configurable settings within SIESTA, allowing it to be used for comparisons with EHTransport. The EHT and NEGF modules of EHTransport are then tested separately in turn. Finally some simple systems are run in EHTransport to test the functioning of the code as a whole.

4.2 Choice of SIESTA Parameters

As the package name ‘Spanish Initiative for Electronic Simulations with Thousands of Atoms’ hints, SIESTA is designed to handle very large systems¹, via a linear scaling DFT algorithm [83]. Unfortunately this algorithm does not result in the construction of a suitable Hamiltonian for use with the Green’s function equations, and so is not used for any calculations in this thesis. To give correct results, it also relies on being able to pick out a energy at the start of a calculation which sits within the bandgap at the start of a calculation and remains there throughout the whole self consistent procedure. This makes it completely unsuitable for materials with no bandgap, such as metallic leads, and very difficult to use for materials without a substantial gap.

SIESTA uses a local basis set, constructed numerically by placing the pseudopotential inside a spherical infinite well and computing the resulting wavefunctions. The size of the well is chosen such that the difference between the energy of the ground state of the system and the first excited state is equal to a certain pre-chosen value. The whole calculation takes place with this basis set, with molecular orbitals becoming linear combinations of them, so although DFT is not formally a LCAO theory in and of itself, the SIESTA implementation is restricted to LCAO solutions. The effect of this is offset by the optional use of multiple-zeta and polarisation orbitals, to allow for additional degrees of freedom in the basis set.

Despite being an ab-initio method, there are still several choices that need to be made when performing a DFT calculation. These range from the choice of exchange-correlation functional

¹‘Large’ here is used comparatively with other DFT packages. Whereas SIESTA may be able to cope with a few thousand atoms in a realistic time frame, compared to a few hundred for an N^3 scaling program, a classical empirical program may be able to process several orders of magnitude more.

and basis set, to a number of adjustable parameters which are a result of the individual implementation of DFT rather than the underlying theory. Before use in real systems, it's important to find appropriate values that give acceptable accuracy for the systems under study without sacrificing computational efficiency. For packages using a localised basis set, previous results [84] suggest that a double-zeta polarised (DZP) basis set provides accurate results for the valence orbitals which will be important to conductance. For the choice of exchange-correlation functional, although SIESTA implements a range of functionals, only pseudopotentials built for use with either the LDA or PBE functionals are supplied with the package. Most of the work here uses the PBE functional, which offers moderate improvements in computed ground state energies compared to the LDA [85] for little extra computational cost.

For SIESTA, the remaining choices that effect the quality of results consist of the number of grid points on a real space mesh upon which some calculations are performed, and the number of k-points used in a periodic calculation. The real space mesh is chosen based on an energy cutoff, where higher energies correspond to a finer mesh². The k-points are chosen by specifying a Monkhorst-Pack 3×3 matrix and a 3 element displacement vector [66]. When choosing simulation parameters, the convergence of the total system energy is often used as a test. Some tests on small systems, consisting of bulk gold and sodium, were made to examine the effects of these choices (Figures 4.1 and 4.2).




Figure 4.1: Convergence of system energy while varying the real space mesh energy cutoff (left) and k-point mesh density (right) of a bulk sodium system in SIESTA. For k-point calculations, the Monkhorst-Pack matrix had diagonal elements n and off-diagonal elements 0. The x axis values correspond to n .

²The basis set chosen in plane-wave based DFT is also often chosen based on an energy cutoff, and this is often the first thought when an energy cutoff is discussed in the context of DFT. The SIESTA cutoff has a different meaning however, and the cutoff energy is typically higher.




Figure 4.2: Convergence of energy while varying the real space mesh energy cutoff (left) and k-point mesh density (right) of a bulk gold system in SIESTA. For k-point calculations, the Monkhorst-Pack matrix had diagonal elements n and off-diagonal elements 0. The x axis values correspond to n .

Bulk metal primitive unit cells will require the most k-points of any system to accurately model, and any selection which can solve these problems reliably should be sufficient for any of the other systems under study, which will primarily involve gold contacts and carbon based molecules, and significantly larger unit cells. An $8 \times 8 \times 8$ Monkhorst-Pack mesh was sufficient for obtaining a converged system energy in the sodium and gold bulk systems. More may be required to obtain e.g. an accurate DOS, but it is not necessary to perform the entire self-consistent calculation with the additional k-points. The real space mesh has a much smaller effect on results, and the smaller than default value of 100 Rydburg appears sufficient for the system energy to converge. It does have effects beyond the energy output however; e.g. smaller cutoffs make it markedly more difficult to get a converged density matrix. For most calculations the mesh cutoff is left at Transiesta’s default of 450 Rydburg.

Additional common configurable parameters include mixing settings such as the number of Pulay matrices to use in mixing and the mixing weight. These do not aim to change the final result³, but do affect the efficiency with which the result is reached. If the mixing parameter is too large for a system, a converged density matrix may never be obtained regardless of the number of iterations. In this thesis these settings have been left at their default values where possible, and the mixing weight has been decreased where required to ensure convergence.

A final option is the toleration for error in the density matrix. This is specified as a maximum

³Although they do not change the global minima of the system, it is possible to get the self consistent loop trapped in a local minima, resulting in differing final results.

alteration of any element of the density matrix between one self consistency loop and the next. If after a loop no element of the density matrix has changed by more than this amount, SIESTA assumes that convergence has been reached, and the self consistency loop ends. This option has no effect on the value of the converged density matrix of a system, but does affect how close SIESTA will try to get to that answer before giving up. In this thesis the slightly smaller than default value of 5×10^{-5} of SIESTA's internal units has been used throughout.

With the SIESTA parameter set chosen, it can now be used to aid in the testing of EHTransport.

4.3 Testing of EHTransport's EHT module

4.3.1 Overlaps of Slater orbitals

The first section of an EHT calculation is to calculate the overlap matrix \mathbf{S} , which is done using the analytic method already outlined. It is important not only that these overlaps should be correct, but that the real space representations of the orbitals are known, or the results of the calculation could not be used to make real space plots of such quantities as the wavefunctions or electron density. The test for orbitals is therefore that for any given set of quantum numbers and positions, the result must be identical to a straightforward volume integral, in both magnitude and sign. As a further check, the result of $\{i|j\}$ should be equal to $\{j|i\}$ ⁴. The equations of the analytic form are not obviously symmetric at first glance.

Volume integrals and analytic overlaps were calculated for all valence l , m and n for a pair of gold atoms with separations of $\vec{R} = (0, 0, 1)$, $(1, 1, 0)$ and $(1, 2, 3)$ Bohr. The orbitals were double-zeta, with the coefficients and exponents taken from the Cerda parameter set [82]. The volume integral was carried out by explicitly computing the value of the orbitals at all points on a discreet grid, and multiplying and summing the results. The grid was cubic, centred on the point midway between the two atoms, of side length 16 Bohr and a constant grid point separation of 0.08 Bohr, which is large enough to capture the bulk of the Slater orbitals and fine enough to provide an accurate result.

The analytic results are correctly symmetric. In terms of magnitude, there was agreement

⁴Actually $\{j|i\}^*$, but as we're working with the real form of the spherical harmonics, the complex conjugation has no effect.

between the analytic and integral expressions to within 4 significant figures. (Table 4.3.1) The difference is due to the naive implementation of the volume integral; finer and larger grids converge towards the analytic result. Most results which are known to be 0 from symmetry are exact in the analytic form. There was a mismatch of sign between some overlaps between the analytic and integral results. It was found that the analytic result needed to be multiplied by an additional term of $-1^{l_i+l_j+m_i+m_j}$ to provide a correct match to the signs of the volume integral results. Finally, for the case of non-zero l_1 and l_2 some overlaps that should have been exactly zero gave small but nonzero ($\approx 10^{-17}$) results. These anomalous results were also antisymmetric ($\{i|j\} = -\{j|i\}$). This was traced to an error in cancellation of the sums of numbers of opposite signs, and to ensure the overlap matrix remains symmetric and positive definite all results below 10^{-15} are clamped to zero.

n_i	l_i	m_i	n_i	l_i	m_i	$\{i j\}$	$\{j i\}$	Integral
5	0	0	5	0	0	3.137308×10^{-1}	3.137308×10^{-1}	3.137308×10^{-1}
5	0	0	5	1	0	1.050678×10^{-1}	1.050678×10^{-1}	-1.050674×10^{-1}
5	0	0	5	1	1	$-6.433546 \times 10^{-18}$	6.433546×10^{-18}	$-2.621861 \times 10^{-18}$
5	1	0	5	1	0	3.264015×10^{-1}	3.264015×10^{-1}	3.263668×10^{-1}
5	1	0	4	2	0	2.181145×10^{-2}	2.181145×10^{-2}	-2.181152×10^{-2}
5	1	1	5	1	0	1.967868×10^{-18}	$-1.967868 \times 10^{-18}$	$-7.869261 \times 10^{-19}$
4	2	2	4	2	1	2.912583×10^{-17}	$-2.912583 \times 10^{-17}$	$-8.677949 \times 10^{-18}$

Table 4.1: A selection of overlaps of double-zeta Slater orbitals separated by 1 Bohr in the Y and Z directions. The sign correction has not been included.

4.3.2 EHT tests of small molecules

The simplest test of the EHT code is a single hydrogen atom. Since the hydrogen atom is represented by a single s-orbital, the overlap and Hamiltonian matrices have dimensionality 1×1 , and the solution is trivial. $\mathbf{S} = 1$ and $\mathbf{H} = I_{Hydrogen}$, so there is a single solution of the eigen-problem of $E = I_{Hydrogen}$, or -13.6 eV, and $\mathbf{Z} = 1$. This system was run through the EHT code with periodicity and self-consistency disabled, to ensure that the solution was indeed -13.6 eV. SIESTA in this case predicts a total energy of -13.4 eV; since hydrogen has no core electrons the approximation of the core with a pseudopotential doesn't have such a great effect, but it does still prevent a perfect result.

As the next simplest possible test, an additional hydrogen atom can be added at varying

distances from the first, for a total of two orbitals in the system. The disassociation energy is unfortunately not accessible under EHT, and the only information contained in the two eigenvalues of this system is the HOMO-LUMO gap, which will be a great overestimate due to the lack of higher energy unoccupied orbitals in the calculation. The only check possible here is simply that the two eigenvalues both tend towards -13.6 eV at large separations, and that they split apart as the molecule is formed. This does indeed happen, with EHT predicting a HOMO-LUMO gap at the equilibrium separation of 21.8 eV compared to 12.7 eV in SIESTA. The disassociation energy is an accessible quantity in SIESTA, as the zero point of energy is fixed if the cell size and the number and types of atoms in the system are also fixed. Using a bond length of 1.40 Bohr, it predicts a disassociation energy of 4.53 eV, which compares very well with the experimental value of 4.52 eV.




Figure 4.3: The HOMO and LUMO energies of a hydrogen molecule as the interatomic spacing is varied.

For the third test, a benzene molecule was geometry optimized in SIESTA, and then the eigenvalues of the optimized molecule computed with SIESTA and EHTransport. There is fairly good match up between the two for the lower energy levels, also extending to the first three unoccupied orbitals before the effects from the limited basis functions start dominating. That results don't instantly decay to the point they are no longer useful as soon as the Fermi level is reached is important for conductance calculations, as the transmission is needed both above and below the Fermi level. (That there are vaguely accurate states above the Fermi level is a trait of the carbon atoms; With only 4 electrons in their outer band, they have one completely unoccupied p orbital. Using molecules with greater p-orbital occupancy will shift the Fermi level

up towards the inaccurate levels.) That the results do rapidly lose accuracy does limit the voltage that can be usefully simulated though. SIESTA predicts a HOMO-LUMO gap of 5.16 eV, while EHT predicts a somewhat smaller gap of 4.42 eV. Both of these are within similar range of the experimental value, which is approximately 4.72 eV [86].




Figure 4.4: The energy levels of a benzene molecule computed with EHTransport and SIESTA. The two diverge significantly above the Fermi level.

4.3.3 Mulliken analysis tests - Tetrafluoromethane

Another important part of the EHT calculations is the Mulliken population analysis, and hence the self consistency loop. To check this, the Tetrafluoromethane (CF_4) molecule was examined. Fluorine has an extremely high electronegativity, and the bonds it forms with carbon have a sizeable amount of ionic character. A single molecule of CF_4 (using a bond length of 1.344 Angstrom) was tested with both SIESTA and EHTransport. In this case, the high valence occupation of the fluorine pushes the Fermi level up, and there are no good states above the Fermi level given by EHT; it predicts a huge band gap of over 20eV. Experimentally, CF_4 has a partial charge on the carbon atom of +0.76 electrons [87]. The self consistent EHT calculation predicts +0.69, which is in good agreement with the experimental value. To highlight the importance of self consistency, (at the least, when examining the Mulliken charges in a system,) performing a Mulliken analysis without enabling self-consistency gives a partial charge on the carbon atom of +2.55. Interestingly, the Mulliken population analysis routine built into SIESTA suggests a partial charge on the carbon of -0.476; that is, that the carbon atom has acquired charge from

the fluorine. SIESTA’s Mulliken population analysis routine appears to not be functional, at least in the version used, and for this calculation type, and is not used further in this thesis.

4.3.4 Tests of EHTransport’s Monkhorst-Pack implementation

To test the Monkhorst-Pack code for generation of k-points, a set of k-points generated for a given grid by EHTransport can be compared to those of SIESTA. Since matching up of k-points is important for comparisons between EHTransport and SIESTA, for the same unit cell dimensions and Monkhorst-Pack matrix and displacement vector, the same set of k-vectors should be generated by both applications. The comparisons should cover the cases of non-Cartesian unit cells, but non-diagonal Monkhorst-Pack matrices are not needed, and the only displacement vectors considered are (0, 0, 0) or (0.5, 0.5, 0.5). In all tested cases EHTransport matched results with SIESTA exactly (Tables 4.2 and 4.3. SIESTA was, however, able to identify additional exploitable symmetry that EHTransport could not discover, and so in some cases could locate and merge additional equivalent k-points to end up with fewer in total than EHTransport. All EHT and NEGF calculations in EHTransport scale linearly with k-point count, so this additional merging stage could result in a speed advantage of the order of 20% for the some systems.

EHTransport				SIESTA			
x	y	z	weighting	x	y	z	weighting
0.000000	0.000000	0.000000	0.125000	0.000000	0.000000	0.000000	0.125000
3.141593	0.000000	0.000000	0.125000	3.141593	0.000000	0.000000	0.125000
0.000000	3.141593	0.000000	0.125000	0.000000	3.141593	0.000000	0.125000
3.141593	3.141593	0.000000	0.125000	3.141593	3.141593	0.000000	0.125000
0.000000	0.000000	3.141593	0.125000	0.000000	0.000000	3.141593	0.125000
3.141593	0.000000	3.141593	0.125000	3.141593	0.000000	3.141593	0.125000
0.000000	3.141593	3.141593	0.125000	0.000000	3.141593	3.141593	0.125000
3.141593	3.141593	3.141593	0.125000	3.141593	3.141593	3.141593	0.125000

Table 4.2: A sample set of k-points generated for a simple cubic unit cell of side length one Bohr, a Monkhorst-Pack matrix with diagonal elements (2, 2, 2), zero off-diagonal elements and a zero displacement vector.

4.3.5 Poisson solver

The Poisson solver code is intended for use in the case of a bias across the system, to account for effects such as screening. The interesting output here is the potential throughout the cell.

EHTransport				SIESTA			
x	y	z	weighting	x	y	z	weighting
0.000000	0.000000	-2.094395	0.037037	0.000000	0.000000	-2.094395	0.074074
-1.047198	1.047198	-1.047198	0.037037	-1.047198	1.047198	-1.047198	0.074074
-2.094395	2.094395	0.000000	0.037037	-2.094395	2.094395	0.000000	0.074074
1.047198	-1.047198	-1.047198	0.037037	1.047198	-1.047198	-1.047198	0.074074
0.000000	0.000000	0.000000	0.037037	0.000000	0.000000	0.000000	0.037037
-1.047198	1.047198	1.047198	0.037037				
2.094395	-2.094395	0.000000	0.037037				
1.047198	-1.047198	1.047198	0.037037				
0.000000	0.000000	2.094395	0.037037				
1.047198	1.047198	-3.141593	0.074074	1.047198	1.047198	-3.141593	0.074074
0.000000	2.094395	-2.094395	0.074074	0.000000	2.094395	-2.094395	0.074074
-1.047198	3.141593	-1.047198	0.074074	-1.047198	3.141593	-1.047198	0.074074
2.094395	0.000000	-2.094395	0.074074	2.094395	0.000000	-2.094395	0.074074
1.047198	1.047198	-1.047198	0.074074	1.047198	1.047198	-1.047198	0.074074
0.000000	2.094395	0.000000	0.074074	0.000000	2.094395	0.000000	0.074074
3.141593	-1.047198	-1.047198	0.074074	3.141593	-1.047198	-1.047198	0.074074
2.094395	0.000000	0.000000	0.074074	2.094395	0.000000	0.000000	0.074074
1.047198	1.047198	1.047198	0.074074	1.047198	1.047198	1.047198	0.074074

Table 4.3: A sample set of k-points generated for a FCC unit cell of side length one Bohr, a Monkhorst-Pack matrix with diagonal elements (3, 3, 3), zero off-diagonal elements and a displacement vector of (0.5, 0.5, 0.5). SIESTA is able to uncover additional symmetries to identify equivalent k-points, and hence merge four pairs of k-points together.

With just a single atom present, (Figure 4.5a) there are no additional orbitals available in which the charge can redistribute, resulting in no charge sources in the unit cell and so the potential remains a linear drop across the system cell. Adding a second atom at a distance sufficient to prevent interaction between the two (Figure 4.5b) causes charge to migrate in order to counteract the electric field, distorting the otherwise linear potential drop in the vicinity of the atoms. Adding periodic boundary conditions in directions perpendicular to the applied field (Figure 4.6a) turns this into a capacitive system; the change in potential is concentrated into the gap between the two atoms. With a periodic chain of atoms, (Figure 4.6b) in the real world would result in a flat permittivity of a system, there is a linear drop in potential across the box aside from the edges of the cell where the boundary conditions interfere. This is all in line with what would intuitively be expected for the real behaviour of these systems, showing that the Poisson code can successfully replicate these expected effects.




