Low-Dimensional Carbon Nanotube and Graphene Devices
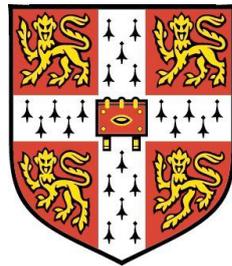
Philip Hedley Scard

Gonville & Caius College

University of Cambridge

A thesis submitted for the degree of

Doctor of Philosophy

March 2010

Declaration

This dissertation is the result of work carried out in the Semiconductor Physics group at the Cavendish Laboratory, from October 2004 to March 2010.

It is my own work and contains nothing which is the outcome of work done in collaboration, except where specified in the text and Acknowledgements. This dissertation has not been submitted, in whole or in part, for any degree or qualification at this or any other university. It does not exceed 60,000 words.

Philip Hedley Scard

Gonville and Caius College, 2010

Abstract

Electronic devices in which the electrons are confined to fewer than three spatial dimensions are an important tool for physics research and future developments in computing technology. Recently discovered carbon nanotubes (1991) and graphene (2004) are intrinsically low-dimensional materials with remarkable electronic properties. Combined with semiconductor technologies they might be used to fabricate smaller devices with more complex functionality. This thesis addresses two routes towards this goal.

The detection of charge transport through quantum dots using a GaAs point contact is a potential tool for quantum computation. This project aimed to fabricate and measure hybrid devices with carbon nanotube quantum dots on top of GaAs point contacts. Dispersion and AFM manipulations of nanotubes on GaAs were studied, revealing comparatively weak binding. Transport measurements indicated that GaAs induces disorder in nanotubes, creating multiple tunnel barriers. Preliminary attempts were made at CVD growth and ink-jet printing of nanotubes directly onto GaAs.

Although only one atom thick, graphene is macroscopic in area and must be patterned to confine conduction; room temperature transistor behaviour requires graphene ribbons only a few nanometres wide. This work fabricated such structures using a charged AFM tip, achieving reliable cutting even on single layer graphene and feature sizes as small as 5 nm. The cutting mechanism was found to be chemical oxidation of carbon by a polarised water layer, with an activation energy determined by the energy of dissociation of water at the graphene surface. The critical variables were the voltage difference between the tip and graphene and the atmospheric humidity. An unstable solid oxide intermediate was also observed. Thermal annealing revealed the presence of a layer of water beneath flakes.

Finally, EFM measurements were made of graphene at 20 mK, enabling estimates of the local carrier density and revealing spatial variations in the electronic structure on a scale consistent with electron and hole puddles.

Contents

List of Figures

LIST OF FIGURES

Chapter 1

Introduction

Advances in microelectronics research are fundamental for many technological developments in the present and near future. The need to reduce the size of components to achieve improved performance means that future electronic devices will involve countable numbers of atoms. On such small scales quantum physics begins to determine device function, and electron motion may be constrained to fewer than three spatial dimensions. Research in such low-dimensional devices is already well established in the field of semiconductors and provides insights into fundamental physics, but new challenges often require new tools.

Until 1985 crystalline carbon was known to exist either as diamond or graphite, both macroscale materials without microscopic electronics device applications. The discovery of buckminsterfullerene [1] initiated the identification of a whole new class of carbon structures with nanoscale dimensions, of which carbon nanotubes and graphene have particularly attractive and promising properties.

Graphene is a single layer of graphite, a hexagonal lattice of carbon only one atom thick and an intrinsically two-dimensional material. Rolled into a cylinder it becomes a carbon nanotube. With diameters of the order of a nanometre these have one-dimensional transport properties, and they can be further divided into zero-dimensional quantum dots.

Carbon nanotubes are an exciting new element to combine with and extend existing microelectronics devices and processes. Their intrinsic transport is already well studied, but industrial use is hindered by the incompatibilities between nanotube synthesis processes and device fabrication techniques. Meanwhile there is much scope for furthering understanding of their interaction with other materials, and expanding the capabilities of nanotube-based devices. This work brings nanotubes together with gal-

1. INTRODUCTION

lium arsenide technology to try to create a new class of hybrid devices, incorporating the low dimensional transport properties of both materials. It examines the physical and electronic interactions between carbon nanotubes and gallium arsenide, and explores the development of a device fabrication process that is compatible with both materials.

Graphene was first isolated after this work was begun, but the coincidence of its properties with the facilities and experimental goals of the Semiconductor Physics Group made it an ideal and exciting material to investigate. Measurements were rapidly being reported with bulk flakes, but incorporating graphene's remarkable properties into functional electronics requires the ability to pattern and manipulate graphene at the nanoscale. AFM (atomic force microscope) lithography was chosen as a novel and practical solution, due to its potentially high resolution and precise control.

The experiments in this thesis analyse the conditions needed for high quality lithography and the mechanism of the cutting process, to define a reliable and repeatable method for patterning single layer graphene. The physical and chemical reactions between graphene and its environment are less well understood than for nanotubes, and the scope of this work encompasses investigations into the role of water on and around graphene, and the oxidation of graphene.

Scanning probe microscopy techniques, of which AFM is one example, are powerful and flexible tools for measuring and probing on the molecular scale. Where AFM gives information about the topography of a sample, EFM (electrostatic force microscopy) measures its electronic properties. It is used here to look for local variations in the electron density of graphene, such as electron and hole puddles.

The projects described connect the substantial and fast moving fields of semiconductors, carbon and scanning probes, as well as many associated topics. The targets of the research retreated over time from measuring complex devices to simply fabricating the devices themselves, and so the focus of the work, and the nature of the problems encountered changed accordingly as it progressed. What was begun as a work on quantum transport became increasingly about the properties of surfaces and their interactions.

The constantly shifting background made it difficult to identify pre-emptively the relevant literature, and consequently much of the analysis for the problems encountered was derived retrospectively from published work. The resulting large number of referenced works provide a broad and connected review of fields of research that will increasingly interact as work on carbon-based electronics progresses.

1.1   Thesis outline

Chapter 2 outlines relevant background physics. The theory of low-dimensional electron systems is given, followed by an explanation of such systems in gallium arsenide heterostructures and graphitic carbon. The structure and properties of nanotubes and graphene are described in detail, along with an explanation of the applications of this work, including quantum computation.

Chapter 3 presents results and observations from a project to integrate carbon nanotubes with gallium arsenide devices. The key stages are developing a reliable dispersing technique, modifications to fabrication processes to incorporate the nanotubes and results from electrical measurements. Background is given to semiconductor device processing.

Chapter 4 contains results from a successful technique developed to cut and pattern graphene using an AFM. Procedures for making graphene devices are described, followed by experiments to establish the best parameters for reliable cutting. A mechanism for the process is deduced, and evidence is given for a water layer trapped beneath graphene. An explanation is provided of the operation of scanning probe microscopes.

Chapter 5 describes preliminary work and results from electrostatic scanning probe measurements of graphene in a low temperature fridge, which are used to make local estimates of the carrier density. Low temperature transport measurements of graphene are included.

Chapter 6 discusses general conclusions that can be drawn from this work, comparing the benefits and potential of carbon nanotube devices on gallium arsenide and AFM lithography of graphene.

1. INTRODUCTION

Chapter 2

Low-Dimensional Electron Systems

2.1 Introduction

As electron systems approach smaller sizes confinement and a reduction in scattering permit quantum mechanical effects to dominate transport behaviour, and cause a reduction in the dimensionality of the system. Semiconductor materials such as gallium arsenide (GaAs) are a well established platform for exploring such low-dimensional electron systems at low temperatures, because the crystal potential can be modified locally in order to manipulate charge carriers. More recently, newly discovered fullerenes and related structures have emerged as alternative materials for such research. These all-carbon structures are intrinsically low-dimensional electron systems at all temperatures, for example graphene (2D) and carbon nanotubes (1D). This chapter outlines the physics of low-dimensional electron transport, GaAs devices, the properties of graphene and carbon nanotubes, and their electronic applications.

2.2 Transport in a crystal

Electron behaviour in a crystal can be approximated by the tight-binding model, in which electronic orbital states of individual atoms combine to form bands of states available to electrons throughout the system [2]. At zero temperature the electrons are confined within bands, such that if a crystal contains only full or empty bands then the electrons cannot contribute to charge flow through the bulk crystal - the material is insulating. Where bands are partly occupied, electrons (conversely modelled as electron vacancies, or holes, in nearly full bands) can move from site to site, thus carrying charge, and so the crystal is metallic.

2. LOW-DIMENSIONAL ELECTRON SYSTEMS

In an intrinsic semiconductor the energy gap (band gap) between the highest occupied band (the valence band) and the lowest unoccupied band (the conduction band) is small enough that at finite temperatures thermal excitations can be of sufficient energy to carry electrons from the valence to conduction band. These electrons behave as if in free space (an electron gas), but with an effective mass different from the free-space mass because of interactions with the underlying crystal lattice. Doping the semiconductor crystal with impurity atoms which can donate additional carriers (electrons or holes) permits tuning of the conduction properties of the material, and enhancement of the carrier density. Semiconductors can thus be made to behave like metals at low temperatures by creating a three-dimensional electron or hole gas.

If the impurity atom has an extra bonding electron in its shell, compared with the host lattice, this is donated to the conduction band, creating an electron gas; it is an n-type dopant. An impurity atom with one fewer outer shell electron removes an electron from the semiconductor valence band, creating a hole gas; it is a p-type dopant. Such three-dimensional gases are not ideal for studying quantum effects because the background of ionised impurities introduces strong disorder - they are better investigated in lower dimensional systems.

The Fermi energy E_f is the energy of the highest occupied electron state at zero temperature, at which the Fermi-Dirac distribution

$$f(E, T) = \frac{1}{e^{(E-E_f)/k_B T} + 1}$$

switches from 1 to 0. The Fermi wavelength is

$$\lambda_f = \frac{h}{\sqrt{2m^* E_f}}$$

where m^* is effective mass of the charge carrier.

The dimensionality of a system of charge carriers is defined as the number of spatial directions in which the carriers are free to move, or with respect to quantum mechanics, the number of dimensions in which extended states are free to evolve. When a confining potential restricts the motion of carriers in a particular direction, the particle wavefunction is reflected off the potential walls and standing waves are formed.

Charge transport in solids is largely described by the properties of carriers at the Fermi energy. Therefore reduced dimensionality is achieved if the energy of the charge carriers is such that only one (usually the lowest) energy state is available, that is, when a dimension of the electron system is comparable to the Fermi wavelength of the electron gas and the gap between energy levels is such that $\Delta E \gg k_b T$. The

transport properties of a system depend on its dimensionality d, since the spectral electron density $n_d(E) = D_d(E)f(E, T)$ depends on the density of states $D_d(E)$, where $D_3(E) \propto \sqrt{E}$, $D_2(E)$ is independent of E, $D_1(E) \propto 1/\sqrt{E}$, and $D_0(E)$ is a series of δ functions.