Figure 4.5: The potential gradient of an isolated cell containing a single hydrogen atom (left), and a pair of hydrogen atoms arranged along the z axis (right) across which a 0.1V bias is applied, taken as a cross-section through the y axis. More rapidly varying colours designate more rapidly changing potential. Bands of the same colour which are not connected do not correspond to the same potential.




Figure 4.6: The potential gradient of a periodic cell containing a pair of hydrogen atoms (top) and a chain of hydrogen atoms (bottom), across which a 0.1V bias is applied, taken as a cross-section through the atoms. More rapidly varying colours designate more rapidly changing potential. Bands of the same colour which are not connected do not correspond to the same potential.

4.3.6 Band structure and orbital PDOS

As a test of the EHT parameters, and of EHTransport's ability to compute band structures and projected densities of states, a bulk gold system was considered. The unit cell chosen was primitive FCC with a side length of 5.452 Bohr to match the experimental structure of gold, and contained a single atom. An initial run with a $16 \times 16 \times 16$ k-point mesh was used to locate the Fermi level and calculate the density of states projected onto the s, p and d-orbitals. No self consistency was used, as within the framework used by EHTransport no charge redistribution is possible with only a single atom in the system.

As expected, the band structure matches up extremely well to SIESTA results. This is only to be expected as the Cerda parameter set used for this calculation was created by fitting to a bulk gold band structure computed from a DFT calculation.

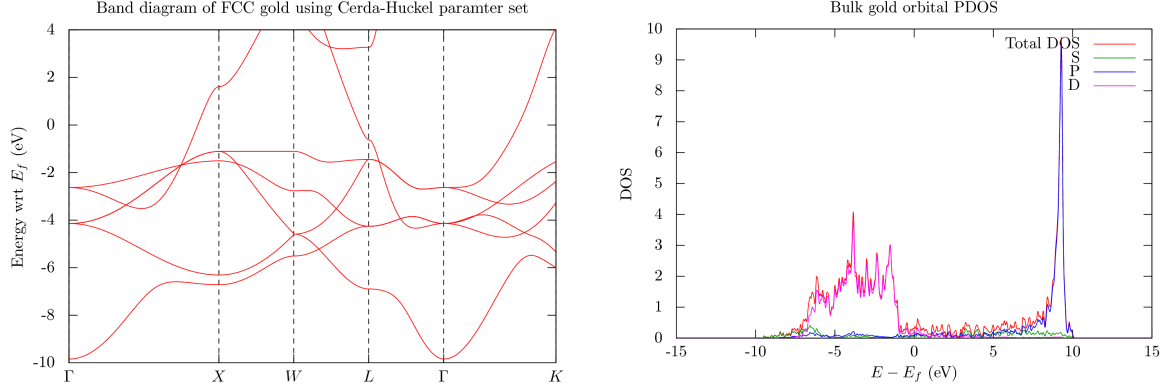


Figure 4.7: The bulk gold band structure (left) plotted along lines connecting the high symmetry points of the FCC Brillouin zone and the density of states projected onto orbitals of each l value (right) computed using EHTransport and the Cerda parameter set.

4.4 Tests of the Green's functions code

4.4.1 Model Hamiltonian input

A standalone test of the NEGF module requires input of a model Hamiltonian and overlap matrix, derived from some system for which an exact analytic result is known. The system of a one-dimensional potential barrier was chosen, as it has a simple analytic result, and is amenable to conversion into an overlap/Hamiltonian matrix form. The leads are chosen to be at a flat potential of 0 eV, and to have a principal layer length of 2 Bohr. The scattering region is 10 Bohr in length, and contains a centred barrier of width 2 Bohr and height 10 eV (Figure 4.8). Since the basis set must be localised, a set of Gaussian functions are used, equi-spaced along the lead and scattering region at a distance of 0.1 Bohr, and with $\sigma = 0.1$.

The $\mathbf{S}_{ij} = \{i|j\}$ and $\mathbf{H}_{ij} = \{i|\rho(x)|j\}$ matrices were obtained by multiplication and summation on a discrete grid of size 2 Bohr and spacing 0.001 Bohr, centred at the point between the two Gaussian functions. These are then used to solve for the transmission and DOS using the NEGF formalism in the normal way. The analytic result for the transmission may be obtained from equation 4.1

$$T(E) = \left| \frac{e^{-ik_1 w}}{\cosh(k_2 w) + (k_2^2 - k_1^2) \frac{i \sinh(k_2 w)}{2k_1 k_2}} \right|^2, \quad (4.1)$$

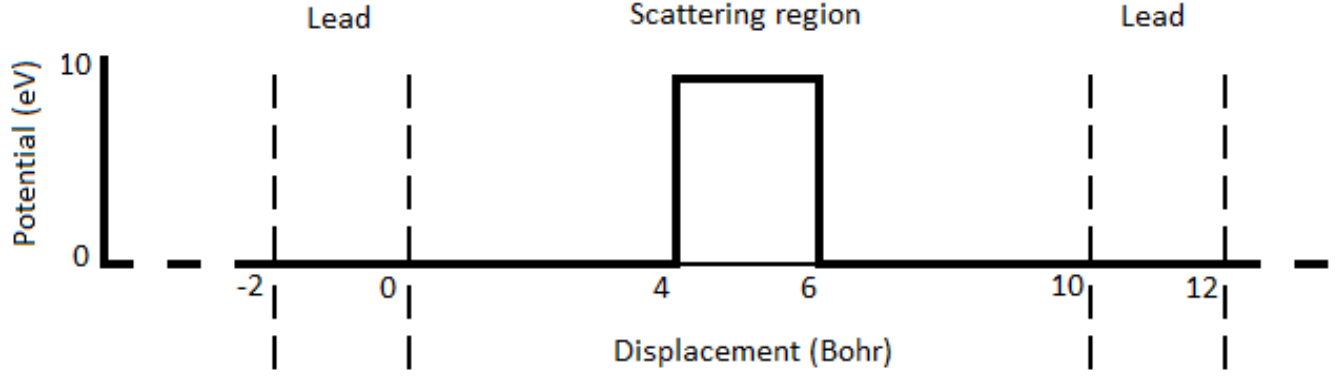


Figure 4.8: The structure of the one dimensional potential barrier.

where w is the width of the barrier in atomic units and

$$\begin{aligned} k_1 &= \sqrt{2E} \\ k_2 &= \frac{\sqrt{2(E-v)}}{i}, \end{aligned} \quad (4.2)$$

where v is the height of the barrier, again in atomic units.

Two versions of the NEGF code are implemented in EHTransport; one that requires real symmetric input, and the other that accepts complex Hermitian matrices. This system was run through them both code paths in order to compare their output to the true analytic result, at 2000 points between 5 and 20 eV.

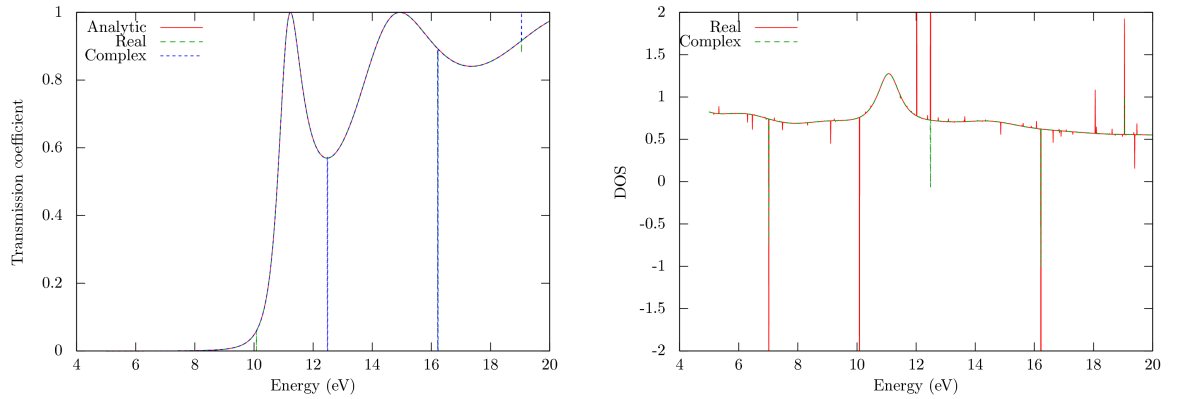


Figure 4.9: The transmission (left) and DOS (right) of one dimensional potential barrier computed using EHTransport. The transmission is also compared to the analytic result.

The transmission spectra computed using both the real and complex versions of EHTrans-

port’s NEGF code agree with the analytic result almost exactly. However, the transmission plots, and particularly the DOS plots, show isolated anomalous points away from the otherwise smooth curve. The problem is worse for the real input code path than for the complex code path. This is attributed to symmetry breaking during processing of the matrices; when a small positive imaginary part is applied to a real and symmetric matrix the result is complex, but still symmetric. When it is added to a Hermitian matrix the result no longer has any special properties. The code path that requires real matrices is thus optimized to expect that matrices at certain points in the calculation are symmetric. However, a combination of numerical errors and the non-perfect block diagonalisation of the leads can result in matrices that are no longer perfectly symmetric, resulting in errors when algorithms are applied to the matrices that assume that they are.

Because of the reduced reliability of the real code path, all future calculations employed the complex code path. This can come at a penalty to speed; making the assumption of symmetry makes several LAPACK functions more efficient. The real path can only be applied to Γ point calculations, so the greatest penalty occurs for large systems which need few k-points.

4.4.2 SIESTA Hamiltonian input

As a second test of the Green’s function code, which also covers non-real input matrices, (i.e. a calculation away from the Γ point,) a Transiesta calculation was performed on a bulk hydrogen system (which is itself one of Transiesta unit tests.). The Hamiltonian generated from Transiesta was fed into both the usual TBTrans post-processing utility and EHTransport, and the DOS and transmission spectra calculated with both and compared. The match up is nearly perfect throughout, both in details and scale. It was also noted that after accounting for the additional k-point folding, the TBTrans utility performed the task in approximately half the time as EHTransport, showing there is still room for improvement in the efficiency of EHTransport’s NEGF module.

A calculation of bond currents shows similar agreement (Table 4.4).




Figure 4.10: The transmission (left) and density of states (right) of a bulk hydrogen system at the Γ point.

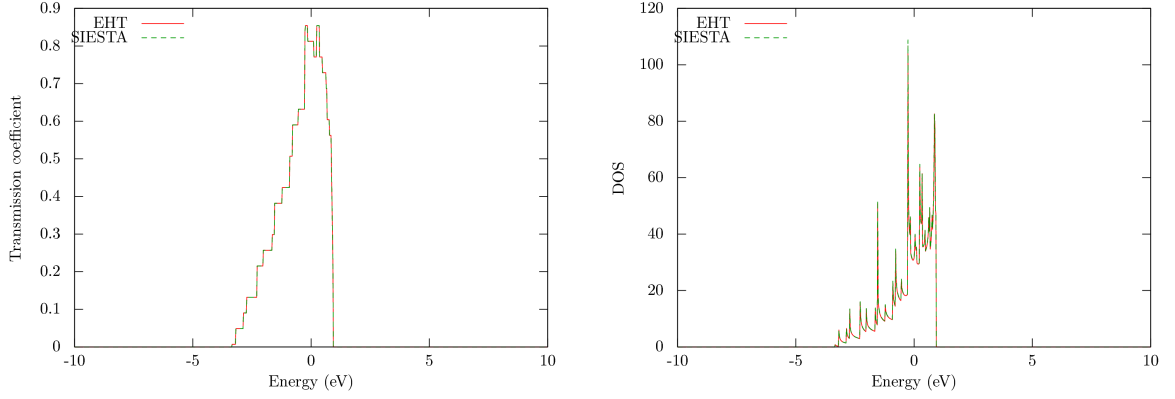


Figure 4.11: The transmission (left) and density of states (right) of a bulk hydrogen system using a Monkhorst-Pack matrix with diagonal elements $12 \times 12 \times 1$.

4.5 Test calculations

4.5.1 Gold capacitor

The gold capacitor system consists of two parallel (111) surfaces. In the [111] direction the gold atoms form planes with a repeating ABC structure. The lead principal layer consists of 3 planes, i.e. one single ABC group, the smallest possible repeating unit. The scattering region consists of 6 gold layers, followed by a vacuum gap, followed by 6 further layers. The layers on each side of the gap are matched, giving the complete structure ABCABC-ABCABC. This allows both sides to link correctly to the leads. A non-Cartesian 1×1 unit cell was used in the directions perpendicular to conduction, giving a total of 12 atoms in the calculation. An 8×8 k-point mesh

EHTransport			SIESTA		
-2.50×10^{-23}	0.999999	0	-1.99×10^{-22}	1.00003	0
-0.999999	3.19×10^{-21}	0.999997	-1.00003	5.87×10^{-21}	0.999996
0	-0.999997	-7.16×10^{-22}	0	-0.999996	2.44×10^{-20}

Table 4.4: The 'bond current' matrix elements corresponding to the first 3 atoms in the bulk hydrogen system, sampled at the gamma point.

was used. Taking the usual interatomic separation as the zero point, a vacuum gap of between 0 and 10 angstroms was tested.

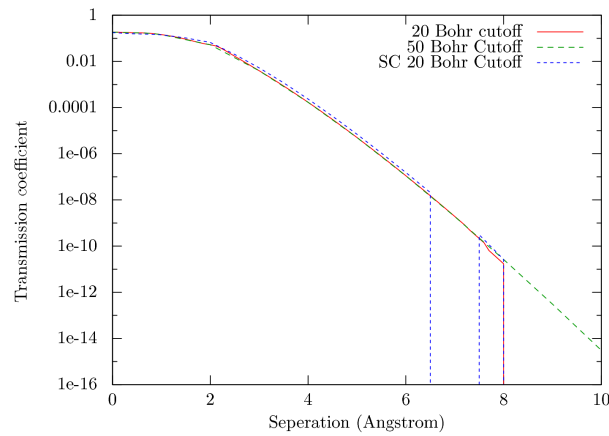


Figure 4.12: A gold (111) capacitor shows exponential scaling of transmission at the Fermi level with plate separation.

For the case where the plate separation is significantly greater than the usual interatomic separation, the results match the known analytical results that tunnelling drops off exponentially with plate separation. For the case where plate separation is less than the normal atomic separation, EHT will not provide a good description of the system, and the results would be invalid⁵. The difference between the 20 Bohr cutoff and 50 Bohr is not significant until the plate separation becomes greater than the cutoff distance, showing that the smaller cutoff is sufficient. All future runs use the 20 Bohr cutoff unless otherwise noted. Self consistency here causes problems; within the region between 6 and 8 angstroms separation, there is an anomalous drop in Fermi level conductance.

⁵The pseudopotential-using SIESTA would also not be capable of describing this system; the problem comes about since the approximation that core electrons remain unaffected would be invalid. Fortunately this system is non-physical, and will not need to be handled.

4.5.2 Monatomic gold wire

The monatomic gold wire is a similar setup to the capacitor, except that a gold wire is placed in the gap instead of leaving it empty. A non-Cartesian 3×3 unit cell was used in directions perpendicular to conduction for the gold leads (Figure 4.13). The 3×3 size was sufficient to minimize interaction between periodic separations of the gold wire. The FCC structure of the gold is continued through the wire, such that the atoms in the wire are in the same positions as if a 1×1 cell had been repeated through throughout the gap region.




Figure 4.13: The structure of the monatomic gold wire system, represented orthographically. The leads are 3 atoms deep.

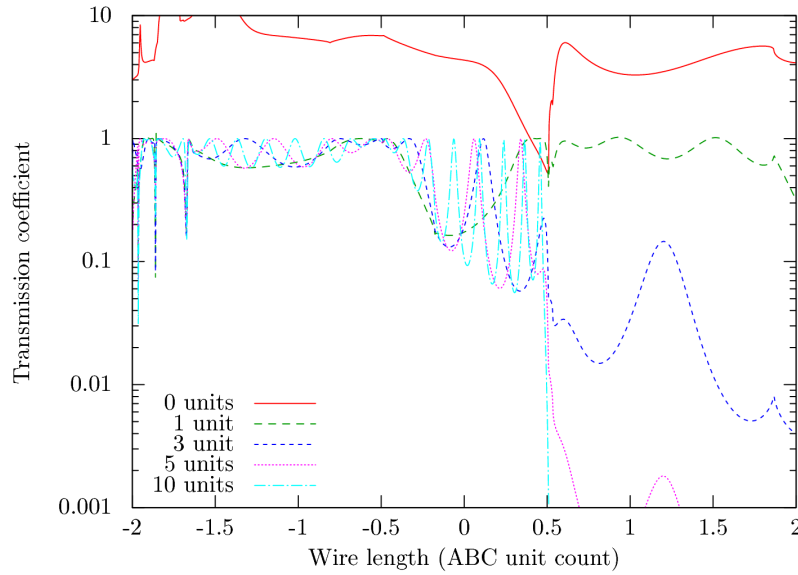


Figure 4.14: Electronic transmission across a monatomic gold wire between two gold (111) electrodes, for different numbers of repetitions of the ABC wire unit.

With a chain length of zero cells (i.e. the two bulk contacts are in direct contact, creating a bulk gold system,) transmission is greater than G_0 per eV across a wide energy range (Figure

4.14). As soon as monatomic chain units are added, the transmission is capped at G_0 per eV, consistent with the theoretical maximum conductance through a single channel. Below the Fermi level, the transmission is largely independent of the chain length. Above the Fermi level, the transmission has a straightforward dependence on chain length, dropping off as the chain length increases. Around the Fermi level the structure is more interesting; although the average transmission varies little, increasing the chain length adds more peaks in this region, aligned such that low bias conductance is predicted to be greater for a wire containing 10 ABC layers than for a wire containing a single layer.

Chapter 5

Conduction in Porphyrin Molecules

5.1 Introduction

In previous chapters I have introduced EHTransport and the underlying mathematics, and tested the code on simple systems. In this chapter, I report on the results of calculations on porphyrin dimer based systems trapped between gold contacts using both EHTransport and the DFT based SIESTA computer codes. The results from EHTransport are compared to those from SIESTA as a demonstration of the effectiveness of EHTransport in probing the density of states and conductance of these systems.

A porphyrin ring is a planar structure containing 20 carbon atoms and 4 nitrogen atoms. The center of the ring may contain a metal atom, or alternatively the dangling bonds can be capped by hydrogen. (Figure 5.1) Although not common as isolated rings, complexes based on porphyrin are ubiquitous in biology. The heme molecule, a vital functional component of the hemoglobin protein that allows our blood to carry oxygen, is a porphyrin ring with an iron atom at the center and several additional side chains [88]. An important class of proteins known as cytochromes also contain a heme group, and have electron transport roles in biology. Although there are currently few applications using these systems outside of biology, there has already been some interest in probing the possible uses of porphyrins in molecular electronics as wires and switches [20, 89].