2.2.1   2D electron systems

Adding a confining potential in one dimension defines a two-dimensional electron gas (2DEG) within the crystal. Motion of electrons through the crystal is strongly influenced by scattering events: elastic scattering by ionised impurities, and inelastic scattering by lattice vibrations. The carrier mobility of the system is defined as

$$\mu = \frac{e\tau}{m^*}$$

where τ is the time between scattering events. On macroscopic scales electron motion is modelled as diffusive. Ballistic motion occurs on length scales equal to or smaller than the elastic scattering events. Scattering can be reduced by removing impurities from regions of electron motion and cooling the lattice to reduce vibrations. The Fermi velocity is $v_f = \hbar k_f/m^*$ and n_{2D} is the 2D carrier density, therefore the electron mean free path between collisions (the elastic scattering length) is

$$l_e = v_f\tau_e = \mu\hbar\sqrt{2\pi n_{2D}}/e$$

The electron mobility is a measure of the response of free charge carriers to an electric field, and along with the carrier concentration is a useful figure of merit for a material [3]. These are derived from measurements of the Hall effect. When a magnetic field is applied perpendicular to the direction of current flow the path of the carriers is curved by the Lorentz force. Carriers accumulate at one edge of the 2D region, establishing an electric field perpendicular to both the current and magnetic field, called the Hall Voltage. The device geometry used to measure this is a channel of width W, length L and with source and drain contacts at either end and side arms with probes to measure voltage both along and across the channel. A Hall bar device is shown in Figure A.1. The Hall voltage is given by

$$V_H = \frac{IB}{qn}$$

therefore the carrier concentration n can be found from the low field magnetoresistance, with $n = 1/R_Hq$ where R_H is the Hall coefficient, equal to V_H/IB. Mobility is

additionally given by $\mu = 1/nq\rho$ where ρ is the resistivity. When considering a 2DEG the channel depth can be ignored, so the diagonal resistivity is

$$\rho_{xx} = \frac{E_x}{j_x} = \frac{V_x W}{IL} = R_{2DEG} \frac{W}{L}$$

The voltage is measured along the channel (V_x) at zero field, to give the 2DEG resistance R_{2DEG} and thus the mobility:

$$\mu = \frac{1}{nqR_{2DEG}} \frac{L}{W}$$

With strong magnetic fields at low temperatures the electrons complete entire cyclotron orbitals before scattering, giving rise to quantised energy levels called Landau levels and the density of states is reduced to delta functions. The Hall conductivity then shows plateaux at integer multiples ν of e^2/h:

$$\sigma_{xy} = \nu \frac{e^2}{h} \quad (\nu = 1, 2, 3 \ldots)$$

This is referred to as the quantum Hall effect [4, 5]. As each Landau level passes the Fermi energy the electrons are free to flow as current, giving oscillations in the diagonal resistivity which are periodic in $1/B$, called the Shubnikov-de Haas effect. In practice both the Landau levels, and hence the quantum Hall plateaux, are smeared by disorder [6]. See Figure A.2 for an example. Electrons near the tails of the peaks are localised and do not contribute to transport, while Landau levels which lie below the Fermi level rise at the material boundaries. The line where such a Landau level intersects the Fermi energy locates a propagating mode called an edge state, and when the Fermi energy lies between bulk Landau levels it is only these modes which contribute to transport.

2.2.2   1D electron systems

If a channel width has $W < l_e$ then transport is quasi-ballistic; scattering off the sides of the channel competes with scattering off impurities. If both W and $L < l_e$ then fully ballistic motion gives quantised levels of conductance through the channel. The confining potential along the channel can be approximated as harmonic. This gives solutions to the Schrödinger equation of parabolic dispersion relations, separated in energy, each one an independent one-dimensional subband. The total current through a ballistic channel due to a source-drain bias voltage V_{sd} is found by summing over N

occupied subbands with $E \leq E_f$. The bias voltage separates the chemical potential of forward moving μ_f and backward moving μ_b electron states such that

$$I = nev = \sum_1^N \int_{\mu_b}^{\mu_f} ev(E)n_{1D}(E)dE$$

Substituting the electron velocity

$$v(E) = \frac{1}{\hbar}\frac{dE}{dk}$$

and the spectral density

$$n_{1D}(E) = \frac{g_s}{2\pi}\frac{dk}{dE}$$

where g_s is the spin degeneracy factor, the current is given by

$$I = \sum_1^N \frac{g_s}{h}e(\mu_f - \mu_b) = N\frac{g_s e^2}{h}V_{sd}$$

Here $g_s = 2$ therefore conductance $G = I/V_{sd}$ increases by steps of $2e^2/h$ in zero magnetic field as each subband becomes populated. If the spin degeneracy is broken, such as from a strong parallel magnetic field, there may be plateaux at Ne^2/h, and at finite temperatures thermal excitations cause these steps to be rounded. The plateaux are fully washed out at temperatures of order $T \approx \Delta E_{sub}/3k_B$ for subband energy spacing E_{sub}. This result corresponds to the Landauer formula, which expresses conductance as a sum over the transmission probabilities $T_{N\sigma}$ for a number of discrete channels N with spin σ:

$$G = e^2/h \sum_{N,|sigma} T_{N\sigma}$$

where for a ballistic system $T_{N\sigma}$ is either 0 or 1 [7]. In practice even ballistic channels will always give some measured resistance, because of reflections where the channel couples to the reservoir.

2.2.3   0D electron systems

A potential minimum defined in all three dimensions creates a quasi-0D system - an artificial atom or quantum dot [8]. For a few electrons in a small space the single-particle energy separation $\Delta E = 2E_f/N$ between levels becomes large, and the spectral energy density a well-defined δ-distribution. In a quantum dot connected by 2 leads to an electron reservoir, if the potential barriers at the leads have the same transmission
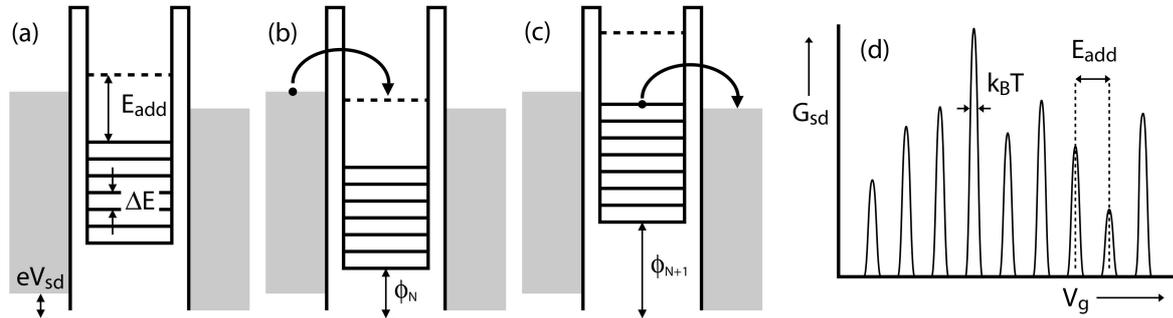
Figure 2.1: Coulomb blockade in a quantum dot. (a) When the thermal (k_BT) and bias (eV_{sd}) energies are smaller than the dot charging energy e^2/C conductance is blocked for certain ranges of the dot electrostatic potential. (b) & (c) Shifting the electrostatic energy appropriately brings an empty level into resonance and conduction through the dot occurs by oscillation between N and $N+1$ electrons in the dot. (d) Peaks in conductance are seen as a function of electrostatic energy, or equivalently gate voltage, separated by $E_{add} = e^2/C + \Delta E$ and of width k_BT (from [10]).

and reflection coefficients, the system behaves in the same manner as a Fabry-Pérot interferometer. This gives oscillations in the transmission probability and the electrons are partially localised between the barriers. A resonance occurs when the phase change between the barriers is an integer multiple of 2π, which is also the point when an eigenenergy of the dot passes the reservoir chemical potentials.

As the gate voltage on the dot is made more negative the barriers eventually rise above the Fermi energy, and the conductance drops below $2e^2/h$, therefore electrons are now localised either inside or outside the dot, and the charge on the dot is quantised to integer multiples of e. Due to the small size of the dot the energy to add an additional electron, $U = e^2/C$ for dot capacitance C, is large [9]. At low temperatures when the charging energy exceeds thermal energies ($T < e^2/Ck_B$) charge transport is suppressed and Coulomb blockade oscillations observed, where the conductance as a function of gate voltage consists of a series of peaks with period $\Delta V_G \sim e/C$. Single electron tunnelling occurs when a positive change in gate voltage lowers the energy of the lowest unoccupied dot level below the highest occupied level in one lead, so an electron tunnels into the dot. This raises the dot's chemical potential above that of the other lead, so the electron tunnels onwards, and the dot's chemical potential returns to below that of the first lead. This allows a further electron to tunnel (Figure 2.1).

2.3 Low dimensional transport in GaAs

2.3.1 Heterostructure wafers

Gallium arsenide heterostructures combined with surface gates can be used to define a variety of low-dimensional electron structures. GaAs is a III-V semiconductor with a zinc blende structure and an f.c.c space lattice, but the primitive basis has two different atoms, at co-ordinates $(0,0,0)$ and $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, centred at each point of the lattice. The creation of a 2DEG in GaAs employs the fact that the lattice spacing in GaAs (5.653 Å) is very similar to that of AlAs (5.660 Å), while the band gaps of the two are different. This permits continuous lattices with little strain to be made with a wide range of band gaps, determined by replacing a fraction, x, of Ga atoms in GaAs with Al atoms. The compound $Al_xGa_{1-x}As$ has a band gap approximately described by $E_{gap} = 1.424 + 1.247x$ eV at room temperature [11].

By using molecular beam epitaxy (MBE), in which individual atomic layers of evaporated source elements are deposited onto a heated GaAs substrate wafer under ultra high vacuum [12], heterostructures can be grown whereby individual layers of $Al_xGa_{1-x}As$ have different doping and mole fractions, since the deposition rate of around 1 monolayer per second gives impinging atoms enough time to migrate across the surface, permitting nearly atomic smoothness [13] . A two-dimensional electron or hole gas can be generated at a distance from the donors by band engineering. This is achieved by complex and non-uniform doping to modify the local carrier densities and chemical potential, and bend the bands. Because the additional donors are not in the conduction region scattering is reduced, and the mobility can be determined such that the mean free path is of the order of microns.

A typical GaAs heterostructure to generate a 2DEG (Figure 2.2a) consists of a thin GaAs cap to protect the Al from oxidation, then Si doped AlGaAs followed by an undoped AlGaAs spacer layer to separate the impurity ions from the next layer, the GaAs quantum well. The substrate is also GaAs. The donor electrons from the Si move to the top of the lower GaAs layer, which has a lower Fermi energy, staying close to the interface because of the attraction of the Si ions, and hence locally bending the conduction band to lower energy. If this triangular well dips below the Fermi level a conducting plane is formed, and if it is sufficiently narrow, the vertical electron wave-functions are restricted to a single state, making transport two-dimensional (Figure 2.2b). To make electrical contact to the 2DEG a local n^+ region is created connecting the device surface to the 2DEG by annealing with a metal alloy. This enables the

11
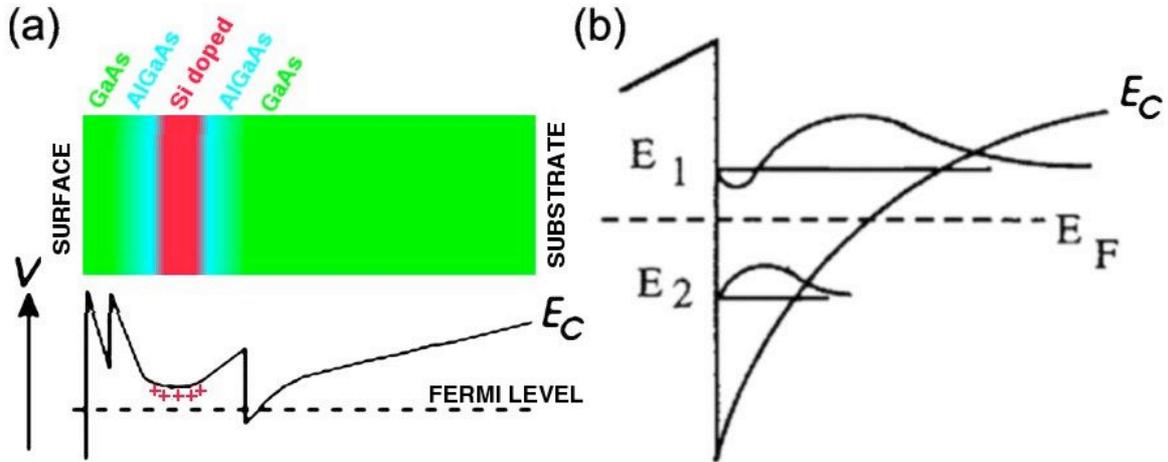
Figure 2.2: GaAs heterostructure (a) A schematic diagram of a typical heterostructure to create a 2DEG, with the associated conduction band. (b) A close-up of the conduction band edge at the interface, showing the first two subbands in the triangular potential well (adapted from [2]).

metals atoms to diffuse down through the top layers of the heterostructure. This is referred to as an 'ohmic contact'.

2.3.2 GaAs devices

To create confinement within the 2DEG the top dopant layer can be etched away, or metal gates are deposited on top of the heterostructured wafer and a negative bias voltage applied. This repels the electrons in the 2DEG until there is total depletion ('definition') beneath the gates. Such a device is a modulation-doped field-effect transistor, or high-electron mobility transistor (HEMT). Without the gate the wafer is often referred to as 'HEMT material'. If two such gates are laid close together such that the width of the gap between them is comparable to the Fermi wavelength of the electron gas, electron transport along the length of the gap becomes one-dimensional [14] (Figure 2.3). Such a construct is known as a split gate or quantum wire and can be used to demonstrate quantised conductance [15]. A quantum dot is defined by adding further narrow surface gates at either end of a short 1D channel to act as tunnel barriers.