Figure 5.1: The structure of a free-base porphyrin ring, without a metal ion at the center.

The porphyrin system studied consists of a pair of Zinc-Porphyrin rings connected together via a carbon backbone, and linked to the gold contacts via thiol groups (Figure 5.2).




Figure 5.2: The structure of the system under study.

5.2 Modelling the contacts

In experiment, gold is a common choice for the contact material. In an STM experiment, a molecule may be attached to a (111) bulk gold surface, and a gold STM tip brought in to make contact with the other side. In a NEGF theoretical calculation semi-infinite leads are required, and it is simple to create a bulk gold structure such that the (111) direction points along the direction of conduction. The interface between molecule and the contact is not simple, however. In experiment, there will very likely be surface reconstructions on the surface of any material. Furthermore, as an STM tip is withdrawn the strong gold-sulphur bond may draw out a short monatomic chain of gold atoms [90,91].

In this chapter, two alternative contacts are investigated. Firstly isolated linear gold chains were studied, followed by an unreconstructed bulk gold (111) surface. The chains have the advantage of containing fewer atoms than a bulk gold surface, allowing faster computation with SIESTA for comparisons. Gold chain structures in physical systems are only stable on very small length scales however [91], so extending them into semi infinite leads may limit their abilities to provide insights into experiment. For this chain structure, the gold atoms were assumed to lie in a straight chain. The atomic spacing was first calculated with SIESTA in a similar way to the bulk system; four gold atoms were placed into a Cartesian unit cell. The four atoms had equal x and y coordinates, and were equispaced along the z axis. The size of the cell in the x and y directions was kept large enough to prevent interactions between periodic images of the chain in these directions, and several variable cell geometry optimizations were carried out with progressively tighter tolerances on the density matrix error and the final stresses on the unit cell. The result was an equilibrium atomic spacing of 2.51Å.

Because this chain structure is so different from bulk gold, neither the Hoffmann nor Cerda

parameter sets were found to be appropriate to accurately describe the contact, and so a new set of EHT parameters was needed. SIESTA was used to calculate the DOS of the gold chain, and then a simplex optimization procedure [92] was employed to match this to the DOS of an EHT calculation. The calculation was seeded with the Hoffmann parameter set. After some experimentation as to which parameters affected the important features in the DOS, the double-zeta d orbital energy, exponents and mixing coefficient were allowed to vary, as was the single zeta p orbital exponent. The s orbital and p orbital energy were held constant. The resulting DOS is not perfect (Figure 5.3), but does reproduce important features of this system such as the Fermi level being at the edge of the d-band, rather than somewhat above it inside of the s-band as it is in the bulk material. (And with the chain, using the Hoffmann parameter set.)




Figure 5.3: The DOS of a monatomic gold chain computed with the different available parameter sets in EHTransport, and compared to SIESTA.

Modelling the contacts as a bulk surface is more realistic deep within the electrodes, although even in this case there may be effects missing from surface reconstruction, especially around the molecule binding site, where the strong thiol-gold bond is capable of pulling atoms from the gold surface completely [90]. The (111) surface is often the one used in experiment, although the (100) surface is sometimes favoured in theoretical calculation due to the smaller repeating unit and typically smoother transmission spectra from the Green’s function method. In the (111) direction the gold atoms form planes with a repeating ABC structure. The lead principal

layer consists of 3 planes, i.e. one single ABC group, the smallest possible repeating unit. The scattering region contains 5 gold layers, in an ABCAB - molecule - BCABC pattern. The repeat of the B layer allows for simple positioning of the molecule with symmetric contacts. A bulk gold calculation was carried out in SIESTA to compute the optimal unit cell size. This calculation utilised a single Cartesian FCC unit cell, containing two gold atoms. This cell was found to have length 4.134\AA .




Figure 5.4: Two ABC layers of bulk gold electrode. The direction of conduction is along the z axis, along the $[111]$ direction, from the left to the right of the image.

5.3 The porphyrin-gold interface

There are several ways of linking porphyrin molecules together to form chains. In this section I consider two similar connectors: a conjugated carbon backbone consisting of either two or four atoms, each connected to the middle of one edge of a porphyrin ring (Figure 5.5). The porphyrin complex is linked to the gold contacts by means of a thiol linker (Figure 5.2). There are several places this thiol linker can be added to the porphyrin chain, and previous work has investigated the effects [24]. In this case, one hydrogen atom from the edge of the ring is replaced with a sulphur atom, at the opposite side from the backbone. In this structure the backbone and both thiol groups are arranged along a single line, and the porphyrin rings can be rotated relative to each other by twisting the backbone.

The geometry of an isolated single zinc-porphyrin ring was initially optimised with the computer code Avogadro, which implements a classical forcefield based geometry optimization tool. UFF classical forcefields [93] were employed. (The more specialised forcefields method implemented by Avogadro for optimization of organic molecules, MMFF94, [94] was not used due to it




Figure 5.5: The structure of the two porphyrin dimers.

not having the capability of handling the zinc atom.) This empirically optimised structure was then used as input to a SIESTA geometry optimisation run, using the PBE exchange/correlation functional and a 300 Rydberg cutoff for the real-space mesh. The sulphur atoms were capped by hydrogen atoms and placed perpendicular to the plane of the molecule. This caused the sulphur to displace slightly out of the plane of the rest of the molecule, but the displacement was far smaller than if the hydrogen atoms were kept in plane. In this second case the hydrogen atoms were positioned close to other hydrogen atoms, which repelled them and pushed the sulphur towards the carbon atom on the opposite side. With the single ring configuration optimized, one sulphur atom was removed and replaced with a second ring, connected by either the 2C or 4C backbone, and this resulting dimer optimised again in SIESTA. For constructing porphyrin chains of more than two units, it was assumed that the spacing between the carbon atoms along each backbone remained the same.

With the molecular geometry optimized, the next requirement was to find the most energetically favourable form of binding to a gold contact, under the assumption that the positions of all atoms in the contact were fixed. The hydrogen cap was removed from one sulphur atom, and this atom placed over a non Cartesian $3 \times 3 \times 5$ gold (111) surface¹, periodic in the x and y directions, in one of 4 different positions; atop, bridge, or the FCC or HCP hollow sites. (Figure 5.6) The height was allowed to vary in order to locate the minimum of energy in each of the 4 possible positions, while the size of the unit cell was held constant. This had to be performed via multiple runs with static geometry and differing initial positions, as a standard geometry optimisation failed due to SIESTA having difficulty moving the whole porphyrin molecule rigidly

¹Where $3 \times 3 \times 5$ refers to the number of atoms in the directions of the 3 vectors of the unit cell, giving a total of 45 atoms. There are 5 atoms (the ABCAB layers) along the third vector, which is orthogonal to the other two, points in the z direction, and is the direction along which conduction occurs.

at once: Everything tended to stay fairly close to the initial positions. The FCC site (Figure 5.7) was found to be the most energetically favourable, in line with previous results [95].




Figure 5.6: The four possible positions on the gold (111) surface at which possible binding was considered. One set of ABC layers are shown, with the C layer (green) being on the surface, followed by the B layer (red) and A layer (yellow).

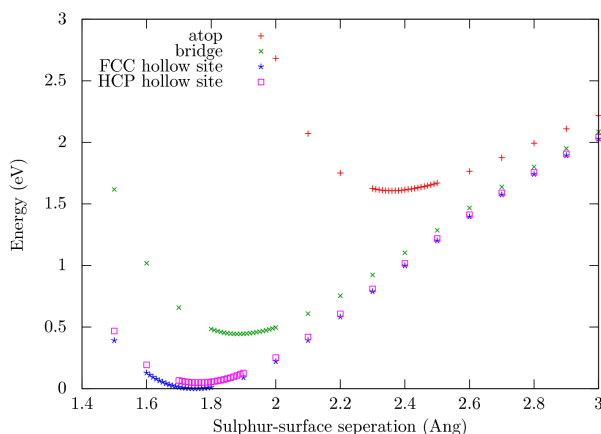


Figure 5.7: Binding energy of the thiol group to various positions on the gold surface. The FCC hollow site was found to be the preferred binding point.

As the final part of geometry optimisation, the hydrogen cap was stripped from the second sulphur, and additional multiple static geometry runs were made with SIESTA with contacts on both sides to determine the optimal gold-molecule distance for the hollow site with a second contact added. This was performed using a $4 \times 4 \times 5$ gold (111) surface, slightly larger in the x and y directions. Both the 2C and 4C configurations were checked, and were found to produce the same minimum of 2.10 \AA (Figure 5.8), which is considerably larger than the separation from a single contact. That the change should be so great from the addition of an extra contact was

unexpected. To rule out error due to the changing unit cell size, (already made unlikely by the fact that 2C and 4C configurations gave the same minimum,) a single contact run was made with an increased z dimension. The energy change was of the order of one part in a million, or 0.001 eV, which suggests the difference is legitimately due to the extra contact and gold atoms. This procedure was also repeated for the gold chain, for which the sulphur was assumed to be on the same line as the gold atoms, yielding an equilibrium separation of 2.27Å.




Figure 5.8: Variation of the energy with the distance between the sulphur atom of the porphyrin dimer and the gold surface. The lowest points of the curves were been chosen to be at zero energy, and are at 2.27Å for the gold chain and 2.10Å for the 111 hollow site.

5.4 Effects of rotation on system energy

In investigating the effect of rotations on system energy, one porphyrin ring of a dimer was rotated around the axis joining the two sulphur atoms. The carbon atoms connecting the porphyrin rings lie on this axis, and do not move. The dihedral angle is measured between the planes of the two porphyrin rings. The systems were run in SIESTA using the PBE exchange-correlation functional and a 450 Rydberg cutoff for the real space mesh². Because of the symmetry of the system, the full range of results is found by probing between 0° and 90°. The gold contacts were held rigid throughout the calculation, and do not rotate along with the molecule.

²total system energy is not an accessible quantity with EHTransport.

5.4.1 Chain contacts

For calculations using a gold chain as the electrode, the gold chain lies on the same axis as the sulphur atoms and carbon backbone. One principal layer of the chain lead contained three gold atoms. A total of 12 gold atoms were present in the scattering region.

A dimer linked by a 2C backbone requires 0.20 eV of energy to rotate from the energetically most favourable position to the least (Figure 5.9). A dimer linked by a 4C backbone requires 0.17 eV, which is not significantly lower given the far greater separation between the rings. The most energetically favourable angle does change from 25° to 0° though, presumably due to reduced interactions between the hydrogen atoms adjacent to the carbon atoms connecting the two rings. This interaction is probably masking a part of the required rotational energy for the 2C backbones, which would otherwise have been >0.2 eV. The 4C result is much higher than the 0.035 eV reported in [21]. Systematic energy fluctuations from the eggbox effect³ and from finite real-space size of the cell are from interactions between periodic images of the molecule (found by rotating the whole molecule at once, rather than the dihedral angle,) are of the order of 0.01 eV, and random energy fluctuations are insignificant.




Figure 5.9: The energy difference against dihedral angle relative to the energy minimum for a zinc-porphyrin dimer with a 2C backbone (left) and a 4C backbone (right) attached to gold chain leads.

³The affect on the calculation caused by the relative location of an atom to the points on the real space mesh on which parts of the computation are performed.

5.4.2 Bulk contacts

The calculations using bulk contacts used a $4 \times 4 \times 5$ gold (111) surface for each contact in the scattering region, which is not quite large enough to avoid interactions between periodic images of the porphyrin dimer, but a $5 \times 5 \times 5$ surface was computationally prohibitive, and the periodic interactions are small enough to be corrected for in the results. Despite the periodicity, due to the size of the large size of the system calculations were performed at the Γ point. One principal layer of the lead consisted of three layers of gold atoms, or a single period of the periodic structure along the (111) lattice direction. There were a total of 160 gold atoms within the scattering region.

Initially, results for a bulk contact are mostly masked by other effects, such as the hydrogen atoms on the porphyrin edges sweeping past gold atoms in the contact. There are also greater random fluctuations, due to the far higher magnitude of the absolute system energy. (The fluctuations of the order of 0.05 eV is still only a change of 1 part in 3×10^{-7} in the total energy.) Subtracting off half of the energy obtained from rotating the whole molecule, (and hence the effects of two contacts, and periodic interactions of the molecule with itself,) gives a result that suggests that the lowest energy occurs at 35° for a 2C backbone and 0° for 4C, with energies of rotation of 0.08 eV and 0.07 eV respectively. These are far lower than the chain results, although the ratio between the two is still similar. Also of note is that the 4C result is in now in better agreement with results in [21]. The bulge at 45° is due to the interactions between the periodic images of the molecule, as this is the point at which they are closest together, with the hydrogen atoms at the edges of the rotating rings sweeping past each other. The peaks at 0° and 60° are due to hydrogen atoms passing close to gold atoms in the contact.

Performing a geometry optimisation at each rotation step with the porphyrin dimer attached between (111) gold surfaces lowers energies by a further 0.5 eV, which is a significant amount compared to the energy of rotation. Despite this, the effects on transmission or relative energies during dihedral angle rotation were small (Figure 5.11). Analysis of forces on atoms during the geometry optimization confirm that it is the exposed surfaces of the bulk gold contacts that undergo the greatest movement, and that the porphyrin dimer is little affected. Since this surface movement has only a small effect on results, and that geometry optimization is computationally expensive, the structure was held rigid and further geometry optimization was not performed




Figure 5.10: The energy difference against dihedral angle relative to the energy minimum for a zinc-porphyrin dimer with a 2C backbone (left) and a 4C backbone (right) attached to bulk gold (111) surfaces.

while varying other parameters.




Figure 5.11: The energy difference against dihedral angle relative to the energy minimum for a zinc-porphyrin dimer with a 4C backbone attached to gold 111 surfaces, performing a geometry optimization at each step.

5.5 Effects of rotation on molecular conductance

5.5.1 Chain contacts

All transmission plots were made with the 4C backbone structure. Using chain contacts, SIESTA predicts a Fermi level around 0.2 eV higher than EHT, assuming a neutral system. The general shape of both plots are in good agreement. Transmission at the Fermi level peaks at 55° in

SIESTA, due to the edge of the high transmission region shifting to a higher energy. Away from this region, transmission peaks at 0° . The minimum of transmission always occurs at 90° . (Figure 5.13) The transmission approaches G_0 in places, but does not exceed it. The position of the Fermi level right at the edge of the d-band causes a strong dependence on the exact position of the Fermi level for the transmission.

There is a high part of transmission below the Fermi level, and a low part above. Since conduction through this sort of planar organic molecule normally takes place through Π -bonding orbitals (i.e. the portion of the p-orbitals pointing perpendicular to the plane of the molecule,) this can be explained by symmetry arguments. The Fermi level is positioned around the edge of the gold d band, so above the Fermi level conduction must take place through the gold s-orbitals. (The gold p-band is too high in energy to play a role near the Fermi level.) From symmetry, the overlap between the s-orbital on any gold atom and the p-orbital pointing out of the plane of the molecule must be zero, preventing coupling between the molecule to the lead at those energies. The d-orbitals however have no such restriction, and so the molecule is more strongly coupled to the gold d-band.

A comparison of figure 5.12b to figure 5.14a shows how vital the self consistent procedure is to obtaining a good match up in results to Transiesta. Without self consistency, the transmission spectrum at the edge of the d-band is altered by approximately two orders of magnitude.

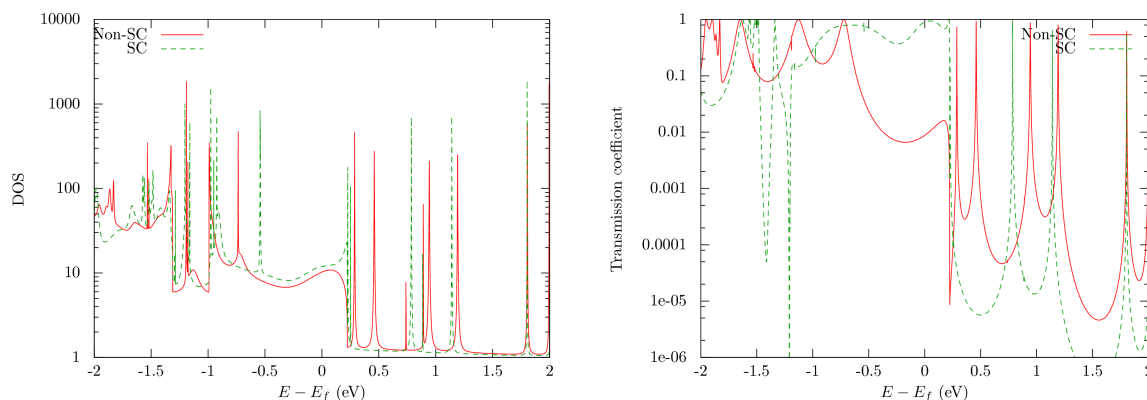


Figure 5.12: The DOS and transmission spectra of a porphyrin dimer between gold chains comparing EHTransport with self consistency enabled and disabled.

A comparison of bond currents between EHTransport and Transiesta (Figure 5.15) shows good agreement. Both codes predict very little current flow through the central zinc atom, and




Figure 5.13: Transmission through a Zn-porphyrin dimer between gold chains as a function of dihedral angle, computed using Transiesta (Left) and EHTransport. (Right)




Figure 5.14: Transmission through a Zn-porphyrin dimer between gold chains at several angles, computed using Transiesta (Left) and EHTransport. (Right)

that the bulk of current flow takes place symmetrically around the outside of the ring.

5.5.2 Bulk contacts

With bulk gold contacts, the Fermi level is above the edge of the gold d-band rather than below it. This makes the transmission curve far more stable with respect to the exact position of the Fermi level; the maximum is now always at 0° (Figure 5.17). The actual value of the maximum still varies, but is of the order of 0.1 compared to 1 in the chain contact case. As with the chain case, the use of self consistency in the computation greatly improves the match up of results with Transiesta.

The difference in Fermi level between Transiesta and EHTransport is increased to 0.4 eV in




Figure 5.15: Bond currents through the porphyrin molecule, computed using Transiesta and Inelastica (Top) and EHTransport (Bottom). Inelastica does not provide currents through contact atoms or directionality information, and labels atoms by index as opposed to the atomic number labels used by EHTransport.

the case of bulk gold contacts. This is large enough that EHTransport is now predicting LUMO conduction, compared to the HOMO based conduction of Transiesta. The transmission spectra produce by EHTransport are also compressed in comparison to Transiesta, particularly above the Fermi level (Figure 5.18).

Since the gold atoms are no longer positioned in a chain pointing through the thiol linkers, the symmetry arguments that prevent the gold s-band from coupling to the porphyrin pi orbitals no longer apply. The result of this is the much higher conduction visible throughout the gold s-band.




Figure 5.16: The DOS (left) and transmission spectra (right) of a Zn-porphyrin dimer between gold (111) surfaces comparing EHTransport with self consistency enabled and disabled.




Figure 5.17: Transmission through a Zn-porphyrin dimer between gold (111) surfaces as a function of dihedral angle, computed using Transiesta (Left) and EHTransport (Right). The Fermi energy has been offset in the EHTransport case in order to line up with the position predicted by Transiesta.

5.6 Computation time

For chain contacts, computation time for SIESTA was 1 hour for the DFT stage and 2 hours 20 minutes for the NEGF stage using 2000 data points. The SIESTA calculation took place on a supercomputer node using 8 CPU cores. The DFT stage memory consumption peaked at 6.4GB. Computation time for EHT was approximately 5 minutes for the SC-EHT stage and 20 minutes for the NEGF stage using 10001 data points. The EHTransport calculation took place on a desktop machine using 2 CPU cores.

For bulk contacts, computation time for SIESTA was approximately 9 hours on 8 CPU cores for the DFT part, and a further 24 hours on 1 CPU core to compute the transmission at 161 energy points. A maximum memory consumption of 16.8GB was recorded during the DFT stage. For EHT the self-consistent EHT stage took 1 hour on 2 CPU cores, and the NEGF stage 10 hours on 2 CPU cores at 2001 data points, and completed successfully on a machine with 4GB of memory.

The bulk electrode calculation employed a total of 236 atoms in the scattering region, compared to 88 for chain electrodes. Furthermore the additional atoms were gold, which has a large number of valence electrons and has occupied d-orbitals, so the effect on the system orbital count is greater than that of the atom count. SIESTA, using a DZP basis set, increased more than 3-fold from 1012 to 3232 orbitals when moving from chain to bulk gold contacts.

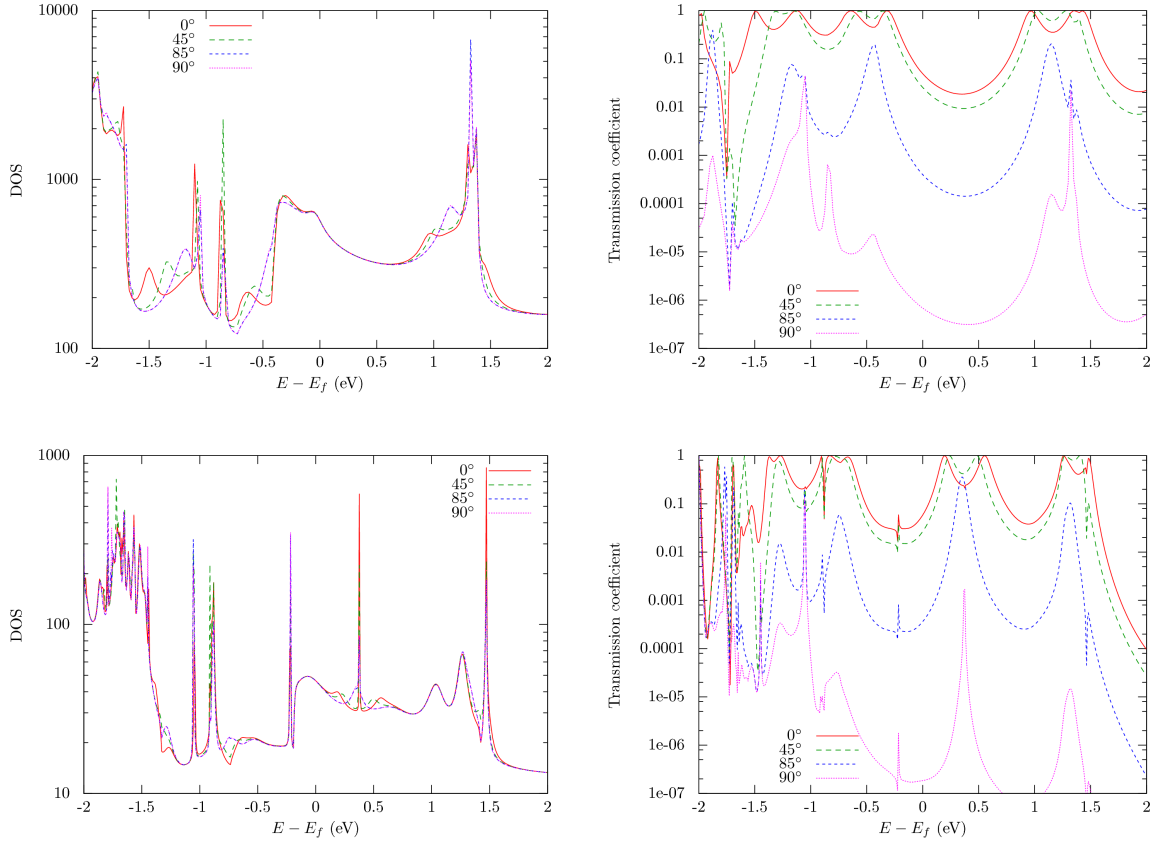


Figure 5.18: DOS and Transmission through a Zn-porphyrin dimer between gold (111) surfaces at several angles, computed using Transiesta (Top) and EHTransport (Bottom).

5.7 Conclusion

EHTransport has been able to provide results which provide a good match to Transiesta. Self consistency during the EHT calculations has proved absolutely vital to obtaining this good match. The degree of match up is sufficient to provide confidence in utilising EHTransport on larger systems for which Transiesta results cannot be obtained.

The very large and reliable change in conductance raises the possibility that the rotation could be exploited to build a switching device. The studied porphyrin dimer has some flaws which prevent this: The energy barrier is too low, and so it could not be operated at room temperature, let alone at the higher temperatures typically experienced by electronic components undergoing resistive heating. There is also no mechanism of reliably flipping between states. It is possible that such a mechanism could be created via the addition of side chains to one of the porphyrin rings. If it were possible to hold it in the 90° position via application of a bias to an additional

third electrode, it could be used as a transistor with an on/off ratio of the order of 10^5 , which is a reasonable ratio compared to current organic transistors [96] but poor compared to conventional doped-silicon based devices. If it were possible to engineer a local energy minimum at the 90° position, the molecule could be used as memory, although the removal of the peak of the energy barrier would increase any issues with stable operation at room temperature. As an alternative, one side of a porphyrin ring could be engineered to remove the symmetry when one porphyrin ring is completely flipped, such that switching from a 0 to a 1 bit would require a 180° rotation of the bond rather than 90° .

Chapter 6

EHT calculations of large Porphyrin Systems

6.1 Introduction

In the previous chapter the results of EHTransport calculations on porphyrin dimers were compared to Transiesta results. The Transiesta calculations took many days to complete, leaving exploration of properties of larger porphyrin systems out of its reach. For example, questions such as behaviour as the length of the porphyrin chain increases, or the effects of solvation cannot be answered using Transiesta with the currently available computing resources. Encouraged by the previous results that EHTransport is able to make useful predictions about the effects of structural changes to a system, I have explored some further properties of porphyrin wires using EHTransport in an attempt to make behavioural predictions of several more complex systems.

If a porphyrin chain is to be used as a molecular wire, systems longer than a dimer will be required. Chains containing up to 5 porphyrin units were analysed. Electrically active porphyrins incorporate magnesium (chlorophyll) or iron (heme) rather than zinc, so an iron based porphyrin was considered and the effect of replacing the zinc analysed. Another important factor that must be taken into account when comparing theoretical calculations to experiment is the effect of any surrounding molecules. Experimental measurements may take place in vacuum, but could also be taken in air or another gas, or a liquid solvent. The effects of the presence of water were analysed by adding a set of water molecules in close proximity to the porphyrin molecule. The final section of this chapter explores the properties of a heme-based cytochrome containing 1750 atoms.

6.2 Effects of wire length on molecular conductance

To investigate the suitability of porphyrin chains for use as molecular wires, the zinc-porphyrin dimer investigated in chapter 5 was replaced by porphyrin chains of varying lengths, from a monomer up to five units. Each unit was connected by a backbone consisting of 4 carbon atoms, and the wire was linked to the (111) surface of bulk gold contacts via the usual thiol linkers. The system was processed with EHTransport to calculate the electronic transmission spectra.

Within the HOMO-LUMO gap conductance is predicted to drop off exponentially with length (Figure 6.1). However, splitting of the LUMO at longer chain lengths creates new states progressively closer to the Fermi level, which is fixed in position by the contacts. The lowest energy

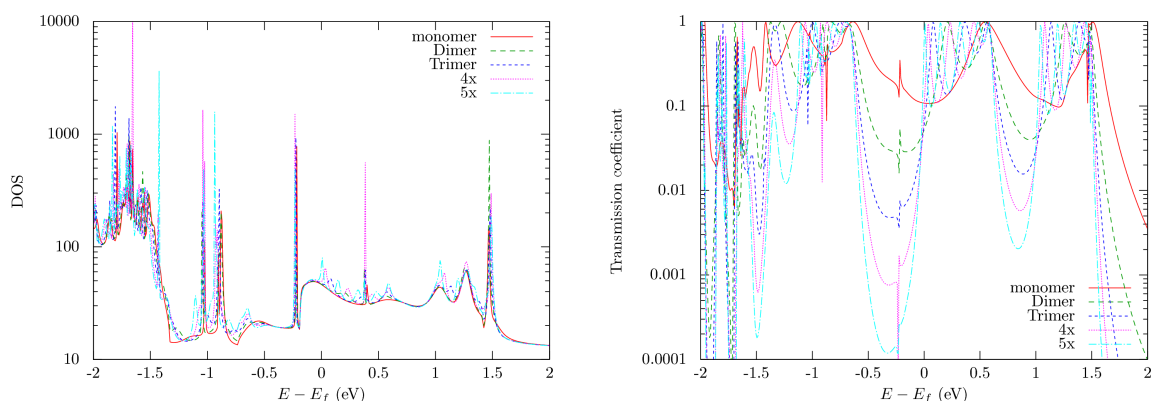


Figure 6.1: The DOS (left) and transmission spectra (right) of a zinc-porphyrin wire of various lengths between gold (111) surfaces.

state attributable to the LUMO shifts from +0.4 eV for the monomer to on top of the Fermi level in a five unit chain. This counteracts the effect of the reduced transmission within the gap, giving the somewhat surprising prediction that the Fermi level conductance of the porphyrin wire increases with wire length for the lengths under study. If the prediction of the Fermi level is incorrect by only a tenth of an eV, however, this prediction would reverse. Slightly below the Fermi level, conductance shows a sub-exponential decay (Figure 6.2). The effects of the LUMO shifting closer to the Fermi level are still felt as the broadened edge of the peak in transmission due to the LUMO makes itself felt.

At 0.15 eV below the Fermi level, the highest energy investigated at which an exponential decay is evident, the decay constant is predicted to be 0.11\AA^{-1} . This is somewhat greater than the 0.04\AA^{-1} seen in experimental work on similar 4C linked Zn-porphyrin chains in [97], and is more on a par with the decay constant they saw when the porphyrin rings are twisted. The experimental structure includes side chains off the porphyrin rings to aid solvation and an additional benzene ring between the porphyrin and gold surface, which makes a direct comparison difficult however. Given this, the match up is reasonable.

The shift in the LUMO was also found by performing calculations on porphyrin wires without the gold contacts or thiol linkers present. The HOMO-LUMO gap¹ was found to shrink from 1.38 eV for the monomer to 0.68 eV for a 5 unit chain (Figure 6.3). This matches up with the

¹In this calculation without contacts the Fermi level is found by filling up available states, and is by definition at the mid point between the HOMO and LUMO, so difference in splitting behaviour between the HOMO and LUMO is not visible.




Figure 6.2: The conductance as a function of wire length of a zinc-porphyrin wire between gold contacts, at several energies. The pink lines are exponential fits to the first 3 data points at each energy.

peaks to unity transmission in figure 6.1. Fitting this to an exponential curve gives the formula $E_{\text{gap}} = 1.90e^{-0.0350L} + 0.665$ eV, where the units of L are Bohr. The fit to the exponential here is very good, with fitting errors of order of one part in a thousand.




Figure 6.3: The energy gap between the homo and luma of a zinc-porphyrin wire as a function of chain length, with no gold contacts or thiol groups present.

From the fitting of the HOMO-LUMO gap, an infinitely long zinc-porphyrin wire would be expected to have a bandgap of 0.665 eV, not significantly smaller than the 5 unit chain. With the effects of the shrinking HOMO-LUMO gap removed, Fermi level conductance may drop off more rapidly at longer chain lengths.

The backbone of four carbon atoms used to link the porphyrin rings here is just one of many

potential linking methods. In addition to the two carbon atom chain used in the previous chapter, porphyrin rings can be fused directly along one edge [98]. The stronger interactions between rings in these cases would be expected to close the HOMO-LUMO gap further, providing further opportunities for use as molecular wires. Indeed, [97] predicts that tape porphyrins would have a decay constant of only 0.02 \AA^{-1} .

6.3 Conductance in the Iron porphyrin system

So far, all porphyrin computations have been carried out using Zinc based porphyrin molecules. Many of the interesting porphyrin based systems found in nature contain different metal centres, such as the iron atom in heme molecules and cytochromes. To investigate the effects of replacing the central metal atom, the two zinc atoms in a porphyrin dimer were replaced with iron atoms. The atomic positions of the remainder of the molecule were left untouched. This new system was then attached to both gold chain and bulk gold contacts, and EHTransport used to compute the transmission spectra of the system.




Figure 6.4: The DOS (left) and transmission spectra (right) of an iron-porphyrin dimer between gold chains.

The primary effect on the system is the shifting of the Fermi level. Interestingly, the effect is in opposite directions for the chain (Figure 6.4) and the bulk (Figure 6.4) contacts. No major changes in the structure of the transmission peaks is visible. In the case of both contacts, Fermi level conductance is predicted to be lower for the Iron porphyrin molecule.

An examination of the conduction pathways show that little transport takes place through

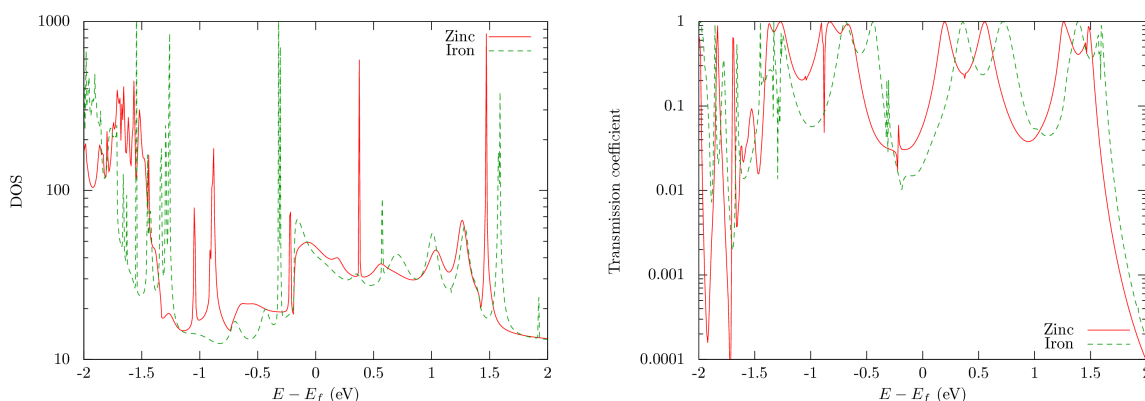


Figure 6.5: The DOS (left) and transmission spectra (right) of an iron-porphyrin dimer between (111) surfaces.

the central metal ion, preferring to move around the outside of the ring. An inspection of the orbital-character of the transport suggests that conduction takes place through p-orbitals. This is expected for this sort of conjugated organic molecule; conduction typically takes place through the π molecular orbitals which are made up of a combination of the p orbitals of each atom pointing out of the molecular plane. More interesting is the conduction through the gold contacts. Despite the earlier symmetry arguments, the gold shows strong s character. Replacing the zinc atom with an iron atom has little effect on the important conduction pathways, but changes the orbital character of the gold orbitals completely. Projecting the DOS onto the gold atoms shows a small shift in the gold d-band relative to the Fermi level, pushing it into range. (Shown also in the figure 6.4, where zinc-porphyrin is LUMO-based but iron is HOMO-based.)




Figure 6.6: Colour coded representation of the orbitals involved in conduction through a zinc-porphyrin dimer. Red shows s character, green p and blue d.

At the end of the self-consistent loop, the iron atom has an excess charge of $+1.1 e$ when connected to the bulk gold contacts. This is a much greater charge than that found on any other atom in the system, which all have magnitudes of less than 0.3 electrons. By contrast




Figure 6.7: Colour coded representation of the orbitals involved in conduction through an iron-porphyrin dimer. Red shows s character, green p and blue d.

the zinc atom in zinc-porphyrin has a self-consistent charge of -0.2 e. Although much greater in magnitude than the charge on any other atom in the system, the charge on the iron atom is still less than the most common $+2$ or $+3$ oxidation states that iron generally forms in experiment.

6.4 Effects of water molecules of molecular conductance




Figure 6.8: The structure of the water porphyrin dimer surrounded by water.

The environment of a molecule can have considerably effects on its properties [99]. Water molecules have ionic nature, with the oxygen atoms gaining negative charge and the hydrogen positive. This may in turn induce additional local charges in a molecule in aqueous solution, and earlier results of this chapter have already suggested that profound changes in behaviour can be achieved with small changes in Fermi level. Previous theoretical work has also suggested that the presence of a solvation shell around a porphyrin wire modifies the length dependence of conductance from exponentially decreasing with wire length to invariant [100].

To investigate the effects of water on the zinc porphyrin dimer, a molecule was placed between bulk gold contacts, and the vacuum space in the unit cell populated with water molecules. To simulate the maximum polarisation that would occur when a voltage was applied, in one test case

the water molecules were aligned to point along the z axis. The water molecules were positioned randomly by picking points between the contacts, checking that no other atoms were within 4 Bohr of an oxygen atom or 3 Bohr of a hydrogen, and placing a water molecule in that position should it fit. This resulted in a water density approximately 60% that of water under standard conditions, with a total of 59 water molecules surrounding the porphyrin dimer.