To observe electron charging and transport in a laterally confined quantum dot a one-dimensional ballistic channel is fabricated a short distance away ($\sim$200 nm [16]) in a separate circuit, with the inter-lying gate set sufficiently negative that no tunnelling
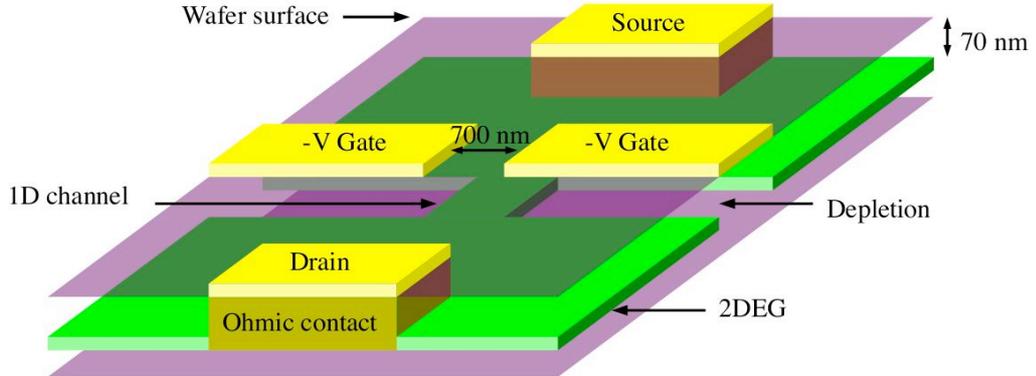
Figure 2.3: 1D split gate. A negative voltage depletes the 2DEG beneath the surface gates, defining a 1D channel in the narrow region in between which is measured through the ohmic contacts.

can occur between the dot and the 1D channel. This permits the single-electron charging energy of the dot to be measured non-invasively (no current flows between the device being measured and the device measuring). When biased in the tunnelling regime the 1D channel is very sensitive to neighbouring electric fields and hence to the potential on the dot, such that there are small dips in the resistance of the channel as the voltage on the plunger gate is swept, corresponding to Coulomb blockade oscillations in the dot. Conversely Coulomb blockade through a quantum dot can be used to infer the local density of states of a higher dimensional electron system [17].

2.4 Carbon nanotubes and graphene

2.4.1 History

The discovery of buckminsterfullerene by Kroto et al. [1] in 1985 opened the door to a whole new branch of solid state physics. This carbon molecule (C_{60}) and related fullerene spheres are structured as a single sheet of graphite (graphene) wrapped into a ball, with some hexagonal rings replaced by pentagons to reduce strain. Hollow graphitic cylinders were first identified by Radushkevich and Lukyanovich [19] in 1952, but it was a 1991 paper by Iijima [20] which brought them to widespread attention [21]. They synthesised multi-walled carbon nanotubes (MWNTs) using an arc-discharge method (an electric arc is introduced between two graphite electrodes in a helium atmosphere), and less than two years later single-walled nanotubes (SWNTs)
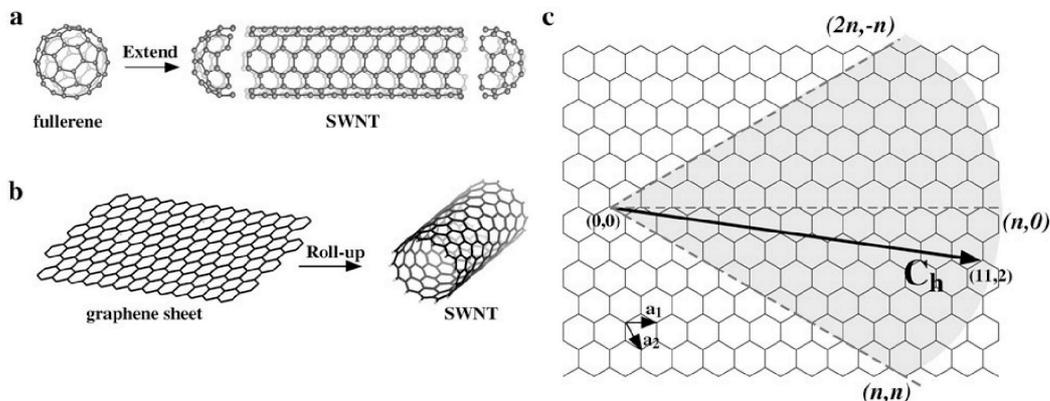
Figure 2.4: sp^2 **bonded carbon structures.** (a) Carbon nanotube structure as related to fullerene spheres and (b) graphene. (c) The 'roll-up' vector $\mathbf{C}_h$ on a graphene sheet defines the circumference, and hence the structure, of any carbon nanotube (from [18]).

were observed independently by both Iijima and Ichihashi [22] and Bethune et al. [23]. SWNTs are a graphene sheet rolled into a cylinder, and capped at either end by a fullerene hemisphere. MWNTs consist of between 2 and 50 such tubes nested coaxially. The length of SWNTs depends on the growth method employed, but the longest recorded length is over 4 cm [24], and their diameters range from 0.4 to around 50 nm, giving them one of the largest aspect ratios of any known material. High-yield production of SWNT ropes by Thess et al. [25] in 1996 paved the way for widespread research. These are bundles of SWNTs tens of nanometres in diameter, bonded by the same strong van der Waals forces which hold layers of graphite together. Nanotubes have been found to have phenomenal mechanical properties, being not only extremely stiff, but also able to buckle elastically under large strains [26], but it is their electronic properties which are of primary interest here.

The small sizes and symmetries of these new materials make them intrinsically low-dimensional electron systems. Thus carbon had been found to form structures with 3D (graphite and diamond), 1D (carbon nanotubes) and 0D (fullerene spheres) properties. The missing element was two-dimensional graphene itself, assumed not able to exist in the free state [27]. Any strictly 2D crystals were believed to be thermodynamically unstable [28], and in graphene itself open carbon hexagons have so much bond energy that they were expected to fold or re-arrange into curved shapes [29]. In 2004 Novoselov et al. successfully prepared isolated thin sheets of graphite, including single

layer graphene, on a silicon substrate using an exfoliation technique employing Scotch tape [27]. Figure 2.4 shows the relationship between these sp^2-bonded structures.

2.4.2 Graphene: electronic structure and properties

Graphene is a one atom thick sheet of carbon atoms arranged in a honeycomb lattice. It is a semimetal, or zero-gap semiconductor, whose valence and conduction bands vanish linearly with no energy gap at six $\mathbf{k}_F$ points [30] at the corners of the hexagonal first Brillouin zone (Figure 2.5a). At low energies near these points the dispersion relation is linear:

$$E = \hbar v_f \sqrt{k_x^2 + k_y^2}$$

The carriers therefore have zero effective mass and behave as relativistic fermions. This occurs because the interaction of the electrons with the periodic lattice potential gives rise to new quasiparticles described by the (2+1)-dimensional Dirac equation [28]. They have an effective speed of light $v_f \sim 10^6 ms^{-1}$ which is independent of energy therefore backscattering is forbidden and transport is intrinsically ballistic. They also exhibit an anomalous Berry phase which introduces a shift of $1/2$ into the quantum Hall effect filling factor, so that $\nu = \pm 4(n + 1/2)$ for integer n [31]. This introduces an additional Landau level at zero energy, shared by both electrons and holes. The prefactor 4 reflects two-fold spin and two-fold valley degeneracies.

Bilayer graphene is also a zero-gap semiconductor, but a band gap is opened in an applied electric field [32]. It does not exhibit the $1/2$-offset quantum Hall effect or the $n = 0$ plateau. Overlap between the conduction and valence bands appears with three layers and multiple layers approach the 3D limit of graphite at ten layers, therefore 3-9 layers are classed as few-layer graphene (FLG) [28]. There are around 3 million layers in a millimetre thickness of graphite.

Graphene has been shown experimentally to exhibit an ambipolar field dependence and room temperature ballistic transport [33]. Its room temperature electron mobility is up to 15,000 $cm^2V^{-1}s^{-1}$ [28], exceeding that of a GaAs HEMT material and approximately temperature independent in the range 10-100 K, therefore defect scattering is dominant [34]. Novoselov et al. [35] have observed the quantum Hall effect in graphene at room temperature. Despite the zero carrier density predicted at the Dirac point graphene's minimum conductivity is found to be of order, but not always equal to, $4e^2/h$. It is subject to temperature and other environmental hysteresis, and ascribed to local ionised impurities and trapped charges in SiO_2 causing conducting puddles [28].
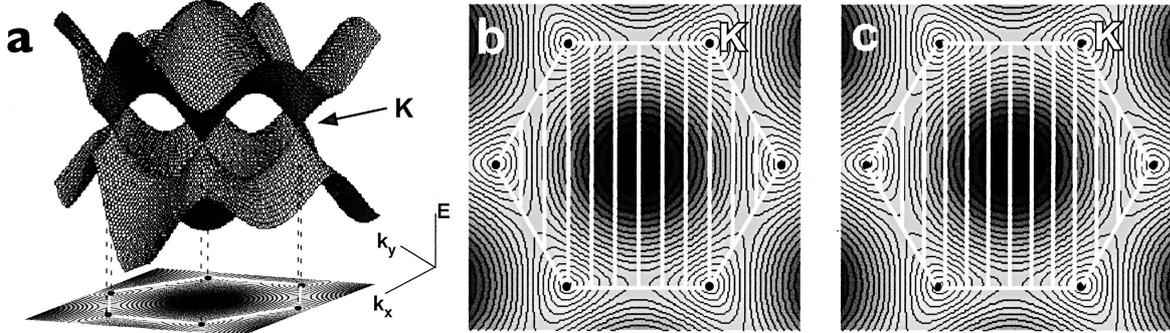
Figure 2.5: Graphene band structure. (a) Conduction and valence band surfaces of graphene, showing the six $\mathbf{k}_F$ points where they touch. (b) Slices through the band surfaces representing allowed subbands in a (9,0) metallic tube. (c) The same for a (10,0) semiconducting tube, where the bands no longer intersect the $\mathbf{k}_F$ points (from [36]).
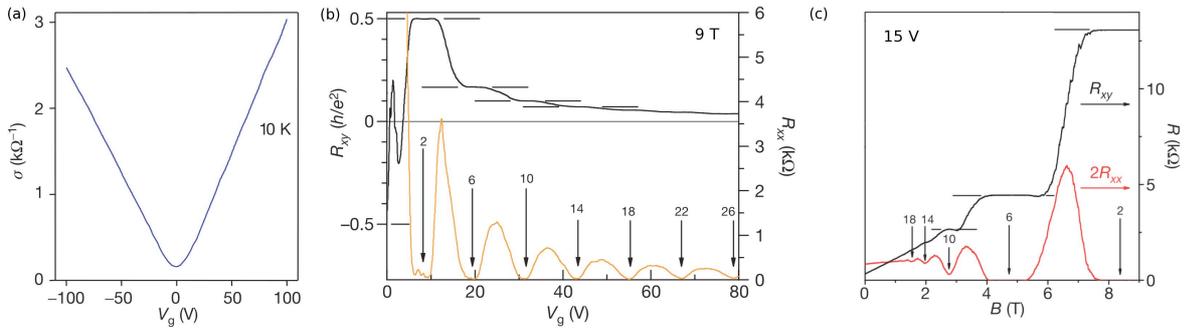
Figure 2.6: Low temperature graphene transport. (a) Graphene conductivity as a function of gate voltage V_g, showing linear gate dependence and a finite minimum conductance at the Dirac point (from [34]). (b) & (c) Quantised magnetoresistance (red/orange) and Hall resistance (black) at 30 mK, with minima and plateaux at expected filling factor intervals (adapted from [37]).

2.4.3 Carbon nanotubes: structure and properties

Although they grow axially [38] (Figure 2.4a), when considering their electrical properties carbon nanotubes are more usefully visualised as a graphene sheet rolled into a cylinder (Figure 2.4b). The structure of individual tubes is defined by the vector joining two lattice points on the graphene sheet, which corresponds to a circumference of the rolled tube. This 'roll-up' vector $\mathbf{C}_h$ is a linear combination of the graphene lattice vectors $\mathbf{a}_1$ and $\mathbf{a}_2$ (Figure 2.4c):

$$\mathbf{C}_h = n_1\mathbf{a}_1 + n_2\mathbf{a}_2 = A(n_1, n_2)$$

Within a viable range of diameters (0.4 to ∼50 nm) it can connect any two lattice points, and even allowing for symmetrical repetitions (only one 12th of the Bravais lattice ($0° \leq \theta \leq 30°$) is needed for all unique structures [39]), there are an enormous variety of different possible structures, and their electronic properties vary accordingly. While a growth process will tend to bias for a range of diameters, the mix of structures and conduction types produced within the range is generally random and not controllable.