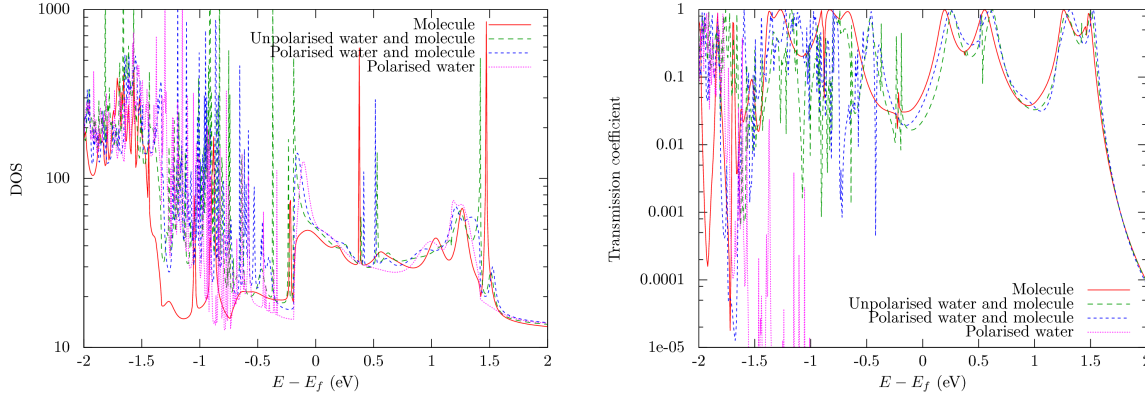


Figure 6.9: The DOS (left) and Transmission spectra (right) of a system consisting of a zinc porphyrin dimer between bulk gold contacts with and without water present, as well as pure water with no porphyrin present.

No effect was visible around the Fermi level after adding water molecules in similar positions in EHT. However, the Fermi level was shifted, by an amount that was small, but which still would have had an effect on the wire properties based on the previous chain length results. Below the Fermi level there are profound effects on conductance. In this region pure water showed a significant level of transmission, so additional effects there are to be expected.

Compared to isolated water molecules, water molecules close to the gold surface became negatively charged (Figure 6.10). Although the gold surfaces as a whole became positively charged, the top layer also held additional negative charge in comparison to the porphyrin system without water present. These changes to the surface of the gold contacts will change the way it links up to the porphyrin. Previous results have shown that the presence of a solvent can modify the work function of a metal surface, altering conductivity behaviour [101], consistent with the modification of the gold surface here.

In addition, since the entire gold electrode has become charged, the principal layer which connected up to the leads has also been affected. The assumption that the link between the final

lead principal layer and the scattering region is identical to that between two principal layers of the leads is not accurate in the situation of a charged electrode.




Figure 6.10: A charge difference plot showing the difference in real space compared to that of the porphyrin system without water present for the contacts and dimer, and an isolated water molecule for all water molecules in the system. Red denotes excess electrons, and blue a deficit.




Figure 6.11: Bond-currents plot at the Fermi level of a zinc-porphyrin dimer between bulk gold contacts, and surrounded by polarised water.

Bond current calculations at the Fermi level, where the transmission spectrum was relatively unaffected by the presence of water, show no change from the dehydrated case. 2 eV below the Fermi level where the transmission spectra is greatly affected by the presence of water the bond current plot also shows more significant effect. Most of the current now enters the porphyrin ring via a water molecule rather than the thiol linker.

6.5 Conductance in the cytochrome b562 protein

Cytochrome b_{562} (IPR009155) is a protein expressed in the bacterium *Escherichia coli*. It is made up of four helical columns, and contains a heme group embedded in its structure. As with cytochromes in general, it is believed to have an electron carrying role within the bacteria, taking part in redox reactions, and so may show interesting electrical conductivity behaviour. Several




Figure 6.12: Bond-currents plot 2 eV below the Fermi level of a zinc-porphyrin dimer between bulk gold contacts, and surrounded by polarised water.

variant structures have been investigated; the one chosen for analysis was a solution structure in an oxidised state [102], a form in which a single protein can be isolated and for which atomic positions were readily available. This structure contains a total of 1750 atoms, made up of carbon, oxygen, sulphur, nitrogen hydrogen and iron (Figure 6.5), and measures 49.6Å from end to end. To make contact with the gold leads, two hydrogen atoms on opposite sides of the protein exterior were replaced with sulphur atoms to form thiol linkers. The part of the gold contact within the scattering region totalled 8 atoms in each lead, with a lead principal layer of 4 atoms. This final scattering region structure comprised 1766 atoms and 4500 orbitals.




Figure 6.13: Cytochrome b_{562} , showing the position of the heme group relative to the protein backbone.

The first stage of the computation, computing the overlap matrix of the system, took approximately 10 hours 45 minutes of CPU time. Computing the self consistent Hamiltonian matrix took a further 3 days of CPU time. The final stage of the computation, calculating the DOS and transmission spectrum at 1001 energy points, took another three and a half days of CPU time.

The full computation consumed a total of 7 days of CPU time running on a single CPU core, and consumed approximately 4GB of memory at its peak.




Figure 6.14: The DOS (left) and transmission spectra (right) of Cytochrome b_{562} between gold chains.

The results show no peaks approaching G_0 , but within the -2 eV to +2 eV window around the Fermi level three broad regions of comparatively high conductance are visible, at positions matching visible molecular states in the DOS (Figure 6.14). Of these, the closest to the Fermi level is centred at around -0.35 eV, and averages of the order of $10^{-6}G_0$. An attempted calculation of bond currents in the middle of this peak failed, with current appearing to vanish inside the gold lead. This non-physical result suggests that the algorithm for computing bond currents suffers from issues with numerical noise on either systems of large size or of low conductance. A modified algorithm with increased numerical stability would be required to identify the conduction pathway of this system.

Previous experimental results suggest a peak Fermi level conductance of the order of $10^{-5}G_0$ across a single protein molecule [103]. EHTransport does not predict transmission of that magnitude at the Fermi level. The Fermi level is fixed by the gold chain contacts however, and would be shifted in a bulk gold calculation. The size of the gold contracts which would be required to create a periodic gold surface while minimising interactions between periodic images for the protein would be too large for such a system to be processed even with EHTransport using computational resources currently available.

6.6 Conclusion

EHTransport was able to provide good quality transmission spectra for larger systems, including a complex system containing thousands of orbitals. However, one auxiliary part of EHTransport broke down when presented with a system of this large size; bond currents could not be calculated.

Porphyrin chains were predicted to have a Fermi level conductance that does not decrease with chain length, making them potentially suitable for use as molecular wires. Iron-porphyrins were not predicted to behave in ways much different from zinc-porphyrin, but the effects of iron ions of different oxidation states were not captured, which may be important as experiments on iron porphyrin show hysteresis in measurements of conductance against voltage that do not show up with zinc-porphyrin [104], providing strong evidence that the iron was switching oxidation states. The presence of water did not greatly affect transmission around or above the Fermi level, but showed most effects below where the porphyrin transmission was additive to that of the water. Since solvation was expected to have a more significant affect, self consistent EHT may not be sufficient for capturing the interactions involved.

Chapter 7

EHT calculations of DNA based systems

7.1 Introduction

Deoxyribonucleic acid (DNA) is possibly the best known of all complex organic molecules, and is at the heart of all known complex multi-cellular life. A DNA molecule is made up of carbon, oxygen, nitrogen, phosphorus and hydrogen atoms and is formed from two individual strands, each of which has a sugary backbone with branches (bases) of four different types jutting off at regular intervals. Each base of one strand pairs up with a complementary base on the other via hydrogen bonds, causing the base pairs to give a ladder like structure. These backbones twist together into a helix structure (Figure 7.1).




Figure 7.1: A DNA strand ten base pairs in length, of sequence ATATATATAT.

There are 4 different naturally occurring bases; adenine (A), cytosine (C), guanine (G) and thymine (T). A fifth base, uracil (U) also occurs in nature, but is typically found incorporated into RNA strands rather than DNA, where it substitutes for thymine. Cytosine always matches up with guanine, and adenine with thymine. When stating a DNA base pair sequence, it is the list of bases attached to a single backbone that is given. (e.g. the sequence AGCT would imply 4 base pairs; A-T, G-C, C-G, T-A) In addition to the base pair sequence, the helix itself take several forms, identified by the labels A, B and Z. Each with differing twist characteristics. (Z-DNA in fact twists in the opposite direction to A-DNA or B-DNA.) The most common form in nature is B-DNA [105], with Z-DNA occurring briefly in response to certain biological actions [106]. A-DNA is most commonly seen when DNA is less hydrated than occurs inside healthy cells.

There have been many experimental reports on the conductivity of DNA [107–111], with results ranging the full spectrum from ohmic conductance through to semi conductance and insulator behaviour. There is no suggestion that the majority of these results are incorrect, and that DNA always exhibits constant behaviour; rather, DNA seems able to take any behaviour,

dependent on base pair sequence, contacts, structure, environment and other properties. Experiments have shown results that suggest a repeating CG unit has higher conductivity than an AT unit, that DNA conducts better in solution, or that a single mismatched base pair can drop the conductivity of a molecule by an order of magnitude or more. Work on modification of base pairs to incorporate metal ions has also resulted in improved conductance [110].

Previous theoretical work on the conductance of DNA has mostly taken place using either tight binding models or, on small systems consisting of a very limited number of base pairs, ab-initio techniques [25]. In this chapter, my self consistent EHT based code is applied as a middle ground, able to cope with larger system sizes than ab-initio techniques, and hopefully able to produce more descriptive results than the simpler tight binding models. I will compare my results to those found by other theoretical techniques, as well as to experiments.

7.2 Structure

DNA structures were generated via the web interface of the application 3DNA [112] (W3DNA). This application can generate DNA structures in several ways. For our calculations, an initial base pair sequence was entered, and a pdb file containing atomic coordinates was created. This pdb file was fed back into W3DNA to generate a base pair parameter file, containing the sequence and positioning of the bases. This was modified as required to introduce mismatched base pairs or other alterations, and then fed back into W3DNA to generate a second pdb file, containing both atomic positions and bond information¹. This file was loaded into Avogadro for further alterations, such as the addition of thiol groups for forming the contacts. Since the output files of W3DNA do not contain the positions of hydrogen atoms, these were also added in via Avogadro's automated algorithm, assuming a neutral pH. The thiol groups were attached to the 4' carbon atom in the backbone, which is bonded to only a single hydrogen, removing ambiguity about positioning of the sulphur atom and ensuring consistency.

Using the semi-empirical geometry optimization algorithms in Avogadro results in noticeable changes in the DNA structure. In particular, base pairs were found to buckle due to the way unusually long bonds in the DNA backbone were handled. For this reason, the structures from

¹Since the structure generated from the parameter file by W3DNA does not exactly match the structure used to create the PDB file the parameter file was created from, it is important to follow this procedure even for simple base-pair sequences which do not require modifications to the parameter file, in order to ensure consistency.

3DNA are used unoptimized, except for the positions of the additional thiol groups. Tests on a CGCGCG system show that this buckling leads to greatly reduced simulated conductivity (Figure 7.2).




Figure 7.2: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads using the original W3DNA generated structure and a version optimized using UFF forcefields.

Another point of note is that the backbone contains complexes made up of one phosphorus atom and four oxygen atoms. Two of the oxygen atoms are bonded to further carbon atoms, and are a part of the backbone. A single double bond is delocalised around the other two oxygen atoms, leading to an ambiguity in the placement of a hydrogen atom, which may be found attached to either oxygen atom. It is these complexes which give DNA its acidic nature; the hydrogen atom in this complex can become detached when the DNA is in solution. To measure the size of this effect, additional calculations were performed with the hydrogen atoms swapped between the oxygen atoms, and again with the hydrogen atoms completely removed. (Figure 7.3) Moving hydrogen atoms from one possible position to another has very little effect. Removing them has a greater effect, and since DNA is, as the names suggests, naturally acidic, some portion of them will detach if the DNA is in solution. For the rest of this chapter, the hydrogen atoms were left in their original, default position.

Conduction through a DNA molecule of repeating base pair sequence CGCG is believed to take place through the stacked π bonds of the base pairs [113]. Since a DNA molecule consists of two interwoven strands, there is a choice available as to which of the two backbones the leads are connected to. Given two chain leads, either both can be connected to the same backbone, or each




Figure 7.3: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads using various possible positioning of hydrogen atoms.

could be connected to a different backbone. Calculating the transmission spectra for the case where two gold contacts are made to the ends of the same backbone to the case where the contacts are attached to opposite backbones shows very little difference. (Figure 7.4) Nor indeed does much change in the regions of high transmission when the contacts are made to both backbones simultaneously, although transmission is improved through regions of low transmission. This provides strong evidence that conduction is indeed taking place through the base pairs, as if the base pairs were non-conductive then far lower transmission would be expected in the case of single gold leads attached to different backbones. Further runs in this chapter make use of two gold leads, attached to the same backbone.

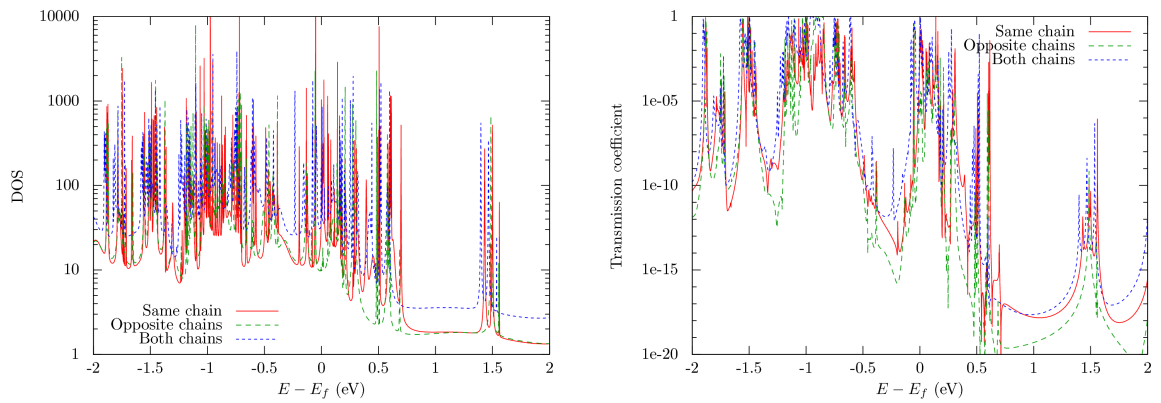


Figure 7.4: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads using several lead positioning variations.

Many of the peaks in these plots have a single point an order of magnitude greater or more

than their neighbours. The smooth curves at the base of the peaks provides evidence that the peaks are real, and not numerical artefacts of the type noted in chapter 4, but the lack of points around the peaks of the curves makes it difficult to accurately determine the heights. An additional run was made with 5001 data points, magnifying the region from -0.3 to $+0.3$ eV (Figure 7.5), but although this allowed identifying individual peaks the tops of these peaks were still only occupied by single data points. A couple of these peaks were analysed further via a script to locate their centres, showing that the peaks do have smooth caps at the top not necessarily at $1G_0$. These peaks correspond exactly to peaks in the density of states, suggesting that conduction is occurring through these particular states despite the states being weakly coupled to the contacts and hence being subjected to very little broadening.




Figure 7.5: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads, exploring a narrow energy range around the Fermi level.

7.3 Effects of DNA sequence on molecular conductance

Experimentally, a sequence of repeating CG base pairs has been shown to have similar conductivity to a chain of repeating AT base pairs [114,115]. However, the addition of an AT sequence into the center of a repeating CG chain has been shown to decrease conductivity by an order of magnitude. Under at least some conditions, the two also have differing length dependence, with CG chains showing conductance scaling as the inverse of length but AT chains showing an exponential decay with length [116], suggesting a different mechanism for transport.

EHTransport predicts that replacing the central CG unit within a CGCGCG sequence with

an AT unit drops conductance around the Fermi level significantly (Figure 7.6), in some regions by considerably more than one order of magnitude. Despite this, the peak at the Fermi level remains relatively unaffected.




Figure 7.6: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads showing the effect of adding replacing the central CG sequence with an AT in a CGCGCG strand.

It should be noted that B-DNA is not normally stable in a vacuum, and experimental results are taken in aqueous solution. As shown in the previous chapter, the charge transfer resulting from the presence of water can effect results. Furthermore, in solution it is expected that a certain percentage of the hydrogen atoms will detach, which also has an effect on the transmission spectra. A further effect however is that dehydrated DNA tends to take the A-DNA form. Experimental studies on DNA in a vacuum have shown a lower conductance than those made in solution [111], and theoretical calculations suggest that this transition between A-DNA and B-DNA is at least partially responsible [117], due to poorer overlap between the Π orbitals in neighbouring base pairs.

There are significant differences in both DOS and transmission between A-DNA and B-DNA (Figure 7.7), with the B-DNA having considerably greater conductance near the Fermi level. For this reason alone, it would be expected that DNA in solution would conduct significantly better than DNA in a vacuum, so the EHT measurements are again consistent with experimental results.




Figure 7.7: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads showing the difference between A-DNA and B-DNA structures.

7.4 Effects of DNA structural failings on molecular conductance

Damage may occur in a DNA strand, resulting in base pairs that are mismatched or even missing completely. Given the large sizes of typical DNA molecules, a single base pair error may not be immediately destructive to the molecule as a whole. Since the bases are different sizes, not all combinations are physically possible. Experimental results have suggested that a single mismatched base pair can cause an order of magnitude drop in conductance [118].




Figure 7.8: The DOS (Left) and electrical transmission spectra (Right) of a DNA molecule trapped between gold leads of a CGCGCGCGCG sequence, in which the 5th base pair in the chain has either been replaced with a C-C combination, or removed completely.

Some drop is seen when using a C-C mismatched pair instead of the original C-G (Figure 7.8), but the difference is not great. One possible explanation is that in reality the presence of a mismatched (or missing) base pair may cause considerable alterations to the backbone structure

around the site of the error.

There is a far greater effect from the removal of a base pair; the overlap between the pi orbitals of one base pair and those of its next nearest neighbour is obviously substantially smaller than that of its nearest neighbour, and this has a substantial effect in the transmission. In both cases the DOS remains relatively unchanged.

7.5 Conclusion

The transmission spectra of DNA molecules contains many more features in comparison with the porphyrins, as might be expected given the larger problem size. Even with the porphyrins the exact position of the Fermi level could have a very significant impact on the conductance of a molecule, but for a given Fermi level the conductance behaviour could be well defined. In DNA based systems, the lack of smooth regions, with transmission fluctuating by orders of magnitude over small energy ranges, makes categorisation of conductance behaviour difficult, but nevertheless the broad behaviour in the region around the Fermi level can be seen and compared between multiple systems.

In all cases of comparison here EHT results are in qualitative agreement with previous experimental results. Some quantitative predictions, e.g. systems containing mismatched base pairs, are smaller than experimentally observed effects. DNA strands of up to 14 base pairs in length have been analysed with little numerical difficulty. Although this cannot match the size of systems amenable to processing by the tight binding models, which typically use Hamiltonian matrices with size of one or two elements per base pair, it compares very favourably to the computational resources which would have been required by traditional DFT techniques, where some research still focuses on structures of only 4 base pairs [119].

Chapter 8

Conclusion

8.1 Summary

In this thesis a completely new computer code named EHTransport has been developed, tested, and then utilised to provide results on porphyrin and DNA based systems trapped between gold electrodes. Performing comparisons against the existing code Transiesta has shown that EHTransport can generate useful results for the porphyrin systems studied, and have shown the importance self consistency in the EHT calculations. The results provided on DNA have further validated the EHT approach by showing agreement with experiments related to the effect of DNA structure on its electronic transmission. Results on a heme-based cytochrome have further shown that the code can cope with basis sets of thousands of orbitals.

Additional support for utilising the self consistent extended Hückel method alongside Green's functions for the calculation of transport properties lies in the commercial product the Atomistix ToolKit, developed by QuantumWise. Alongside DFT techniques, this toolkit has also recently added support for a self consistent extended Hückel approach [63]. This has resulted in few published works in comparison to the older and more familiar DFT component of the code, but a comparison of the two on one system has concluded that there is good agreement [120].

Extended Hückel is not a first principles method, and is reliant on an empirically chosen parameter set. Although the addition of self consistency improves the transferability of this parameter set, care still needs to be taken that the parameters chosen are suitable for the system under study. This can include comparisons with first principles methods for reduced systems of similar type. The method used here for computing the self consistent parameters has shown to result in correct charge transfer behaviour, but as other parameters constructed for use in non-self consistent calculations were used directly, further improvement is possible in this area.

A zinc-porphyrin dimer was found to have some suitability for use as an organic transistor with a 1 to 10,000 on/off ratio, should a method of attaching a gate contact in such a way as to provide control over the dihedral angle. It was also found to have some potential as a molecular wire, with conductance above the Fermi level being largely independent of wire length. The zinc-porphyrin dimer has potential for encoding a single bit of information within its dihedral angle, given appropriate structural modifications to ensure stability.

Transmission through DNA molecules was found to be complex, and structure and base pair sequence changes had very significant results. However even if it was feasible to switch out base

pairs at a rate high enough for a useful transistor, the on/off ratio would be poor. Although DNA computers have been investigated, these are biological in nature. Similarly, any attempts to use DNA as a memory storage medium would be likely to focus on its base pair sequence, rather than its electrical properties.

8.2 Further work

EHTransport is functional in its current form, but opportunities for improvement exist. Some parameters, including those that control mixing during self-consistency, are currently hard coded and cannot be modified via the input file. Replacing the linear mixing with a Pulay mixing scheme would speed convergence of the self consistent EHT stage. A more efficient implementation of the NEGF stage is possible, following the procedure used by Transiesta. As the target application of EHTransport includes large systems, for which k-points are unlikely to be required, additional use could be made of the fact that Hamiltonian matrices will be real and symmetric, providing further efficiency improvements. This work used the standard Hoffman parameter set for all atomic species other than gold. Making corrections to this set to account for the self consistency in use may provide more accurate results.

A mechanism for fixing the charge on specific atoms during self consistency would have allowed investigation of the effects of the oxidation state of an iron-porphyrin dimer. With such a feature in place, comparisons of different oxidised states of iron could be made for the iron-porphyrins, and analysed for the potential to account for the hysteresis seen in experiment. By fixing the final principal layer of the lead to zero excess charge, any issues with linking the scattering regions to the leads could also be avoided.

Although some work has been done to parallelize EHTransport through the use of OpenMP, further work is required to allow for its optimal use on shared memory systems with many processors. Run on the same hardware as was used for Transiesta in this work should allow for the processing of systems containing of the order of ten thousand orbitals in time scales of the order of a week. This will require parallelization of further parts of EHTransport which are currently run serially. Similarly, some of the most time consuming operations are manipulations of large matrices, and are amenable to stream processing algorithms [121]. This could be exploited to provide a further speed up on desktop class equipment equipped with suitable specialised

compute hardware or graphics cards.

Although EHTransport uses the more computationally efficient EHT in the place of DFT in the construction of a Hamiltonian, the full NEGF formalism is used in computing the transmission spectra. This is made somewhat more efficient by the reduction in the size of the basis set, but the same could be achieved with DFT by using a non-polarised basis set. Approximations in the NEGF formalism could be sought to reduce the computational complexity of this part of the problem.

Additional interesting investigation of porphyrins would include different wire length dependence with differing linkers between the porphyrin rings. Experiments often include additional side chains to aid solvation, and groups between the porphyrin ring and thiol linker for ease of synthesis, which may have effects on conductance, and could be simulated.

Appendix A

EHTransport manual

A.1 Introduction

EHTransport is a C++ code designed to perform transmission spectra calculations of molecules or other finite length structures connected between two semi-infinite leads. It is designed to be invoked from the commandline, and takes a single commandline argument giving the path to the input file. The input file uses the same FDF format as SIESTA, and it is a design goal for the same input files to be transferable between the two programs; although the two parsers do not have identical behaviour, and the options required for a given calculations may vary between the two, the features of the FDF file format still allow for the files to be transferable.

The FDF file format consist of a list of keywords and values, one per line. A `#` on a line signifies the beginning of a comment, and any text between it and the beginning of the next line is ignored. Keywords are not case sensitive. The basic types of a value include strings, which can be any arbitrary text sequence, (when a string appears in a block it may not contain whitespace, and no string may contain the `#` character) integers or floats. Where an input value has units, it is specified by a float followed by the name of the unit. For energies, permissible units are 'ryd' or 'ev'. For lengths, 'bohr' or 'ang'. For temperatures 'k'. Some keywords take a further keyword as their argument; these are again case insensitive. As a special case of this are options which can only be enabled or disabled, in which case 'true', '.true.', '1', 'yes', 't' or 'y' are used to enable the option, and 'false', '.false.', '0', 'no', 'n', and 'f' will disable it.

Where a keyword requires multiple arguments, it is specified as a block. '%block <keyword>' must appear on a line of its own. In the following lines the set of parameters are given, separated by whitespace. Finally, '%endblock <keyword>' must appear on a line of its own.

The following is an example of a correctly formatted input file, as an example of the syntax. The names and effects of the available keywords will be described in the following sections.

```
#Perform an EHT calculation on a monotomic gold chain
```

```
JobType Huckel
```

```
SystemLabel Lead
```

```
SelfConsistency None
```

```
CellCutoff 8 Bohr
```

```

NumberOfSpecies 1

%block Chemical_Species_label
    1 79 Au
%endblock Chemical_Species_label


AtomicCoordinatesFormat NotScaledCartesianBohr

NumberOfAtoms 6

%block AtomicCoordinatesAndAtomicSpecies
    0.000000000000000 0.000000000000000 2.25510159197654 1
    2.76192410899580 1.59459763004823 6.76530477592962 1
    2.76192410899580 -1.59459763004823 11.27550795988271 1
    0.000000000000000 0.000000000000000 15.78571114383579 1
    2.76192410899580 1.59459763004823 20.29591432778887 1
    2.76192410899580 -1.59459763004823 24.80611751174195 1
%endblock AtomicCoordinatesAndAtomicSpecies


LatticeConstant 1.0 Bohr

%block LatticeVectors
    5.52384821799160 0.000000000000000 0.000000000000000
    2.76192410899580 4.78379288340561 0.000000000000000
    0.000000000000000 0.000000000000000 27.06121910371850
%endblock LatticeVectors

```

The following are the global keywords, applicable to all job types:

JobType (Huckel | Transport | UnitTest | SiestaTransport | PrecalcLeads)

Default: None

Set the type of calculation which EHTransport will perform. Huckel is used to perform an EHT calculation. Transport is used to perform NEGF calculations on the output of an earlier EHT calculation. UnitTest is used to perform some tests for successful compilation and functionality of the EHTransport code. SiestaTransport is used to perform NEGF calculations on the output

of a transiesta run. PrecalcLeads is used to calculate and store to disk the self-energy matrices of an earlier EHT calculation.

SystemLabel <string>

Default: label

Sets the name of the job. This is used in the names of any output files, to allow multiple runs in a single directory without overwriting each others results. For example, if this option is specified as 'test', running an EHT calculation will result in the creation of a file called 'test.EHT' in the working directory containing the calculation results.

Any output files are written to EHTransport's working directory, which will not be changed from the initial invocation of the program. The exact set of output files produced will vary based on the job type and the selected options. Calculation progress is written to standard output, to provide feedback on the progress of a calculation.

A.2 Performing an Extended Hückel calculation

A.2.1 Setup, input and output

An EHT calculation requires no input beyond the FDF file. This file must at the least specify the positions and types of every atom in the system. As output, the file '<systemlabel>.EHT' is created. This file is in a binary format, and stores the Overlap and Hamiltonian matrices for each k-point, along with atomic charges if calculated and the parameters the system was set up with. Other output files may also be generated depending on the selected options.

During a long calculation, you may occasionally want to interrupt EHTransport without killing it. Once per self consistent loop, EHTransport will check for the file 'override.txt' in its working directory, and if found will read from it the first line. This can contain a command with which to alter EHTransports behaviour. Currently the only supported command is 'shutdown', which will immediately terminate the self-consistency loop as if the limit on the maximum number of iterations had been hit.

A.2.2 Valid keywords

The following are the list of keywords applicable to EHT calculations:

UseSaveData <bool>

Default: False

If this option is enabled, EHTransport will not perform an EHT calculation, but rather will attempt to read in an existing .EHT file. The file that will be read will have the same name as the file that would have otherwise been created. If this file could not be found, or could not be read successfully, EHTransport will not attempt to perform the EHT calculation from scratch, but rather the job will fail.

NumberOfSpecies <integer>

Default: None

Specify the number of different atomic species taking part in the calculation. Note that if multiple parameter sets are in use for a single atom, each variant counts as a separate atomic species. This option is required if UseSaveData is not enabled, or ignored if it is.

ChemicalSpeciesLabel <block>

Default: None

Specify the atomic species used in the calculation. This must be provided as a block, with a number of values equal to three times the number of species in the calculation. For each species, two integers and a string must be given; the id of the atom in the calculation (numbered from 1 up to NumberOfSpecies), then the number of electrons in the atom, followed by the atomic symbol. Note that the number of electrons field is actually ignored in EHTransport, and is only required to maintain SIESTA compatibility. This option is required if UseSaveData is not enabled, or ignored if it is.

NumberOfAtoms <integer>

Default: None

Specify the number of atoms in the system. This option is required if UseSaveData is not en-

abled, or ignored if it is.

AtomicCoordinatesFormat (NotScaledCartesianBohr | NotScaledCartesianAng)

Default: NotScaledCartesianBohr

Provides the coordinate system in which the atomic coordinates will be given. Only Cartesian options are supported, with the results that numbers given for atomic coordinates are multiplied either by 1 Bohr or by 1 Angstrom to give the final coordinates. Note that the scaled coordinate options available in SIESTA are not supported in EHTransport. This option is required if UseSaveData is not enabled, or ignored if it is.

AtomicCoordinatesAndAtomicSpecies <block>

Default: None

Provides the coordinates and species of each atom in the system. This must be provided as a block, with a number of values equal to four times the number of atoms in the calculation. For each atom, three floats and an integer must be given, specifying the position of the atom in the system, and the id of the atomic species of the atom. This option is required if UseSaveData is not enabled, or ignored if it is.

LatticeConstant <Length>

Default: 1 Bohr

This is a scaling factor for the unit cell size. All values in the unit cell matrix are multiplied by this length to give the final values. This option is ignored if UseSaveData is specified.

LatticeVectors <block>

Default: See description

This provides the size of the unit cell for the calculation. It can be specified either as three float values, which will be interpreted as the lengths of the x, y and z vectors of an orthogonal unit cell, or as nine float values which are interpreted as an orthogonal unit cell. If the block is not present, EHTransport will use an orthogonal unit cell of a size large enough to fit in all atoms of the system with no atoms within 10 Bohr of each face of the cell. This option is ignored if

UseSaveData is specified.

NumKPoints <integer>

Default: 0

If greater than zero, this specifies the number of explicit k-points that will be provided. (See 'KPoints' keyword) This option is ignored if UseSaveData is specified.

KPoints <block>

Default: None

Specifies the set of k-points at which the EHT calculation will be performed. This must be provided as a block, with a number of values equal to four times the number of k-points in the calculation. For each k-point, 4 floats must be given, the first 3 providing the location of the k-point in reciprocal space (values between 0.0 and 1.0) and the fourth value giving the weighting of the k-point. This option is required if NumKpoints is specified, or ignored if it is not.

KGridMonkhorstPack <block>

Default: Gamma point calculation only

The Monkhorst-Pack mesh to use for generating the set of k-points for a system. This must be provided as a block containing nine integers. The set of k-points created will be identical to those of SIESTA for a given Monkhorst-Pack mesh and unit cell, aside from SIESTA picking up on more available symmetry and merging some k-points. This option is ignored if UseSaveData or NumKPoints are specified. If neither this option nor NumKPoints are specified, the calculation will be performed at the gamma point.

AtomicModes <block>

Default: See description

Specifies the parameter set to use for each atomic species. This must be provided as a block containing two integers for each atomic species whose value you wish to change. The first integer is the atomic number of the species you wish to change, and the second is the id of the parameter set you wish to use. The default is the Hoffmann parameter set for all atoms other than gold,

whose default is the Cerda parameter set. This option is currently supported only for gold, for which the valid values are 1 (Hoffman) 2 (Cerda) or 10 (Custom parameter set, fitted to a monatomic gold chain DOS). This option is ignored if UseSaveData is specified.

SelfConsistency (None | Simple | Poisson)

Default: None

Specifies the type of self consistency which will be employed in the EHT calculation. Note that Poisson mode is not supported for a non-orthogonal unit cell. This option is ignored if UseSaveData is specified.

TS.Voltage <float>

Default: 0

Specifies the bias across the unit cell in the z direction in volts. This option is ignored if SelfConsistency is not Poisson, or if UseSaveData is specified.

CellCutoff <length>

Default: 8 Bohr

Specifies how many neighbouring unit cells are considered when constructing the overlap matrix of the system. Any cell from which a line of less the specified length can be drawn with one end in the central unit cell and the other in the neighbouring cell will be included in the calculation. If set to 0 periodicity will be disabled. (Note that in this case, it makes no sense to perform the calculation anywhere other than the gamma point.) This option is ignored if UseSaveData is specified.

InternalSupercell <block>

Default: None

Specifies the exact size of the internal supercell to use when constructing the overlap matrix of the system. This must be provided as a block containing three integers, representing the number of additional cells to include in the x, y and z directions. This extends in both directions; for example, 0 would represent no periodicity in that direction, and 1 would result in a total of

3 unit cells in the given direction. This option is ignored if CellCutoff or UseSaveData is specified.

OverlapCutoff <length>

Default: 20 Bohr

Specifies the range at which all overlaps between orbitals is assumed to be 0. Reducing this option will decrease computation time, but will result in poorer calculation results. Too low a value may result in the calculation failing with an error that the overlap matrix was not positive definite. This option is ignored if UseSaveData is specified.

DumpSupercell <bool>

Default: False

If this option is enabled, a file called 'supercell.xyz' will be created, containing the complete internal supercell used by EHTransport. (i.e. the unit cell for which the calculation was performed, along with all neighbouring cells used in calculation of the overlap matrix.) The file will be in xyz format, with lengths in Bohr.

SaveEnergies <bool>

Default: False

If this option is enabled, an output file called 'energies' will be created containing the values of all eigenvalues of the system Hamiltonian. The values for all k-points will be mixed together and sorted into ascending order. Values are in eV, and relative to the Fermi level.

SaveAtomicCharges <bool>

Default: False

If this option is enabled, an output file called 'charges' will be created containing the relative charge on each atom in the system. Charges are given in fractions of an electron, one per line, in the same order as which the atoms were specified in the AtomicCoordinatesAndAtomicSpecies section.

BandLines <block>

Default: Don't save any bandlines

If specified, make a band plot for the system. Two output files will be created; '`<systemlabel>.bands`', which contains a list of energies at each point in k-space along each band, and '`<systemlabel>.plotbands`', which contains a gnuplot script which can be executed to generate a plot of the band structure. This must be provided as a block containing 5 values for each path in k-space for which the bands are to be plotted. For each path, the first value must be an integer specifying how many points to pick along the line. The next 3 should be floats, containing the end point of the line in reciprocal space. (Valid values are between 0.0 and 1.0) The final value is a string, which will be used for the labelling of the point in the gnuplot script. The start of each line is taken to be the end of the previous line. To avoid point duplication, the first point on the line is skipped, thus the first entry should normally be '1 0 0 0 \Gamma' in order to get the gamma point.

BandPlotCount <integer>**Default: Use all eigenstates**

This option specifies the number of bands to include in the bandplot. This allows for removing the upper bands, which will be very poorly computed, and due to high energies will mask the interesting part of the plots. This option is ignored unless BandLines is specified.

BandPlotFermi <energy>**Default: The calculated system Fermi level**

This option allows overriding the Fermi level of the system for the purposes of band plotting. This option is ignored unless BandLines is specified.

ProjectedDensityOfStates <block>**Default: Don't save projected density of states**

If specified, an output file called '`<systemlabel>.pdos`' will be created, containing the DOS of the system and potentially a number of PDOS projected over different atoms or orbitals. This must be provided as a block containing three floats followed by an integer. The first float is the minimum energy at which the DOS will be sampled. The second is the maximum energy, and

the third is the amount of broadening to apply to each eigenvalue when computing the DOS. All energies must be provided in eV. The final integer is the number of data points at which the DOS is to be sampled.