Extensive modelling has been done of the electronic structure of carbon nanotubes using tight-binding calculations [38–40]. For large diameters the electronic structure is approximated by that for a graphene sheet, but with an additional boundary condition imposed by rolling the sheet into a tube. While graphene is two-dimensional, the small diameters and additional symmetries of nanotubes make them one-dimensional systems, and their bands must also be continuous around the tube circumference. Therefore a nanotube bandstructure consists of periodic vertical slices through the graphene band surfaces, with allowed wave vectors $\mathbf{k}$, denoting the 1D subbands, given by $\mathbf{C}_h.\mathbf{k} = 2\pi m$ (for integer m). If one of these subbands passes through one of the $\mathbf{k}_F$ points (if $(n_1 - n_2)/3 = q$, for integer q), then the nanotube remains a zero-gap semiconductor (Figure 2.5b), i.e. first-order metallic. Otherwise it has a band gap and is semiconducting (Figure 2.5c). This would imply that 1 in 3 tubes are metallic.

When the curvature of the tubes is also taken into account there is some mixing of the π/σ bonding and π^*/σ^* anti-bonding orbitals. Therefore the π-only approximation above is no longer valid, and the $\mathbf{k}_F$ points shift slightly from the vertices of the first Brillouin zone. This breaks the periodicity and introduces small band gaps in all the first-order metallic tubes except achiral 'armchair' (n, n) tubes, which remain truly metallic [40]. The band-gap of the other, moderate band-gap tubes has been predicted [41] and shown experimentally [42] to be inversely proportional to the tube diameter.

Contact with the substrate usually results in disorder in one-dimensional metallic wires, which localises the electrons. White and Todorov [43] predicted that because of the stability of the nanotube structure the disorder experienced is averaged over the circumference of the tube, implying localisation lengths increase as the tube diameter increases, potentially approaching 10 μm compared with 40 nm for copper. Also, ballistic transport in conducting polymers and other 1D systems is impaired by a Peierls instability, where an energetically favourable periodic lattice deformation breaks the 1D symmetry. However, in nanotubes any such deformation must be spread around the whole circumference and is resisted by their stiffness - the lattice cost of rearranging atoms around the whole circumference is low compared to the gain in electronic energy from only 2 available subbands [44] - so the transition from Peierls-distorted to high temperature metal occurs below room temperature [38]. Tans et al. [45] have demonstrated experimentally single electron ballistic conduction in metallic tubes, and observed extended, phase coherent molecular orbitals as long as 140 nm (the distance between their contacts).

The mismatch between the small number of discrete nanotube states (modes) and the continuous states of metal electrodes leads to a quantised contact resistance

$$R_Q = \frac{h}{2e^2 M}$$

for a number of modes M lying between the Fermi levels of the electrodes. For a metallic SWNT $M = 2$ therefore $R_Q = h/4e^2 = 6.45$ kΩ, the theoretical minimum resistance in a 2-terminal measurement, given ballistic transport and transparent contacts [46].

2.5 Applications

GaAs and related materials have been widely used as tools to further explore low dimensional and quantum phenomena, but are used in only limited electronics and computing applications. Carbon nanotubes and graphene meanwhile have broad potential for both fields. They have a substantial advantage by being inherently low dimensional and exceptionally small structures, and, if and when the hurdles associated with their fabrication and handling are overcome, could lead to functionally complex devices with a compactness and structural simplicity not feasible with semiconductor technology alone. They also exhibit low-dimensional properties at temperatures far above those required with GaAs; strong confinement around the nanotube circumference and their resistance to Peierls distortions leads to large spacing between 1D subbands and

a 1D nature preserved even to beyond room temperature [44]. Fabrication techniques are already sufficiently established for the integration of carbon nanotubes into device research, where they are being used for electronics and optoelectronics experiments previously involving semiconductor 1D wires and quantum dots (see Section 3.1.2).

A key direction of research is the development of both carbon nanotube and graphene transistors, to address the problems currently facing the micro-processor industry due to the size and efficiency limitations of silicon technologies. Some new approaches adopt novel environments such as spin-based devices or photonics, while others replace key components in existing devices with new technologies and materials, for which carbon is particularly promising. Tans et al. [47] first recorded room temperature transistor behaviour in a nanotube and they have since been made into working logic circuits [48].

Unlike GaAs quantum dots, carbon nanotubes can have ferromagnetic or superconducting leads [49], enabling the study of novel physical phenomena such as superconducting correlations and spin injection into quantum dots [50]. Kasumov et al. [51] have shown proximity induced superconductivity in a tube suspended between superconducting electrodes and Hueso et al. [52] have used spin polarised contacts for molecular spintronics using nanotubes.

A further application is solid state quantum computing [53, 54]. Quantum computation has applications for solving certain problems which are currently very difficult with conventional, classical computers, but has yet to be realised on a useful scale. Classical computers handle information in strings of binary bits, each with value 0 or 1, and can perform a sequence of operations on only a single string of bits at one time. The quantum computing analogue is the qubit, a quantum state defined as the superposition of two orthonormal basis states $|0\rangle$ and $|1\rangle$. Therefore, instead of the two discrete states of a bit, a qubit can be in any of an infinite number of possible states, defined on and within the Bloch sphere : $\psi = cos\theta|0\rangle + e^{i\phi}sin\theta|1\rangle$. This enables a quantum computer to perform the same sequence of operations on a superposition of many input strings; a quantum computer with N qubits is capable of 2^N simultaneous operations.

However, although the output register of such a computer is left in a superposition of all possible results generated from each state in the input register, quantum measurement can only ever read out a single answer, and this will not necessarily be the result sought. A quantum computer is therefore not deterministic, and only becomes a useful system when repeating an algorithm the necessary number of times to ensure a statistically likely correct result is still faster than attempting the same calculation
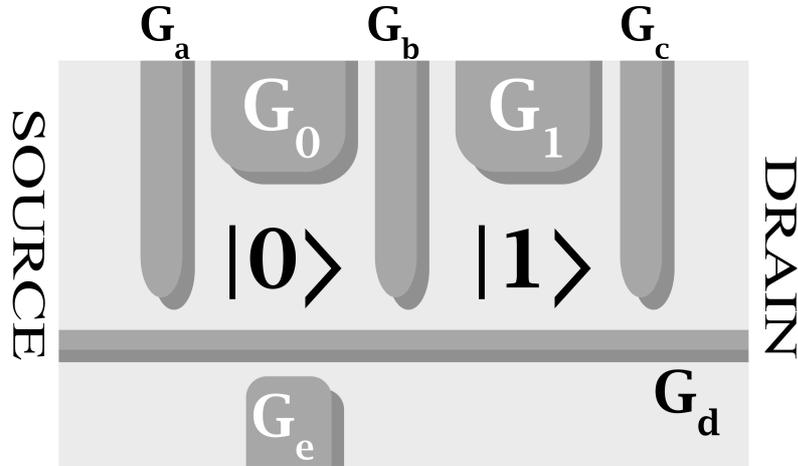
Figure 2.7: A solid state qubit device, using gates to define two lateral quantum dots in series, with four tunnel barriers G_{a-d}, two plunger gates G_0 and G_1, and G_e together with G_d for a voltage probe on the left-hand dot.

on a classical computer. This is true for applications such as the factorisation of large numbers [55] and searching an unstructured database [56], and complex quantum computers will be particularly valuable for investigating quantum systems which test the limits of classical modelling, such as protein folding and molecular orbitals. Therefore the first quantum computers will also most likely be application-specific.

The key hurdle in the development of quantum computation is finding an effective physical implementation which will also be scalable, in the same way that integrated circuits facilitated the development of classical computers. While it looks increasingly as though truly useful quantum computers will have to wait for some as yet unknown technology, systems under investigation include nuclear spin resonance, quantum optics and solid state electronics. A solid state charge qubit can be defined in two GaAs dots in series (a double quantum dot), where the left and right dots correspond to the states $|0\rangle$ and $|1\rangle$ respectively (Figure 2.7). The left-hand barrier is lowered to allow one additional electron into the left-hand dot (this electron is then in state $\psi = |0\rangle$), then the central barrier is lowered to permit Rabi oscillations of the electron between the two dots. Raising the barrier again leaves the electron in a superposed state ($\psi = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$). A one-qubit operation such as a phase rotation can now be performed by changing the voltage on one of the plunger gates to change the energy of a dot. If the central barrier is again lowered then raised, the electron is left in the state $\psi = cos\theta|0\rangle + e^{i\phi}sin\theta|1\rangle$. A measurement is made on the system by determining

the polarisation of the qubit with the non-invasive probe [57]. Such qubits can be combined into logic gates by using one pair of polarised dots to polarise adjacent pairs in rows or in branching channels. Two qubit operations have been demonstrated with 4 dots arranged as a square [58].

One of the chief obstacles to creating a multi-qubit computer is quantum decoherence. A qubit state cannot be completely isolated from the macroscopic world, and as it starts to couple with the near infinite number of degrees of freedom in the surrounding environment the coupling within the superposition is dissipated and the coherence is lost. The qubit reverts to a simple classical bit. The time for this decay to occur is called the 'decoherence time' and is the maximum time available for the completion of a calculation. Decoherence times are typically very short on a human time scale, hence the scarcity of readily observable quantum effects; the decoherence time for a molecule in a laboratory vacuum is around one hundredth of a femtosecond.

Decoherence processes in solid state qubits are dominated by electron-phonon scattering, which is enhanced because GaAs is piezoelectric. Decoherence times vary inversely with temperature and have been measured as 1 ns [59] at 20 mK, but are predicted to saturate at around 3 ns [60]. Phase coherence times in nanotubes are predicted in the range 0.01 to 4 ps [61] and measured up to 10 ps at 1K [62], which is substantially lower than in GaAs, but they are expected to have longer spin decoherence times. A double dot can be used as two interacting spins, and since ^{12}C has no net spin and ^{13}C is rare, hyperfine coupling is reduced [63], so that spin decoherence times of tens of microseconds are expected [64]. Together with their small dimensions and structural simplicity this makes nanotube quantum dots a promising architecture for solid state qubits.

Chapter 3

Carbon Nanotube Devices on Gallium Arsenide

3.1   Introduction

The integration of carbon nanotubes and GaAs heterostructures has potential for improved function and complexity in low-dimensional electronic devices. This chapter describes successful fabrication of nanotube-on-GaAs devices. The interaction between nanotubes and the GaAs surface is explored with different dispersion fluids and techniques and the first reported AFM manipulations of nanotubes on GaAs. Transport measurements are presented from completed devices, along with experiments employing concurrent advances in the field to improve device fabrication.

The dispersion testing was carried out from February to November 2005, the first generation of devices was fabricated and tested from December 2005 to October 2006, and the second generation from September 2007 to February 2008. The introductory sections describe the state of the field at the time the work was begun and later sections reference reported developments as it progressed.

3.1.1   Nanotube synthesis

The nanotubes produced by Iijima [20] were found amongst the products of an arc discharge process used to make fullerenes, where an electric arc is introduced between two graphite electrodes in a helium atmosphere at temperatures as high as 3000-4000℃. They were first produced on a macroscopic scale by Ebbesen and Ajayan [65] and this process achieved up to a 90% yield of SWNTs [66]. An alternative, laser ablation

technique gave 95% SWNTs, with diameters determined by the reaction temperature [67]. A pulsed laser vapourises a graphite target in a high temperature reactor under inert gas, and nanotubes condense on the cooler reactor walls [68]. Chemical vapour deposition (CVD) was first used to grow nanotubes by José-Yacamán et al. [69] in 1993 and developed into the 'high pressure carbon monoxide' (HiPCO) process by Smalley et al. [70]. They added iron pentacarbonyl to a carbon monoxide gas flow, where it decomposes on heating, forming catalytic irons clusters on which SWNTs nucleate and grow in the gas phase. This technique has a high yield (up to 97 mol %) and they currently make around 1 g per hour for commercial use.

All these techniques produce nanotube soot in the synthesis vessel, consisting of bundled and knotted ropes of tubes. To be useful for single-nanotube electronic devices nanotubes must be separated from the bundles and applied to the device substrate, without contamination from residual catalyst and amorphous carbon. A significant improvement was made by Kong et al. [71] in 1998, who employed CVD to grow high quality SWNTs directly onto a silicon substrate chip using patterned iron-based catalyst islands as growth seeds for a methane feedstock. At 1000℃ this gave a sparse dispersion of long and largely defect and impurity free nanotubes at predictable locations, suitable for contacting to make devices. Further developments achieved wholly isolated and even aligned SWNTs [72] at chosen sites, and a low temperature bound on silicon of 750℃ [73]. Good yields have been achieved below 600℃ on a zeolite substrate [74].

3.1.2 Nanotube quantum dots

Quantum dots are readily defined in carbon nanotubes because of the ease with which potential barriers can be introduced. The mismatch between the finite number of modes in a nanotube and the contact metal continuum creates a minimum contact resistance (see Section 2.4.3). Schottky barriers are also often formed at the contacts (see Section 3.4.5) so back-scattering occurs at the metal-nanotube interface, making the nanotube a finite length, quantised object [49]. The intrinsic elastic mean free path of carbon nanotubes is of the order of microns; 10 μm and 0.5 μm at 1.6 K and 300 K respectively, compared with 100 μm and 0.06 μm in GaAs [62]. However, length-independent resistance is typically only seen under 200 nm [78] because of disorder due to structural defects and contact with the substrate. Ballistic transport and single quantum dots are observed in suspended metallic nanotubes up to 1 μm long [79]. Clean quantum dot behaviour is rarely achieved in semiconducting SWNTs, as they
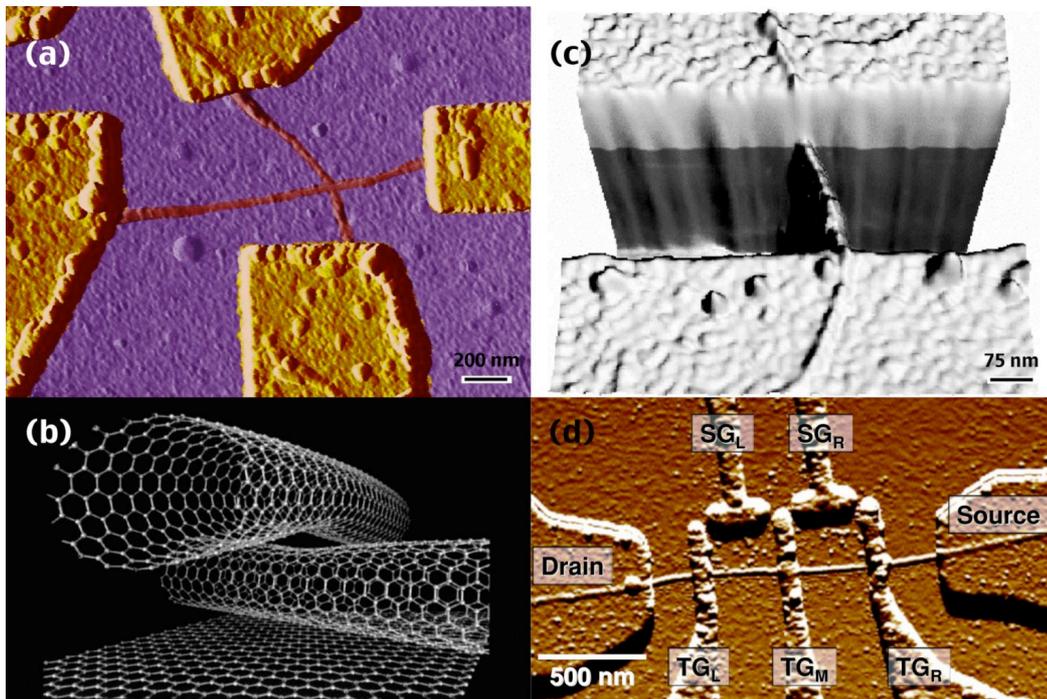
Figure 3.1: Tunnel barriers in nanotubes: (a) a crossed nanotube device (no. E2t4-2/3), (b) compression in crossed nanotubes (from [75]), (c) a nanotube lying across an etched trench (from [7]) and (d) a nanotube double dot (from [76]).
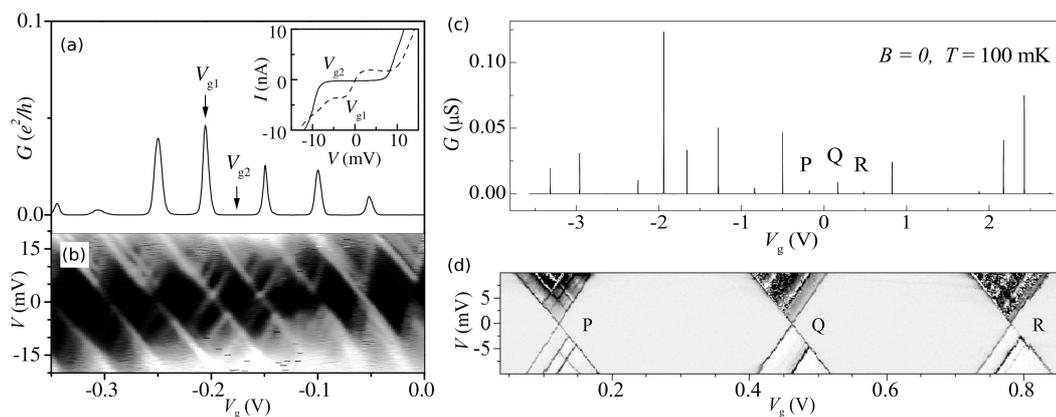
Figure 3.2: Quantum dot transport in a nanotube: (a) Regular Coulomb blockade peaks in the gated conductance of a nanotube at 4.2 K. (b) Differential conductance (lighter = more positive) showing Coulomb diamonds with changing bias voltage for the same gate range. (c) & (d) A similar device measured at 100 mK, showing sharper peaks and more regular diamonds (from [77]).

are more susceptible to disorder, resulting in multiple islands rather than a single well-defined dot [80].

Tunnel barriers can be defined at chosen locations by kinking or compressing the nanotube. Laying gates over a nanotube divides it into segments [81], used to define and measure double quantum dots [50]. A nanotube laid over an edge or trench sags [82] and barriers can be accurately defined with AFM manipulations (see Section 3.3).

Nanotubes are sufficiently strongly attracted to substrates that they bend to conform to the substrate topology. With crossed nanotubes this creates an inter-tube force of ∼35 nN (with a pressure of 10 GPa [83]) compressing both nanotubes to create two pairs of quantum dots [84] (Figure 3.1). Gating can be done with a conventional back gate, another nanotube [85] or an AFM tip [86].

3.1.3 Experimental aims

Nanotubes have significant potential as components within electronic circuits, especially when integrated into other technologies such as CMOS devices. The best route to high quality tubes is in situ CVD growth but the temperatures required exceed 500℃ and would destroy CMOS circuitry. They could be dispersed on the substrate from solution, but silicon CMOS devices need to be patterned in advance and there is little chance of an isolated nanotube falling onto exactly the correct location. It is possible to move nanotubes across the surface (see Section 3.3) but very time-consuming for many devices and far from reliable. Although GaAs heterostructures also do not survive temperatures above 500℃ (morphology distortion can even begin at 300℃ [87]) they do have the advantage that electronic structure within the substrate can be created after the chip has been fabricated, by defining channels in the 2DEG with metal top-gates. Nanotubes can therefore be deposited arbitrarily on the surface and the circuitry designed around them. In particular the charge-sensing capability of a point contact in the 2DEG might be used to measure electron movement through a nanotube quantum dot on the wafer surface, with potential for use in a qubit system, while conversely a SWNT might be used as a 1D channel to detect charge transport through a quantum dot in the 2DEG.

There are several challenges involved in fabricating and measuring such devices. The nanotubes must be dispersed on the GaAs surface suitably isolated, long and clean, and in appropriate locations for subsequent metal patterning of contacts and gates. The metal contacts need to be well aligned and make reliable and good contact to the tubes. Gating of the nanotube by the 2DEG must be demonstrated, and vice-

versa, in order for charge detection to be possible. Finally all elements will need to be tuned and balanced to allow both control and measurement of a quantum dot.

3.2 Dispersing SWNTs

Since in situ CVD growth of carbon nanotubes was not possible on GaAs it was necessary to deposit tubes onto the surface. Purified HiPCO SWNTs were purchased from the Smalley Group at Rice (now handled by Unidym [88]), specified as <5% by weight amorphous carbon and catalyst material, with a mean length of 1 μm. They are supplied as dry nanotube soot, which consists of tightly tangled and knotted ropes of nanotubes, from which isolated SWNTs need to be separated and applied to the substrate at an appropriate density. The extended π electron system is highly polarisable, creating attractive van der Waals forces, and the linear nanotube structure leads to co-operatively enhanced attraction, for strong side-by-side aggregation into ropes. While ropes can be untangled relatively easily with a shear stress, such as by ultrasonication, it is harder to separate individual nanotubes; their low entropy of mixing and high molecular weight necessitates strong dispersing interactions.

The geometry of the devices needed for this experiment is too precise for the arbitrary dispersion of nanotubes over pre-laid contacts, therefore a dispersed density is needed such that circuitry can be designed around individual, isolated nanotubes. The key stages are suspension of the soot in a suitable fluid to lift the tube network, ultrasonication to separate the tubes, application of the suspension to the substrate, and removal of the suspending fluid, leaving the nanotubes in place. Reported work uses a very wide variety of chemicals and methods to suspend and disperse nanotubes. Even where equivalent materials or techniques are used, reported quantities and conditions rarely match. The references below give a broadly complementary survey of the better established and understood techniques.

3.2.1 Background

A simple approach is suspension in a solvent, with dichloroethane most commonly used [62, 75–77, 83, 84, 89–94]. Chloroform [95], ethanol [96] and acetone [97] are also reported as effective. Ausman et al. [98] have done a comparative study and find the best solvents for nanotubes to be amides, in particular non-hydrogen bonded Lewis bases such as N-methylpyrrolidone (NMP). Around 1 mg of nanotubes can be suspended well in 10 ml of solvent [95]. Nanotube solubility in organic solvents is

enhanced by functionalising them with amides, attached at opened ends or sidewalls, but this shortens the tubes to around 300 nm [99].

An alternative methodology uses surfactants. These bind to the nanotube surfaces and separate them by a steric repulsive force. The most common are anionic surfactants in aqueous solution, such as sodium dodecyl sulphate (SDS) [100–106], the primary agent in soaps and shampoo. The effectiveness of suspension is determined by the concentration of both surfactant and SWNTs. Low surfactant concentrations leave large dense clusters of tubes after sonication, because there is not enough coating for the electrostatic repulsion between surfactant molecules to counterbalance the van der Waals attraction between the tubes. Intermediate surfactant concentrations give a black homogeneous phase of mostly isolated tubes that does not coarsen over several days [106]. High surfactant concentrations give a combined dispersion of single tubes and some light clusters, due to the formation of surfactant micelles; these are unable to fit between close bundles and the resulting osmotic pressure creates an effective attraction (known as a depletion attraction). The optimal composition is reported to be just above the critical micellar concentration, when the surfaces are saturated and the surfactant alone is beginning to aggregate into micelles [105], which for SDS is at 0.35 weight % SWNTs and 1 wt% SDS [101]. The range of stable concentrations have a nanotube:surfactant ratio of 1:5 to 1:10, with no more nanotubes than 0.5 mg/ml [104].

The surfactant is removed by washing in water (Bonard et al. [105] recommend 30 minutes, Li et al. [107] 5 hours) and annealing in air (Islam et al. [104] suggest 4 hours at 180℃). This results in a 2 nm decrease in measured nanotube diameter. Additionally a polyelectrolyte matrix aids dispersion in water, by noncovalently clinging to the neutral nanotube surfaces and providing electrostatic colloid stabilisation [103]. Used together with surfactants they keep the nanotubes stably dispersed for several months.

During ultrasonic agitation cavities are generated in the liquid which cause shock-waves when they collapse. This cuts nanotubes and releases them from ropes and bundles [103]. Because this shortens the tubes a compromise is needed between adequate length and enough isolated tubes; over-shortening the nanotubes enhances aggregation irreversibly [7]. Sonication times needed with dichloroethane are quite long, from 30 minutes [76] to 8 hours [90], and Nygård et al. [77] suggest that dichloroethane enhances the cutting process. Heating the bath helps, but the long times mean that resulting fragments are only 100–300 nm long [98]. The enhanced nanotube separation achieved with surfactants allows for much shorter sonication times; 10 minutes [106] gives a stable suspension, or 5 minutes with an ultrasonic tip [102]. The shorter

sonication time reduces fragmentation, giving a maximum length around 1 μm [103]. The power input also has an effect. A low power, high frequency (12 W, 55 kHz) bath sonicator after 16 hours gave twice the length and twice the yield of a high power (6 W, 22.5 kHz) tip sonicator after only one.