PDOSMergeOrbitals <bool>

Default: True

If specified, the PDOS will be collapsed from a per-orbital array to something that is stored per-atom. This has an effect on atom ids in the PDOSAtoms block; if a system contains 2 gold atoms and this option is specified 'range 1 2' will calculate the complete DOS of the system. If this option is not specified, it will compute the PDOS projected onto the s orbital and first p orbital of the first gold atom. Ignored if ProjectedDensityOfStates is not specified.

PDOSAtoms <block>

Default: See description

Specifies the list of projections on which to calculate the PDOS. This must be provided as a block. The first item in the block is the number of projections to be made. For each projection the block first contains a float which is the factor by which to scale the projection, followed by a keyword which may be one of 'all', 'range', 'ranges', 'atom', or 'natom'. 'all' requires no further values, and computes the complete DOS. 'range' requires two further integers representing the first and last atom (or orbital, if PDOSMergeOrbitals is false,) on which to project the DOS. 'ranges' requires an integer that specifies the number of ranges to include, followed by that many pairs of integers to specify the ranges. 'atom' and 'natom' project over an explicit list of atoms, and require one integer to specify the number of atoms (or orbitals) to project over, followed by that many further integers to specify the atoms. 'atom' will sum the PDOS of each of these atoms, whereas 'natom' will average it. Note that the first entry in the output file is always the full DOS, so there is a use case for creating this block with a number of projections of zero. Ignored if ProjectedDensityOfStates is not specified.

PDOSStart <integer>

Default: 1

If PDOSAtoms is not specified, a single PDOS can be computed using the PDOSStart, PDOSEnd and PDOSMulti keywords. This is equivalent to a PDOSAtoms block of '1 <PDOSMulti>range <PDOSStart><PDOSEnd>'. Ignored if PDOSAtoms is specified, or ProjectedDensityOfStates is not specified.

PDOSEnd <integer>

Default: Last atom of the system

See PDOSStart description. Ignored if PDOSAtoms is specified, or ProjectedDensityOfStates is not specified.

PDOSMulti <float>

Default: 1.0

See PDOSStart description. Ignored if PDOSAtoms is specified, or ProjectedDensityOfStates is not specified.

DOSMulti <float>

Default: 1.0

Since the first entry in the pdos file is the full DOS of the system, even if it does not appear in the PDOSAtoms block, this option exists to allow its scaling. Ignored if ProjectedDensityOfStates is not specified.

SaveGrid <block>

Default: 50 50 50

Specifies the number of points on the real space grid on which to save real space properties, such as potential and charge. Ignored if no other options are specified which would result in the calculation of a real space property. Note that real space properties are saved in .cube file format, and few programs can correctly open such a file if it contains a non-orthogonal unit cell. Furthermore, saving the potential requires invoking the Poisson solver, even in Poisson self-consistency mode was not selected, which will not function correctly for a non-orthogonal unit cell.

SaveRho <bool>

Default: False

Save the charge density of the system, computed on a real space grid, to the output file '`<systemlabel>.RHO.cube`'. This is computed using the Gaussian charge projections of the Poisson solver, rather than Slater orbitals.

SavePotential <bool>

Default: False

Save the potential of the system, computed on a real space grid, to the output file '`<systemlabel>.POT.cube`'. An additional output file '`<systemlabel>.POT.txt`' will also be created, containing the averaged potential across each plane in an increasing z direction. This option is not supported for non-orthogonal unit cells.

SaveDRho <block>

Default: Do not save

Save the difference in charge between this system and another system, which may be different, and may contain different numbers of atoms. This must be provided as a block. The first value is the path to a text file from which the charges on the system to compare with will be read. Following this may be any number of groups of four integer values, specifying the initial atom from the system to compare, the initial atom from the current system, the number of atoms, and the number of times to repeat. For example, if the old system consisted of a single water molecule in isolation, and the new system consisted of a 12 atom gold cluster surrounded by 4 water molecules with the water molecules following the gold atoms in the input file, and you wish to extract how the charge has changed on the water molecules, the input would be '0 12 3 4'.

A.3 Performing a NEGF calculation

A.3.1 Setup, input and output

When performing a transport job, EHTransport uses the .EHT output files created from a minimum of two EHT runs in order to calculate a transmission spectra. This is saved to the output file '<systemlabel>.avtrans'. Note that there is a difference in format of the .avtrans output file between EHTransport and SIESTA. In SIESTA the file contains four columns in the order energy, transmission, DOS, PDOS of molecule without leads. In EHTransport it contains three columns and is ordered energy, DOS, transmission; i.e. the DOS and transmission columns are transposed. The input file expected for the scattering region is always '<systemlabel>.EHT'.

A.3.2 Valid keywords

The following are the list of keywords applicable to NEGF calculations.

TS.HSFileLeft <string>

Default: None

Specifies the file path to the .EHT file to use for the left hand lead. This is a required option.

TS.HSFileRight <string>

Default: Duplicate left lead.

Specifies the file path to the .EHT file to use for the right hand lead. If not set, the same file will be used for both the left and right hand leads.

TS.TBT.Emin <energy>

Default: -1.0 eV

Specifies the minimum energy at which to calculate the transmission and DOS.

TS.TBT.Emax <energy>

Default: +1.0 eV

Specifies the maximum energy at which to calculate the transmission and DOS.

TS.TBT.NPoints <integer>

Default: 100

Specifies the number of data points at which to evaluate the transmission and DOS.

TS.Voltage <float>

Default: 0

Specifies the bias applied across the system, in volts.

CutContactCoupling <integer>

Default: See description

Specifies which element of the scattering regions overlap matrix and Hamiltonian to explicitly set to zero before beginning the NEGF calculation. By default, any orbitals which overlap with the leads are zeroed, but since more than one periodic layer of the leads are normally included in the scattering region, this value can often be safely increased. Note that this option works by the number of orbitals, rather than the number of atoms.

StoredLeadPath <string>

Default: Do not use stored lead self energy matrices

Specifies the path to a folder containing saved self-energy matrices. For each energy at which the transmission is to be evaluated, at the point at which EHTransport would normally calculate the lead self energy, it will check in this folder for a file which contains the self energy matrix and if found will load from this instead rather than recalculating. If this path does not end in a directory separator character, the final part of the path is taken as a prefix with which to label the saved files.

CalcBondCurrents <bool>

Default: False

Save the bond currents through the system. Two output files are generated, 'bondcurrents.txt' which contains an $n \times n$ matrix where n is the number of atoms in the system, listing the bond

currents between every pair of atoms, and 'obondcurrents.txt' which is the same, except in terms of orbitals rather than atoms.

BondCurrentEnergy <energy>

Default: 0 eV

Specifies the energy offset from the fermi level at which to calculate the bond currents. Ignored if CalcBondCurrents is not specified.

A.4 Other available job types

A.4.1 UnitTest

If this job type is selected, EHTransport will run some internal calculations and check the results against what is expected. Tests include comparing a range of analytic Slater orbital overlaps with volume integrals, comparing hydrogen atom EHT calculations with analytic results and testing NEGF output on a model Hamiltonian with analytic results. Some files will also be generated in the working directory, containing more information about the results of some of the tests. There are no further keywords associated with this job type.

A.4.2 SiestaTransport

This job is used to perform NEGF calculations on Transiesta output files. Behaviour is similar to the Transport job type, except that the input file for the scattering region is expected to be '<systemlabel>.TSHS', and the lead input files are also expected to be in the same Transiesta TSHS format. All keywords from Transport job type are valid and have the same meanings here, with the exception of 'StoredLedPath', which is not available. In addition, the 'NumKPoints', 'KPoints' and 'KGridMonkhorstPack' keywords from the Huckel job type are valid and have their same meanings. Note that the z coordinate of all k-points should be 0.

A.4.3 PrecalcLeads

This job type is used to compute the self energy matrices of leads at a range of energies. Where the number of orbitals in a principal layer of the lead is significant compared to the number of orbitals

in the entire scattering region, this can make a significant difference to the speed of the calculation. Of course, this is only a win when the same lead is being used for multiple different scattering regions to calculate transmission on the same energy grid, since otherwise the time taken to compute and store the lead self energy matrices would more than erase this gain. The files generated are named '`<label: %s>.<kpointid: %d>.<energy: %.6f>.<voltage: %.2f>.LEAD`'. The following keywords are recognised, and have the same function as in the Transport job type: '`TS.HSFileLeft`', '`TS.HSFileRight`', '`TS.TBT.EMin`', '`TS.TBT.EMax`', '`TS.TBT.NPoints`' and '`TS.Voltage`'. The '`StoredLeadPath`' keyword is also valid, but now specifies the location to save to rather than read from.

A.5 Code layout

A.5.1 File structure

AnalyticEHT.h, AnalyticEHT.cpp

Contains code for analytically computing overlaps between Slater orbitals

DebugUtils.cpp

Contains code for printing matrices to files in text form, and analysing matrices for properties such as symmetry

EHTData.h, EHTData.cpp

Contains the parameter sets required for performing EHT calculations.

EHTransport.cpp

Contains the startup and program initialisation code

FDF.h, FDF.cpp

Contains code for parsing fdf input files

GausPot.h, GausPot.cpp

Contains code for constructing model Hamiltonian's for a 1d potential barrier using a Gaussian function basis set. Also computes analytic results.

initkgrid.cpp

Contains code for generating a set of kpoints from a given Monkhorst-Pack mesh.

Input.h, Input.cpp

Contains code for constructing a system from the fdf input file and launching the appropriate other modules.

NEGF.h, HEGF.cpp

Contains code for running NEGF calculations based on a real and symmetric input Hamiltonian

OverlapCache.h, OverlapCache.cpp

Implements a caching mechanism to avoid having to recalculate overlaps between Slater orbitals unnecessarily.

PeriodicEHT.h, PeriodicEHT.cpp

Contains code for performing an EHT calculation, including band structure and PDOS analysis.

periodicNEGF.h, periodicNEGF.cpp

Contains code for running NEGF calculations based on a complex hermitian input Hamiltonian

Poisson.h, Poisson.cpp

Contains code for computing the potential at each atom site given a set of charges on each atom, and for making real space plots of charge distribution and potential.

stdafx.h

Contains shared definitions required by all other files

structs.h, structs.cpp

Contains utility structures, including those that represent unit cells, atoms and orbitals, and performs simple operations on those structures.

TSHS.h, TSHS.cpp

Contains code for reading a Transiesta output file in TSHS format and converting it into EHTransports internal representation.

A.5.2 Adding additional atomic species

In an EHT calculation, the input fdf file provides a mechanism for selecting between alternate parameter sets for a given atomic species, but does not offer a way of including completely new parameters. To add support for a new atomic species, two files in EHTransport need to be modified; structs.h and EHTData.cpp. The positions at which new code needs to be inserted should be obvious on inspection, and many of the required parameters such as atomic number and valence electron count are readily available. What is perhaps more complex is the orbital parameters required. Standard non-self consistent EHT parameters are tabulated in several locations, but the additional parameters required for self consistent EHT are not. The values inbuilt into EHTransport are computed from all electron DFT calculations using the program ATOM with a PBE exchange correlation functional.

Bibliography

- [1] J.D. Plummer and P.B. Griffin. Material and process limits in silicon VLSI technology. *Proceedings of the IEEE*, 89(3):240–258, mar 2001.
- [2] M. A. Reed, C. Zhou, C. J. Muller, T. P. Burgin, and J. M. Tour. Conductance of a molecular junction. *Science*, 278:252–254, 1997.
- [3] Ali Yazdani, D. M. Eigler, and N. D. Lang. Off-resonance conduction through atomic wires. *Science*, 272(5270):1921–1924, 1996.
- [4] Sander J. Tans, Alwin R. M. Verschueren, and Cees Dekker. Room-temperature transistor based on a single carbon nanotube. *Nature*, 393:49–52, May 1998.
- [5] V M Garcá-Suárez and C J Lambert. Non-trivial length dependence of the conductance and negative differential resistance in atomic molecular wires. *Nanotechnology*, 19(45):455203, 2008.
- [6] Jeremy M. Beebe, Vincent B. Engelkes, Larry L. Miller, and C. Daniel Frisbie. Contact resistance in metal-molecule-metal junctions based on aliphatic SAMs: Effects of surface linker and metal work function. *Journal of the American Chemical Society*, 124(38):11268–11269, 2002. PMID: 12236731.
- [7] Cai-Juan Xia, De-Sheng Liu, Han-Chen Liu, and Xue-Jun Zhai. Large negative differential resistance in a molecular device with asymmetric contact geometries: A first-principles study. *Physica E: Low-dimensional Systems and Nanostructures*, 43(8):1518 – 1521, 2011.
- [8] Adrian Bachtold, Peter Hadley, Takeshi Nakanishi, and Cees Dekker. Logic circuits with carbon nanotube transistors. *Science*, 294(5545):1317–1320, 2001.
- [9] K. Hansen, E. Laegsgaard, I. Stensgaard, and F. Besenbacher. Quantized conductance in relays. *Physical Review B*, 56(4):2208–2220, 1997.
- [10] John G. Simmons. Generalized formula for the electric tunnel effect between similar electrodes separated by a thin insulating film. *Journal of Applied Physics*, 34(6):1793–1803, 1963.
- [11] John G. Simmons. Low-voltage current-voltage relationship of tunnel junctions. *Journal of Applied Physics*, 34(1):238–239, 1963.
- [12] R. Landauer. Spatial variation of currents and fields due to localized scatterers in metallic conduction. *IBM Journal of Research and Development*, 1(3):223–231, july 1957.
- [13] M. Büttiker. Absence of backscattering in the quantum Hall effect in multiprobe conductors. *Phys. Rev. B*, 38:9375–9389, Nov 1988.

- [14] Supriyo Datta. Electrical resistance: an atomistic view. *NANOTECHNOLOGY*, 15:S433–S451, 2004.
- [15] Yongqiang Xue, Supriyo Datta, and Mark A. Ratner. First-principles based matrix Green’s function approach to molecular electronic devices: General formalism. *Chemical Physics*, 281:151–170, 2002.
- [16] P. S. Damle, A. W. Ghosh, and S. Datta. Unified description of molecular conduction: From molecules to metallic wires. *Physical Review B*, 64:201403, 2001.
- [17] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136:B864–B871, Nov 1964.
- [18] J. C. Slater. A simplification of the hartree-fock method. *Phys. Rev.*, 81:385–390, Feb 1951.
- [19] James B. Robinson and Peter J. Knowles. Breaking multiple covalent bonds with Hartree-Fock-based quantum chemistry: Quasi-Variational Coupled Cluster theory with perturbative treatment of triple excitations. *Phys. Chem. Chem. Phys.*, 14:6729–6732, 2012.
- [20] J R Reimers, T X Lü, M J Crossley, and N S Hush. Molecular electronic properties of fused rigid porphyrin-oligomer molecular wires. *Nanotechnology*, 7(4):424–429, 1996.
- [21] Aleksey A. Kocherzhenko, Sameer Patwardhan, Ferdinand C. Grozema, Harry L. Anderson, and Laurens D. A. Siebbeles. Mechanism of charge transport along zinc porphyrin-based molecular wires. *Journal of the American Chemical Society*, 131(15):5522–5529, 2009. PMID: 19331354.
- [22] H Wende, M Bernien, J Luo, C Sorg, N Ponpandian, J Jurde, J Miguel, M Piantek, X Xu, Ph Eckhold, W Kuch, K Baberschke, P M Panchmatia, B Sanyal, P M Oppeneer, and O Eriksson. Substrate-induced magnetic ordering and switching of iron porphyrin molecules. *Nat Mater*, 6(7):516–520, 2007.
- [23] Neil Bennett, Gengzhao Xu, Louisa J. Esdaile, Harry L. Anderson, J. Emyr Macdonald, and Martin Elliott. Transition voltage spectroscopy of porphyrin molecular wires. *Small*, 6(22):2604–2611, 2010.
- [24] Yanwei Li, Jinhuan Yao, Shengkui Zhong, and Zhengguang Zou. Theoretical investigations on the orientational dependence of electron transport through porphyrin molecular wire. *Current Applied Physics*, 11(6):1349 – 1353, 2011.
- [25] Sairam S. Mallajosyula and Swapan K. Pati. Toward DNA conductivity: A theoretical perspective. *The Journal of Physical Chemistry Letters*, 1(12):1881–1894, 2010.
- [26] Roald Hoffmann. An extended Hückel theory. i. hydrocarbons. *The Journal of Chemical Physics*, 39(6):1397–1412, 1963.
- [27] C. Garrido, A.E. Aliaga, J.S. Gómez-Jeria, J.J. Cárcamo, E. Clavijo, and M.M. Campos-Vallette. Interaction of the C-terminal peptide from pigeon cytochrome C with silver nanoparticles. a Raman, SERS and theoretical study. *Vibrational Spectroscopy*, 61(0):94 – 98, 2012.

- [28] Volodymyr Babizhetskyy, Hansjürgen Mattausch, Arndt Simon, Régis Gautier, and Jean-François Halet. New members of ternary rare-earth metal boride carbides containing finite boron-carbon chains: RE₂₅B₁₄C₂₆ (RE=Pr, Nd) and Nd₂₅B₁₂C₂₈. *Journal of Solid State Chemistry*, 184(7):1671 – 1681, 2011.
- [29] Michaël Magoga, Fabien Archambault, and Jorge I. Cerdá. Nt_STM: A step forward in scanning tunneling microscopy (STM) simulations. *Computer Physics Communications*, 183(6):1246 – 1249, 2012.
- [30] Louis Chopin Cusachs and James Ward Reynolds. Selection of molecular matrix elements from atomic data. *The Journal of Chemical Physics*, 43(10):S160–S164, 1965.
- [31] Robert Rein, Nubuo Fukuda, Htain Win, George A. Clarke, and Frank E. Harris. Iterative extended Hückel theory. *The Journal of Chemical Physics*, 45(12):4743–4744, 1966.
- [32] Supriyo Datta. Electrical resistance: an atomistic view. *Nanotechnology*, 15(7):S433, 2004.
- [33] S. Sanvito, C. J. Lambert, J. H. Jefferson, and A. M. Bratkovsky. General Green’s-function formalism for transport calculations with spd hamiltonians and giant magnetoresistance in Co- and Ni-based magnetic multilayers. *Phys. Rev. B*, 59:11936–11948, May 1999.
- [34] M P Lopez Sancho, J M Lopez Sancho, J M L Sancho, and J Rubio. Highly convergent schemes for the calculation of bulk and surface Green functions. *Journal of Physics F: Metal Physics*, 15(4):851, 1985.
- [35] Wang, J.-S., Wang, J., and Lü, J. T. Quantum thermal transport in nanostructures. *Eur. Phys. J. B*, 62(4):381–404, 2008.
- [36] L. H. Thomas. The calculation of atomic fields. *Mathematical Proceedings of the Cambridge Philosophical Society*, 23(05):542–548, 1927.
- [37] V Fock. Näherungsmethode zur lösung des quantenmechanischen mehrkörperproblems. *Zeitschrift für Physik A Hadrons and Nuclei*, 61(1):126–148, 1930.
- [38] M.P. Nightingale and Cyrus J. Umrigar. *Quantum Monte Carlo methods in physics and chemistry*. Springer, 1999.
- [39] M Ya Amusia, A Z Msezane, and V R Shaginyan. Density functional theory versus the Hartree-Fock method: Comparative assessment. *Physica Scripta*, 68(6):C133, 2003.
- [40] M Born and R. Oppenheimer. Zur quantentheorie der molekeln. *Ann Phys., Lpz.*, 84:457–484, 1927.
- [41] Nikitas Gidopoulos. Kohn-Sham equations for multicomponent systems: The exchange and correlation energy functional. *Phys. Rev. B*, 57:2146–2152, Jan 1998.
- [42] T. Kreibich and E. K. U. Gross. Multicomponent density-functional theory for electrons and nuclei. *Phys. Rev. Lett.*, 86:2984–2987, Apr 2001.
- [43] N. David Mermin. Thermal properties of the inhomogeneous electron gas. *Phys. Rev.*, 137:A1441–A1443, Mar 1965.
- [44] H. Englisch and R. Englisch. Hohenberg-Kohn theorem and non-V-representable densities. *Physica A: Statistical Mechanics and its Applications*, 121(1–2):253 – 268, 1983.