Sonication may damage nanotubes even when they are not cut. The cavitation energy is sufficient to fracture and strip single layers from a MWNT, and sonicated tubes show substantially more buckles and bends than untreated tubes, with an effect relative to the sonication time. The Raman 'D' band, which indicates dislocation defects (at $\sim$1300 cm^{-1}), doubles in relative intensity after only 2 minutes, with a substantial increase by 60 minutes [108]. Electron spin resonance (ESR) broadening, indicating a loss of structural symmetry, also increases after five minutes of sonication, and the temperature of onset of oxidation reduces from 600℃ to 500℃ [93].

Residual bundles and catalyst are removed by centrifugation [98, 105, 106], typically 10 minutes at 10,000 g. Even with a stable suspension it is best to disperse the tubes immediately, before aggregation begins. Nygård et al. [77] find that deposition is improved if the substrate is wet in the suspension first. Techniques include spin coating [82, 84, 90], which achieves good results, dipping the chip [109], followed by a rinse and nitrogen gas blow dry, and pipetting onto the chip, again followed after a few seconds by a rinse and blow dry [7, 76]. The rinse is with de-ionised (DI) water for aqueous surfactants or isopropanol (IPA) for organic solvents. The tubes precipitate out during the time of contact between the suspension and substrate, and stick via van der Waals forces, but if the droplet is left to dry then nanotubes precipitate as small grains [94, 106]. Once on the surface there is no known way to chemically remove them [7]. More elaborate techniques include flow-induced alignment of the nanotubes [101], and depositing the suspension onto a chip with surface acoustic wave (SAW) transducers. The SAWs align nanotubes in the fluid to within 25° and 45° degrees of the direction of propagation due to a combined electric field and mechanical fluidic effect [106].

Several techniques are used to assess the quality of a suspension. Light and X-ray scattering [100] give information about the structure and nanotube dispersion within the suspension, and near-infra-red [108] or UV [98] spectroscopy measure relative concentrations. Dispersion on a substrate is judged by counting and measuring large numbers of nanotubes with an AFM [104].

3.2.2   GaAs texture

During the first tests dispersing nanotubes it became apparent that the top surface of GaAs is heavily textured. Where silicon is nearly monatomically flat, HEMT material exhibits corrugations up to 10 nm deep and ∼150 nm wide. The sharpest features observed had curvature approaching 30° and with radius <10 nm (the resolution limit of the AFM tip). Since nanotubes conform tightly to the substrate topology, this verges on curvature sufficient to induce kinks [110]. While this would be unlikely, the features appear to the AFM of a similar size to nanotubes and the constant height variation would hamper certain modes of AFM manipulations (see Section 3.3). Seventeen wafers with different growth parameters were profiled by AFM, all grown within the Semiconductor Physics Group. When scanning along long features the AFM adjusts adjacent scan lines to a plane-fit and so height variations are smoothed, but the feedback tends to result in a positive adjustment to sudden up-steps. Therefore scanning was done at 45° to the corrugations. A 4μm field of view gave sufficient resolution to see the sharpest features, and also a large number of corrugations for averaging. Noise and dirt was filtered by averaging each pixel to the median of a 20 nm radius. The root mean square height variation was averaged over several cross-sections, perpendicular to the corrugations. The error between measurements was around 10%.

The AFM images reveal two competing morphologies in the heterostructure surface texture. There is a clear trend that increased growth temperature increases the height and density of corrugations (see Figures 3.3 and 3.4). Increased As flux results in much smoother long range texture, but at smaller sizes gives a mixture of scale-like features and dimples. As a result the measured RMS height decreases, as the corrugations are smoothed, but the texture is visibly rougher. Hasegawa et al. [111] have previously recorded a transition from surface ridge structures to mounds as the As flux increases in heterostructure growth, due to abundant As_2 suppressing metal diffusion. Wafer mobility also rises with increasing temperature, approximately doubling from 550℃ to 650℃ (Figure 3.4-inset). The lowest mobility measured (over $2{\times}10^5 m^2$/Vs) is sufficient for the transport measurements intended. These trends are not strong because the wafers were not grown specifically for this test, and other parameters also vary.

Plain doped GaAs also shows short ridges, which increase in height with either increasing growth temperature or As flux - there are not enough data points to be certain which. Yamamoto et al. [112] describe GaAs top-layers grown onto flat (775)B-orientated GaAs substrates at different temperatures. Between 540℃ and 580℃ they appear very similar to those measured here, but by 640℃ the short ridges have morphed
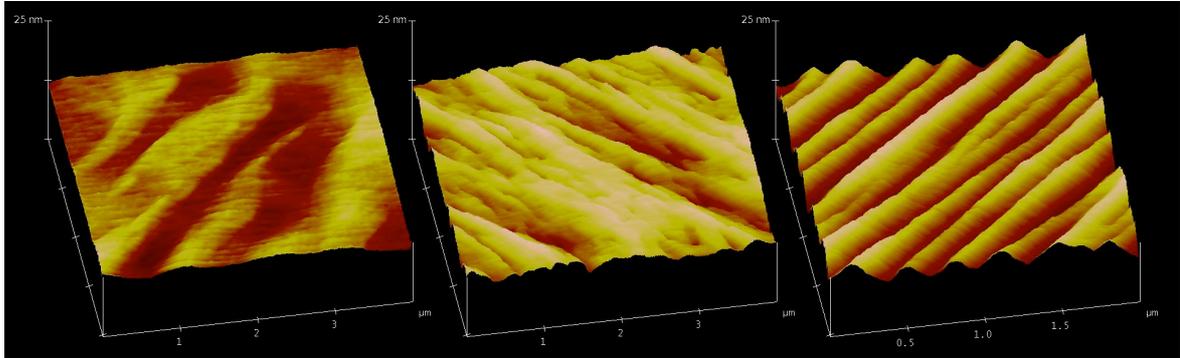
Figure 3.3: GaAs surface morphologies. AFM images, left to right: i) low growth temperature, low As flux ii) low growth temperature, high As flux, iii) high growth temperature. Horizontal and vertical scales differ.
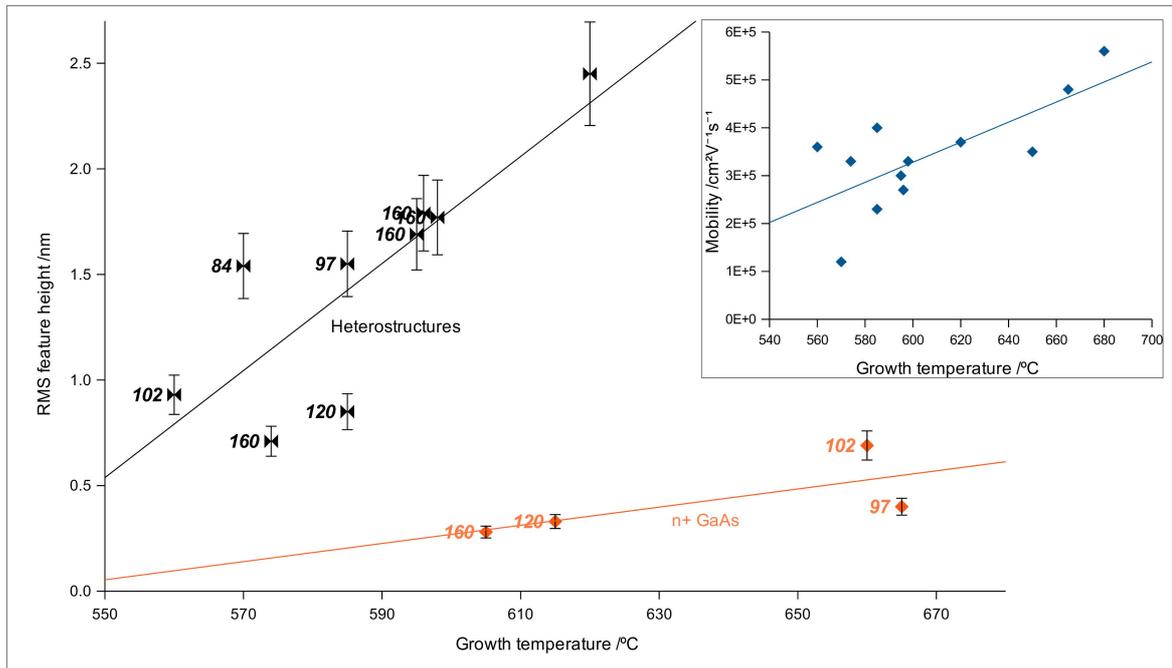
Figure 3.4: Temperature dependence of GaAs surface roughness. Measured RMS feature height for doped and heterostructured GaAs; the number labels are the arsenic flux (nA). (Inset) Dependence of electron mobility on growth temperature in heterostructures. Trend lines are added as a visual guide only.

into long corrugations like the HEMT material, around 4 nm high and 35 nm wide.

3.2.3 Experimental outline

When designing devices suitable nanotubes were to be found by AFM, since unlike SEM this gives confirmation of their diameter. A relatively low final yield was anticipated, and so a high throughput of devices desirable. A 10 μm AFM field of view is approximately the largest for reliably finding nanotubes, since one pixel corresponds to 20 nm, a little less than the AFM tip width and hence the minimum lateral nanotube diameter measured. Such a scan takes around 10 minutes, and assuming that the majority of nanotubes will be unsuitable for devices, being too short or bent, an ideal dispersion would give a couple of nanotubes in each scan; a day's scanning would record 50-60 nanotubes over the entire mesa, of which $\sim$5 might be good for devices.

The only assaying tool available was also AFM, therefore fluid suspensions were judged by eye according to colour and opacity (Figure 3.5a), and substrate dispersions by counting and measuring tubes, over as many AFM scans as was practical. Over 100 chips were prepared from different suspensions comparing different parameters and techniques, and with the substrate dispersion often very uneven in many cases only a qualitative assessment could be made. Safety concerns when disposing of nanotubes and solvents required that all the suspension fluids be collected, and considering also economy of materials, suspensions were usually made up to only 10 ml.

The transfer of a single nanotube from a bundled rope in dry soot to being isolated on a substrate chip is a multi-stage process. The number of potential variables is large, and this work reveals that the more readily reported quantitative parameters do not always have the most significant effect. Initial tests repeated the process previously used successfully in the group; nanotubes were suspended in chloroform with up to 90 minutes of sonication, and a few microlitres of the suspension were pipetted onto the chip and left to dry. The suspension did not visibly darken at all with sonication, and only short ropes with diameters greater than 5 nm were found on the chip, unevenly distributed in rings where droplets had dried.

Although the colour of the fluid did not change, it was apparent that some tubes were being suspended. With light agitation the visible specks move as if in jelly, but with a sharp impulse they start to move freely. This suggests that even in suspension the nanotubes can form a loose extended network, perhaps simply expanded from the initial soot particles. This supports the observations of Schaefer et al. [100] who find that suspended nanotube networks are mass fractal objects, exhibiting similar structure

amongst ropes at large scales as amongst individual tubes at smaller scales. Ethanol and acetone suspensions were even less successful and rubbing a textured GaAs chip against smoother quartz covered in CVD-grown nanotubes did not transfer any.

3.2.4   Sodium dodecyl sulphate

Aqueous solutions of sodium dodecyl sulphate according to the optimum recipe given by Vigolo et al. [101] (0.35 wt% SWNTs and 1 wt% SDS) also did not yield acceptable results. The suspension darkened after 30 minutes of sonication, but nearly half the volume of the fluid was occupied by a large fluff of tangled nanotubes, and the resulting dispersion showed only ropes. If the supernatant was decanted and more water added to the solid residue then nanotubes spontaneously suspended and darkened the fluid, suggesting that the initial volumes of solids were far too high. Had a centrifuge been available this could have removed most of the higher mass particles in the suspension.

Alternatively the sonication may have been deficient. The sonicator available for this work was a Ney Ultrasonik 14h. Its 115 W output is substantially larger than the 12 W advocated by Ausman et al. [103], but the sonication power varies across the bath (also reported by Niyogi et al. [108] with a similar bath). Sample jars always drifted from regions of higher power to areas where nanotube clumps were not even agitated. Some groups use test-tube carousels (see Section 3.6.3) which ensure the sample always experiences a consistent power. The fluid height also had an effect; substantially higher agitation was achieved when the level in the jar matched the surrounding bath. If the jar was held by a clamp then the ultrasound was dissipated through the clamp. The only reliable method was holding the jar by hand, via a lump of Blu-Tack to prevent the ultrasound being transferred to the hand tissue. This also placed a practical limit on the length of time for which a sample could be sonicated.

The technique used to transfer nanotubes from suspension to the chip has a substantial influence on the quality of the dispersion, and it is difficult to ensure repeatability. Spin-coating gives uniform dispersions, but using SDS or nanotubes with the group's wafer spinners was not permitted due to the risk of contaminating other materials. Leaving fluid to dry on the chip resulted in very patchy distributions with surface densities not proportional to the suspension density, useless for device fabrication or studying methodology. Removing the suspension droplet before it fully dried yielded very few tubes on the chip, contrary to the results of Nygård [7]; SDS-coated nanotubes may preferentially remain in suspension rather than precipitate onto the surface.

A solution was to coat the substrate in photo-sensitive resist, and to expose and
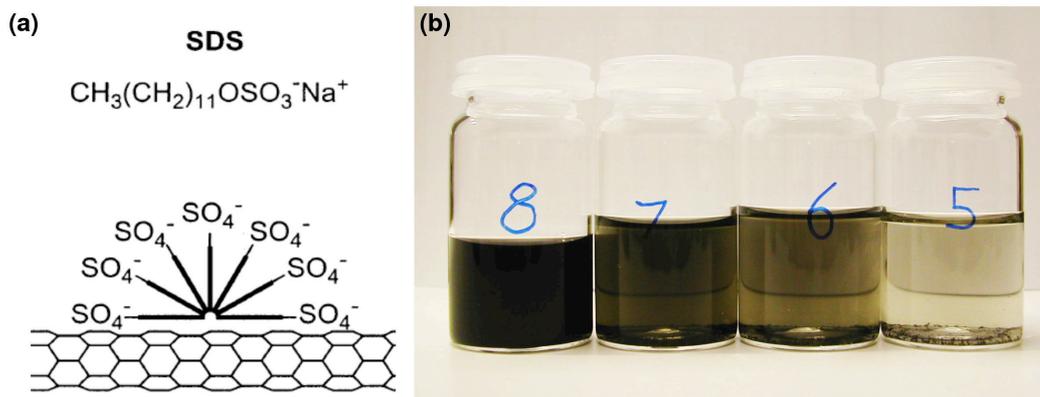
Figure 3.5: Nanotube suspensions in aqueous SDS. (a) SDS molecules bonding to the surface of a nanotube (from [104]). (b) Suspensions one week after preparation, with the same quantities of SDS, and with SWNTs in the ratio (left to right) 10:5:2.5:1.

develop an array of windows in the resist. Approximately 1 μl of suspension per mm^2 was pipetted onto the chip (large droplets cannot be held by surface tension and run off the edge of the chip). As the fluid on the surface dried, small droplets were pinned by surface tension in each of the pits defined by the windows, and all the nanotubes within deposited there. The SDS dried as a several nanometre thick crystal layer which was removed, along with the photoresist mask, by 2 minutes of sonication in IPA. Subsequent sonication even in acetone did not appear to remove any of the deposited nanotubes, which are very strongly bonded [7], nor the residual SDS. This is discussed in more detail in Section 3.3.

If the pits were too large then the droplets, and hence nanotubes, dried only into the edges (see Figure 3.6a & b). Circles with diameter of 100 μm gave a reliably even nanotube density, and matched the constraints of the device pattern to be used. The clean surface outside the pits eliminated any risk of short-circuits from stray nanotubes and soot away from the device area, but dispersion results still varied between pits at the centre and the edge of the chip, and residual SDS coating the nanotubes and surface made accurate quantitative assessment very difficult. Various parametric studies were attempted, varying the relative quantities of SWNTs and SDS, taking suspension samples from different heights in a settled sample jar, and comparing sonication times, but only weak conclusions were possible.

Suspensions with varying quantities of SWNTs were left to settle, and then the fluid decanted and the solid residue dried. This removed excess SDS and the mean
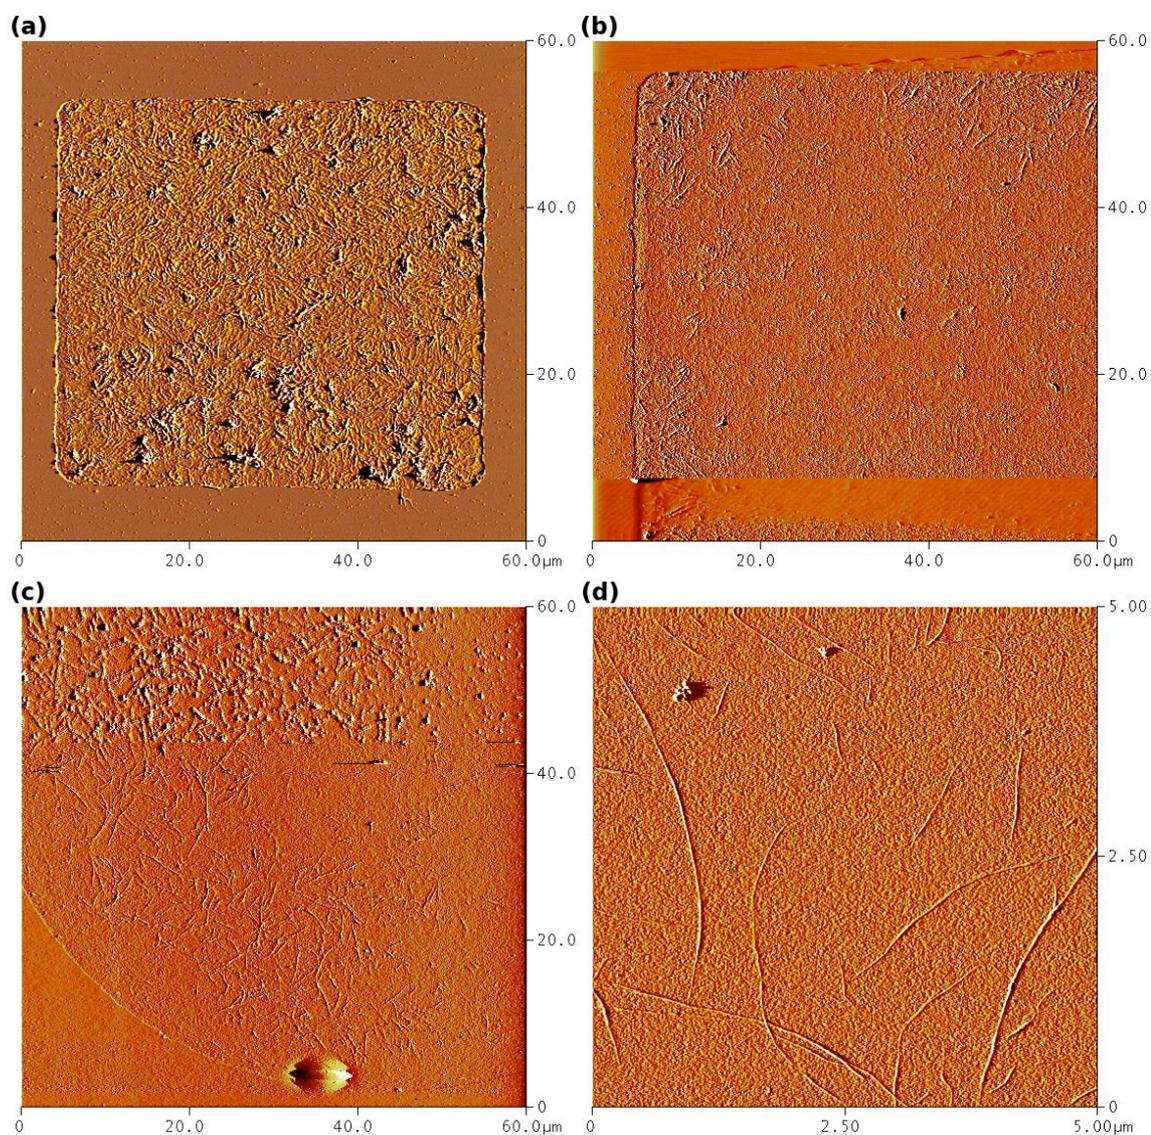
Figure 3.6: Nanotube dispersions from aqueous SDS. The AFM signal is amplitude shift; the colour scales are arbitrary. (a) High density, even dispersion across a 50μm wide square pit. (b) When the size is increased to 250 μm the tubes are drawn to the edges of the pit, and are sparse in the centre. (c) A more appropriate dispersion in a circular pit; 0.01 wt% SWNTs, 0.20 wt% SDS. (d) Mostly single nanotubes, but still at too high a density for device fabrication; 0.01 wt% SWNTs, 0.25 wt% SDS.

35

residual mass found to be 3.4±0.6 times the mass of the original nanotube material. The SWNTs and SDS were therefore bonding in a ratio ∼1:2.5, substantially below the 1:5-1:10 range established by Islam et al. [104]. However when increasing the amount of SDS in suspensions with 0.01 wt% SWNTs, they reached a maximum blackness at 0.25 wt% SDS at a ratio of 1:25. As dispersed, this was the optimum recipe observed, with 5 minutes sonication (the short time was necessary with the jar being held by hand). The suspension were stable over several days (see Figure 3.5), compared with chloroform, where the soot re-agglomerated immediately after sonication. At best a good surface density was achieved (see Figure 3.6c & d) but very few, if any isolated tubes. Nanotube/rope lengths were all 1-1.5μm, and though a couple of diameters were measured under 2 nm (expected for a single tube plus attached SDS), none pervaded the whole length; most were nearer 5 nm.

3.2.5 Dichloroethane

In comparison, dichloroethane suspensions gave immediately better results. When the same quantity of SWNTs were added to each of dichloroethane and aqueous SDS the former gave a denser and more uniform dispersion with a greater frequency of narrow diameters, from only 1 minute of sonication (compare Figures 3.6d and 3.7a). The nanotubes are also cleaner and longer. Diluting the suspension caused large flecks of nanotubes to precipitate out and no isolated tubes were then seen in on the chip. A similar effect was seen transferring suspensions between vessels, and with SDS suspensions, though dilution with a matched SDS solution reduced the amount of precipitation. It was therefore desirable to minimise the number of stages in preparing the suspension, and the contact time between the suspension and substrate.