- [45] S Datta. Relativistic extension of the Hohenberg-Kohn theorem. *Pramana*, 28:633–639, 1987. 10.1007/BF02892864.
- [46] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133–A1138, Nov 1965.
- [47] Abdel-Raouf E. Mohammed and V. Sahni. Density-functional-theory gradient expansion approximation for the screened-Coulomb exchange energy. *Phys. Rev. B*, 29:3687–3690, Mar 1984.
- [48] O. Gunnarsson, M. Jonson, and B. I. Lundqvist. Descriptions of exchange and correlation effects in inhomogeneous electron systems. *Phys. Rev. B*, 20:3136–3164, Oct 1979.
- [49] John P. Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Phys. Rev. Lett.*, 77:3865–3868, Oct 1996.
- [50] Axel D. Becke. A new mixing of Hartree-Fock and local density-functional theories. *The Journal of Chemical Physics*, 98(2):1372–1377, 1993.
- [51] R. P. Feynman. Forces in molecules. *Phys. Rev.*, 56:340–343, Aug 1939.
- [52] P. Pulay. Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules. *Molecular Physics*, 17(2):197–204, 1969.
- [53] M. Lüders, M. A. L. Marques, N. N. Lathiotakis, A. Floris, G. Profeta, L. Fast, A. Continenza, S. Massidda, and E. K. U. Gross. *Ab initio* theory of superconductivity. i. density functional formalism and approximate functionals. *Phys. Rev. B*, 72:024545, Jul 2005.
- [54] M Pearson, E Smargiassi, and P A Madden. Ab initio molecular dynamics with an orbital-free density functional. *Journal of Physics: Condensed Matter*, 5(19):3221–3240, 1993.
- [55] Erich Runge and E. K. U. Gross. Density-functional theory for time-dependent systems. *Phys. Rev. Lett.*, 52:997–1000, Mar 1984.
- [56] G. Stefanucci, C.-O. Almbladh, S. Kurth, E.K.U. Gross, A. Rubio, R. van Leeuwen, N.E. Dahlen, and U. von Barth. Time-dependent transport through single molecules: Nonequilibrium Green’s functions. In Miguel Marques, Carsten Ullrich, Fernando Nogueira, Angel Rubio, Kieron Burke, and Eberhard Gross, editors, *Time-Dependent Density Functional Theory*, volume 706 of *Lecture Notes in Physics*, pages 479–492. Springer Berlin / Heidelberg, 2006.
- [57] Erich Hückel. Quantentheoretische beiträge zum benzolproblem. *Zeitschrift für Physik A Hadrons and Nuclei*, 70:204–286, 1931. 10.1007/BF01339530.
- [58] Erich Hückel. Quantentheoretische beiträge zum benzolproblem. *Zeitschrift für Physik A Hadrons and Nuclei*, 72:310–337, 1931. 10.1007/BF01341953.
- [59] Erich Hückel. Quantentheoretische beiträge zum problem der aromatischen und ungesättigten verbindungen. iii. *Zeitschrift für Physik A Hadrons and Nuclei*, 76:628–648, 1932. 10.1007/BF01341936.
- [60] Erich Hückel. Die freien radikale der organischen chemie. *Zeitschrift f/ur Physik A Hadrons and Nuclei*, 83:632–668, 1933. 10.1007/BF01330865.

- [61] R. S. Mulliken. Electronic population analysis on LCAO-MO molecular wave functions. i. *The Journal of Chemical Physics*, 23(10):1833–1840, 1955.
- [62] Alan E. Reed, Robert B. Weinstock, and Frank Weinhold. Natural population analysis. *The Journal of Chemical Physics*, 83(2):735–746, 1985.
- [63] Kurt Stokbro, Dan Erik Petersen, Søren Smidstrup, Anders Blom, Mads Ipsen, and Kristen Kaasbjerg. Semiempirical model for nanoscale device simulations. *Phys. Rev. B*, 82:075420, Aug 2010.
- [64] Mads Brandbyge, José-Luis Mozos, Pablo Ordejón, Jeremy Taylor, and Kurt Stokbro. Density-functional method for nonequilibrium electron transport. *Phys. Rev. B*, 65:165401, Mar 2002.
- [65] Felix Bloch. über die quantenmechanik der elektronen in kristallgittern. *Zeitschrift für Physik A Hadrons and Nuclei*, 52:555–600, 1929. 10.1007/BF01339455.
- [66] Hendrik J. Monkhorst and James D. Pack. Special points for Brillouin-zone integrations. *Phys. Rev. B*, 13:5188–5192, Jun 1976.
- [67] José M Soler, Emilio Artacho, Julian D Gale, Alberto García, Javier Junquera, Pablo Ordejón, and Daniel Sánchez-Portal. The SIESTA method for ab initio order-N materials simulation. *Journal of Physics: Condensed Matter*, 14(11):2745, 2002.
- [68] Emilio Artacho, E Anglada, O Diéguez, J D Gale, A García, J Junquera, R M Martin, P Ordejón, J M Pruneda, D Sánchez-Portal, and J M Soler. The SIESTA method; developments and applicability. *Journal of Physics: Condensed Matter*, 20(6):064208, 2008.
- [69] Daniel Sánchez-Portal, Pablo Ordejón, Emilio Artacho, and José M. Soler. Density-functional method for very large systems with LCAO basis sets. *International Journal of Quantum Chemistry*, 65(5):453–461, 1997.
- [70] Pablo Ordejón, Emilio Artacho, and José M. Soler. Self-consistent order- n density-functional calculations for very large systems. *Phys. Rev. B*, 53:R10441–R10444, Apr 1996.
- [71] Javier Junquera, Óscar Paz, Daniel Sánchez-Portal, and Emilio Artacho. Numerical atomic orbitals for linear-scaling calculations. *Phys. Rev. B*, 64:235111, Nov 2001.
- [72] Eduardo Anglada, José M. Soler, Javier Junquera, and Emilio Artacho. Systematic generation of finite-range atomic basis sets for linear-scaling calculations. *Phys. Rev. B*, 66:205101, Nov 2002.
- [73] Kurt Stokbro, Jeremy Taylor, and Mads Brandbyge Pablo Ordekón. TranSIESTA: A spice for molecular electronics. *Annals of the New York Academy of Sciences*, 1006(1):212–226, 2003.
- [74] J. C. Slater. Atomic shielding constants. *Phys. Rev.*, 36:57–64, Jul 1930.
- [75] James D. Talman. Expression for overlap integrals of Slater orbitals. *Phys. Rev. A*, 48:243–249, Jul 1993.
- [76] D.R Bowler and M.J Gillan. An efficient and robust technique for achieving self consistency in electronic structure calculations. *Chemical Physics Letters*, 325(4):473 – 476, 2000.

- [77] Richard H. Summerville and Roald Hoffmann. Tetrahedral and other M₂L₆ transition metal dimers. *Journal of the American Chemical Society*, 98(23):7240–7254, 1976.
- [78] Sanshiro Komiya, Thomas A. Albright, Roald Hoffmann, and Jay K. Kochi. Reductive elimination and isomerization of organogold complexes. theoretical studies of trialkylgold species as reactive intermediates. *Journal of the American Chemical Society*, 98(23):7255–7265, 1976.
- [79] P. Pyykkö and Lawrence L. Lohr. Relativistically parameterized extended Hückel calculations. 3. structure and bonding for some compounds of uranium and other heavy elements. *Inorganic Chemistry*, 20(7):1950–1959, 1981.
- [80] Noel J. Fitzpatrick and George H. Murphy. Double zeta d radial wave functions for transition elements. *Inorganica Chimica Acta*, 111(2):139 – 140, 1986.
- [81] Pekka Pyykkö and Leif Laaksonen. Relativistically parameterized extended hückel calculations. 8. double-zeta. parameters for the actinoids thorium, protactinium, uranium, neptunium, plutonium, and americium and an application on uranyl. *The Journal of Physical Chemistry*, 88(21):4892–4895, 1984.
- [82] J. Cerdá and F. Soria. Accurate and transferable extended Hückel-type tight-binding parameters. *Phys. Rev. B*, 61:7965–7971, Mar 2000.
- [83] Pablo Ordejón, David A. Drabold, Richard M. Martin, and Matthew P. Grumbach. Linear system-size scaling methods for electronic-structure calculations. *Phys. Rev. B*, 51:1456–1476, Jan 1995.
- [84] Abhik Ghosh, Jan Almlöf, and Paul G. Gassman. Ab initio SCF studies of basis set effects in free base porphyrin. *Chemical Physics Letters*, 186(1):113 – 118, 1991.
- [85] Matthias Ernzerhof and Gustavo E. Scuseria. Assessment of the Perdew–Burke–Ernzerhof exchange-correlation functional. *The Journal of Chemical Physics*, 110(11):5029–5036, 1999.
- [86] Gang Zhang and Charles B. Musgrave. Comparison of DFT methods for molecular orbital eigenvalue calculations. *The Journal of Physical Chemistry A*, 111(8):1554–1561, 2007.
- [87] David M. Lemal. Perspective on fluorocarbon chemistry. *The Journal of Organic Chemistry*, 69(1):1–11, 2004. PMID: 14703372.
- [88] M.F. Perutz. Hemoglobin structure and respiratory transport. *Scientific American*, 239, 1978.
- [89] Gita Sedghi, Victor M. Garcia-Suarez, Louisa J. Esdaile, Harry L. Anderson, Colin J. Lambert, Santiago Martin, Donald Bethell, Simon J. Higgins, Martin Elliott, Neil Bennett, Emyr J. Macdonald, and Richard J. Nichols. Long-range electron tunnelling in oligoporphyrin molecular wires. *Nat Nano*, 6(8):517 to 523, 2011.
- [90] Guan M Wang, William C Sandberg, and Steven D Kenny. Density functional study of a typical thiol tethered on a gold surface: ruptures under normal or parallel stretch. *Nanotechnology*, 17(19):4819, 2006.
- [91] E.Yu. Zarechnaya, N.V. Skorodumova, S.I. Simak, B. Johansson, and E.I. Isaev. Theoretical study of linear monoatomic nanowires, dimer and bulk of Cu, Ag, Au, Ni, Pd and Pt. *Computational Materials Science*, 43(3):522 – 530, 2008.

- [92] Katta G. Murty. *Linear programming*. John Wiley & Sons Inc., New York, 1983. With a foreword by George B. Dantzig.
- [93] A. K. Rappe, C. J. Casewit, K. S. Colwell, W. A. Goddard, and W. M. Skiff. UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. *Journal of the American Chemical Society*, 114(25):10024–10035, 1992.
- [94] Thomas A. Halgren. Merck molecular force field. i. basis, form, scope, parameterization, and performance of MMFF94. *Journal of Computational Chemistry*, 17(5-6):490–519, 1996.
- [95] Hannu Hakkinen. The gold-sulfur interface at the nanoscale. *Nat Chem*, 4(6):443–455, 2012.
- [96] L A Majewski, R Schroeder, M Voigt, and M Grell. High performance organic transistors on cheap, commercial substrates. *Journal of Physics D: Applied Physics*, 37(24):3367, 2004.
- [97] Gita Sedghi, Louisa J. Esdaile, Harry L. Anderson, Santiago Martin, Donald Bethell, Simon J. Higgins, and Richard J. Nichols. Comparison of the conductance of three types of porphyrin-based molecular wires: β ,meso, β -fused tapes, meso-butadiyne-linked and twisted meso-meso linked oligomers. *Advanced Materials*, 24(5):653–657, 2012.
- [98] Katsunori Tagami, Masaru Tsukada, Takuya Matsumoto, and Tomoji Kawai. Electronic transport properties of free-base tape-porphyrin molecular wires studied by self-consistent tight-binding calculations. *Phys. Rev. B*, 67:245324, Jun 2003.
- [99] E. Leary, H. Höbenreich, S. J. Higgins, H. van Zalinge, W. Haiss, R. J. Nichols, C. M. Finch, I. Grace, C. J. Lambert, R. McGrath, and J. Smerdon. Single-molecule solvation-shell sensing. *Phys. Rev. Lett.*, 102:086801, Feb 2009.
- [100] C Lambert.
- [101] V. Fatemi, M. Kamenetska, J. B. Neaton, and L. Venkataraman. Environmental control of single-molecule junction transport. *Nano Letters*, 11(5):1988–1992, 2011.
- [102] Fabio Arnesano, Lucia Banci, Ivano Bertini, Jasmin Faraone-Mennella, Antonio Rosato, Paul D. Barker, and Alan R. Fersht. The solution structure of oxidized escherichia coli cytochrome b562. *Biochemistry*, 38(27):8657–8670, 1999.
- [103] Eduardo Antonio Della Pia, Qijin Chi, D. Dafydd Jones, J. Emyr Macdonald, Jens Ulstrup, and Martin Elliott. Single-molecule mapping of long-range electron transport for a cytochrome b562 variant. *Nano Letters*, 11(1):176–182, 2011.
- [104] G Sedghi.
- [105] A.G.W. Leslie, Struther Arnott, Rengaswami Chandrasekaran, and R.L. Ratliff. Polymorphism of DNA double helices. *Journal of Molecular Biology*, 143(1):49 – 72, 1980.
- [106] Stefan Rothenburg, Friedrich Koch-Nolte, and Friedrich Haag. DNA methylation and Z-DNA formation as mediators of quantitative differences in the expression of alleles. *Immunological Reviews*, 184(1):286–298, 2001.
- [107] Hans-Werner Fink and Christian Schonenberger. Electrical conduction through DNA molecules. *Nature*, 398(6726):407–410, 1998.

- [108] Erez Braun, Yoav Eichen, and Uri Sivan. DNA-templated assembly and electrode attachment of a conducting silver wire. *Nature*, 391(6669):775–778, 1997.
- [109] Danny Porath, Alexey Bezryadin, and Simon de Vries. Direct measurement of electrical transport through DNA molecules. *Nature*, 403(6770):635–638, 1999.
- [110] A. Rakitin, P. Aich, C. Papadopoulos, Yu. Kobzar, A. S. Vedenev, J. S. Lee, and J. M. Xu. Metallic conduction through engineered DNA: DNA nanoelectronic building blocks. *Phys. Rev. Lett.*, 86:3670–3673, Apr 2001.
- [111] P. Tran, B. Alavi, and G. Gruner. Charge transport along the λ -DNA double helix. *Phys. Rev. Lett.*, 85:1564–1567, Aug 2000.
- [112] Guohui Zheng, Xiang-Jun Lu, and Wilma K. Olson. Web 3DNA—a web server for the analysis, reconstruction, and visualization of three-dimensional nucleic-acid structures. *Nucleic Acids Research*, 37(suppl 2):W240–W246, 2009.
- [113] Christopher R. Treadway, Michael G. Hill, and Jacqueline K. Barton. Charge transport through a molecular π -stack: double helical DNA. *Chemical Physics*, 281(2–3):409 – 428, 2002.
- [114] K.-H. Yoo, D. H. Ha, J.-O. Lee, J. W. Park, Jinhee Kim, J. J. Kim, H.-Y. Lee, T. Kawai, and Han Yong Choi. Electrical conduction through poly(dA)-poly(dT) and poly(dG)-poly(dC) DNA molecules. *Phys. Rev. Lett.*, 87:198102, Oct 2001.
- [115] Hanfei Wang. *Sequence and Length Dependence of the Conductivity of Individual DNA Duplexes and Applications in Protein Detection*. PhD thesis, Columbia University, 2012.
- [116] Xu, Zhang, Li, and Tao. Direct conductance measurement of single DNA molecules in aqueous solution. *Nano Letters*, 4(6):1105–1108, 2004.
- [117] Kittusamy Senthilkumar, Ferdinand C. Grozema, Célia Fonseca Guerra, F. Matthias Bickelhaupt, Frederick D. Lewis, Yuri A. Berlin, Mark A. Ratner, and Laurens D. A. Siebbeles. Absolute rates of hole transfer in DNA. *Journal of the American Chemical Society*, 127(42):14894–14903, 2005.
- [118] Jason D. Slinker, Natalie B. Muren, Sara E. Renfrew, and Jacqueline K. Barton. DNA charge transport over 34 nm. *Nat Chem*, 3(3):228–233, 2011.
- [119] Takayuki Tsukamoto, Yasuyuki Ishikawa, Yasuo Sengoku, and Noriyuki Kurita. A combined DFT/Green’s function study on electrical conductivity through DNA duplex between Au electrodes. *Chemical Physics Letters*, 474(4–6):362 – 365, 2009.
- [120] Xinqian Li, Aleksandar Staykov, and Kazunari Yoshizawa. Orbital views on electron-transport properties of cyclophanes: Insight into intermolecular transport. *Bulletin of the Chemical Society of Japan*, 85(2):181–188, 2012.
- [121] Mark Harris. Mapping computational concepts to GPUs. In *ACM SIGGRAPH 2005 Courses*, SIGGRAPH ’05, New York, NY, USA, 2005. ACM.