Dipping the chip briefly into the suspension gave an improved dispersion (Figure 3.7b) and diameters measured were less than 1 nm over lengths of 700–1000 nm; these are single tubes long enough for use in devices. The chips were dipped 'face-down', horizontally - dipping them edge on gave extremely few dispersed nanotubes, possibly due to a fluid-flow effect. Brief (2 minutes) sonication in IPA immediately after dipping removed larger bundles and specks but not single nanotubes (unlike acetone, which did). The chip was then immediately blown dry with nitrogen (Figure 3.7c). Reducing the amount of SWNT material used to the smallest speck possible with tweezers (<1 mg) gave ideal dispersions, sufficiently uniform across the whole chip (Figure 3.7d), and if there were two few nanotubes the chip could be dipped repeatedly. The suspensions gave consistent dispersions after at least 3 days.
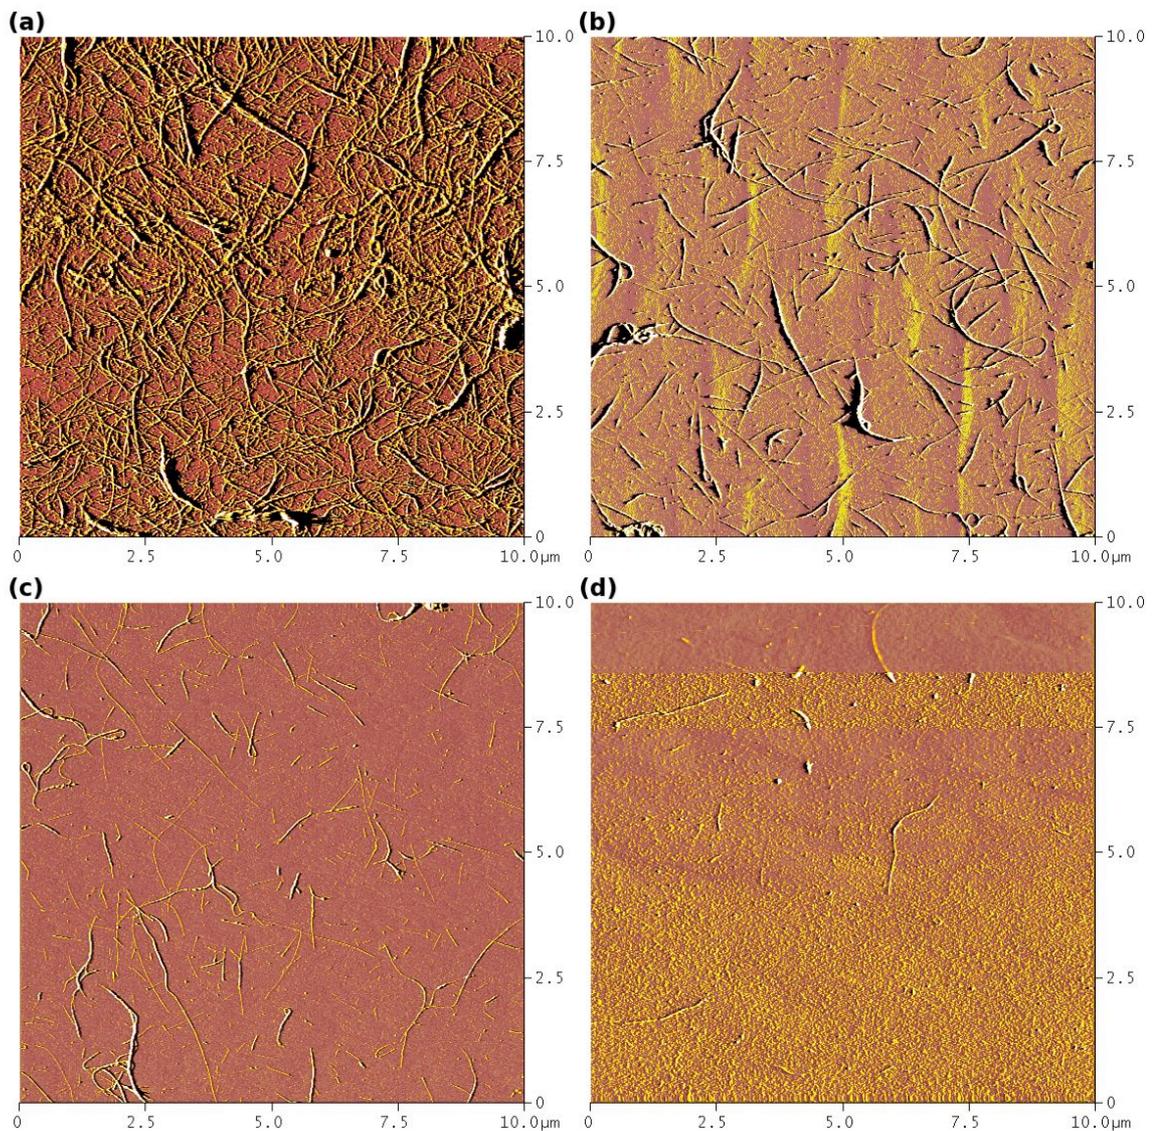
Figure 3.7: Nanotube dispersions from dichloroethane. The AFM signal is amplitude shift; the colour scales are arbitrary. (a) High density dispersion with 0.01 wt% SWNTs. (b) Improved results from dipping the chip. (c) Better dispersions after cleaning in IPA. (d) Minimal quantities of SWNTs gave ideal dispersions.

The best-performing suspension (<1 mg nanotubes sonicated in 10 ml dichloro-ethane for 1 minute, the chipped dipped horizontally then sonicated for 2 minutes in acetone and blown dry) was applied to a chip already part-processed for device fabri-cation (see Section 3.4.1 and Figure A.3). The results were disappointing, with a far higher ratio of ropes and clumps to isolated nanotubes than on the test chip. The surface was also covered in a few nanometres of patchy and rough PMMA. Though not anticipated, this is normal for such fabrication techniques, and not a problem for typical devices where the chip surface is of little concern. Soaking another processed chip in heated Posistrip cleaned the metal well but did not return the GaAs surface to a pristine state, and dispersing nanotubes on this chip did not yield improved results.

As a test a single dispersion was applied to GaAs, quartz and silicon using the same process. Each chip showed qualitatively different results with regard to amounts of isolated tubes, ropes, and large bundles. The SDS and dichloroethane processes developed above were also used with MWNTs on GaAs and achieved very poor re-sults, and with nanotube peapods, of which none were found on the substrate. These anecdotes hint that the nanotube material and target surface also have a determining role. The success of these processes with SWNTs on clean GaAs, contrary to some of the literature, suggests that not all the determining factors are established and it is then not surprising that they should fail with GaAs that was etched and dirty.

Though not ideal, the dispersions achieved were potentially adequate and repeating the testing process using processed wafers was not practicable. The remaining device chips were bathed in Posistrip and nanotubes dispersed.

3.3 Manipulating nanotubes

Since carbon nanotubes cannot presently be grown on GaAs and dispersion from sus-pension does not allow for precise placing or alignment, in situ manipulation of the nanotubes may be required for devices with specific geometric properties such as bend-induced tunnel barriers [113] or multiple tubes. The size and movement precision of an AFM tip are a good match for nanotube dimensions, making it an ideal apparatus for attempting manipulations. A description of AFM function is given in Section 4.1.2.

3.3.1 Background

During tapping mode imaging the AFM tip is only in instantaneous contact with the surface and no lateral forces are applied, therefore nanotubes on a surface are not

affected [7]. If the oscillation amplitude is reduced to be comparable to the diameter of a nanotube and feedback is turned off the SPM head will no longer raise to lift the tip over a nanotube, which it then strikes laterally. Postma et al. [114] have thus used a conventional AFM to bend, cut, move, cross and unravel nanotubes on silicon substrates, with oscillating amplitudes between 1 and 10 nm (a factor of 10–20 reduction from normal imaging). The force against the tube is enhanced by increasing the lateral speed, from 1–5 μm/s for imaging to 20–80 μm/s [84]. In this way feedback can be left on, partially reducing the wear to the tip. Although Lefebvre et al. [84] claim otherwise, contact mode manipulations also move nanotubes if the tip deflection and/or lateral speed are increased [83]. The Van der Waals forces bonding the tubes to the substrate surface are stronger than the rigidity along the length of the tube; when pushed with the tip the tubes will bend, sometimes substantially, before being universally translated. This makes the process extremely laborious and time-consuming. To translate or rotate a nanotube it must be nudged short distances in increments, at intervals along its length. Moving a nanotube by 1 μm can require thousands of steps [84], though creating a tunnel barrier only requires a single tug [115].

Cutting nanotubes to shorten them or remove a short-circuit requires violent force, likely to also translate the tube [83]. They can be cut more easily using local anodic oxidation (see Section 4.1.3), where a potential difference is applied between an SPM tip and the substrate. SWNTs have been cut with an STM in ultra-high vacuum [116] at voltages greater than $\pm$3.8 V. Since it is voltage dependent the process is attributed to injected electrons breaking bonds in the nanotube (compare with bond strengths of 3.6 eV for C–C and 6.3 eV for C=C). Park et al. [117] cut large diameter SWNTs above a threshold of -4 V (tip negative relative to sample only) using an AFM in air. Between -3 and -4 V they nick the tube, creating increased resistance and Coulomb barriers without cutting through it entirely. Kim et al. [90] cut MWNTs in contact mode AFM with a cutting voltage threshold dependent on the speed of the AFM movement. The lowest they record is -6 V at 100 nm/s - a 300 ms interaction time with a 30 nm diameter tube. There is no reported work on GaAs or related substrates.

3.3.2 Preliminary experiment

The first work was done with a Veeco Dimension 3100 AFM with a NanoScope III controller, and version 5 of the control software using custom lithography scripts. Scanning in contact mode with increased deflection and scan speed ($\sim$40 μm/s) overrode feedback, rubbing away dirt, residual SDS and nanotubes and exposing a cleaner GaAs
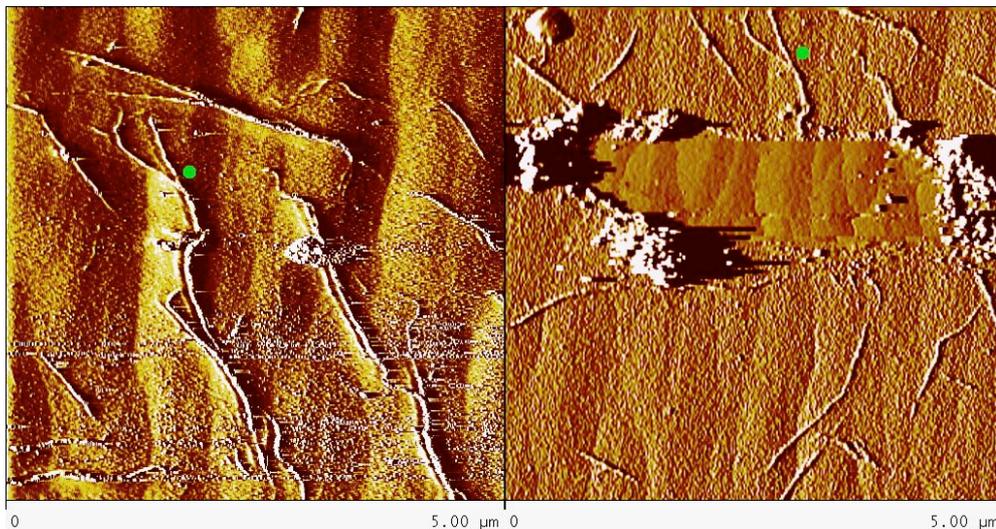
Figure 3.8: Cleaning a surface with AFM. After rapid contact mode sweeping residual SDS has been piled up and nanotubes sheared. There has been drift between the two images, shown by the green dot.

surface (Figure 3.8). The nanotubes were sheared at the edges of the swept area with no sign of the untouched sections being dragged. Moving the tip in single straight lines in tapping mode had no effect until the oscillation amplitude was reduced to 1% of the scanning value, which so close to the surface that the tip often crashed into the surface and sheared. On a quartz substrate a nanotube was easily nudged at several places along its length and a rope was separated into two single nanotubes. The same manipulations on GaAs caused the entire chip to move, with the target nanotube unaffected on the surface. With the chip fixed in place by vacuum suction nanotubes could be moved but the large force required always sheared the tip; only small nudges or cuts were possible. This AFM lacked positional feedback, therefore the path actually drawn could lie anywhere in the scan field, irrespective of the programmed coordinates, and the line length was subject to a 100% error. Coupled with the resistance of the nanotubes to movement this made even moderate translations on GaAs impossible.

3.3.3   Surface binding energy

Nanotubes on GaAs had shown resistance to cleaning, in comparison with quartz where they washed off with sonication in water. Their bonding to the surface was apparently stronger than the friction between the chip and the AFM plate, and than the AFM

tip. It was certainly not possible to repeat the manipulations described on silicon substrates above. All of this suggested an exceptionally strong surface attraction.

There have been isolated theoretical studies of the bonding strength of nanotubes on particular surfaces. All diameters given here are calculated with a nanotube modeller applet [118]. Using density functional theory Orellana et al. [119] predict the binding energy of a (6,6) SWNT (diameter $d = 8.2$Å) to be 0.21 eV/Å on Si(001), mediated via a C-Si chemical bond. With a molecular simulation using empirical force potentials Su et al. [72] calculate binding energies of 0.42 eV/Å for a (10,10) SWNT ($d = 13.7$Å) on Si(111) and 0.33 eV/Å on Si(100), and 0.34 eV/Å for a (18,0) SWNT ($d = 14.2$Å) on Si(100). Hertel et al. [75] deduce the force from observed compression of nanotubes (as they maximise surface contact) which at $\sim$0.3 eV for a (10,10) nanotube agrees with the model above. Kim et al. [120] model binding on InAs, a III-V material similar to GaAs, using a total energy method. They find it highly dependent on the substrate morphology, varying from 0.4 eV per binding site on (110) and (111) surfaces to 0.2 eV/site on ($\bar{1}\bar{1}\bar{1}$), which is As-exposed. Considering the (110) surface, which shows strongest binding, they find 0.3 eV/site for (10,0) ($d = 7.8$Å), (17,0) ($d = 13.3$Å) and (18,0) ($d = 14.2$Å) nanotubes.

Allowing for more than one binding site across the nanotube diameter since the tube-surface binding energy is roughly linear with diameter for small tubes [121], and converting all energies to per Å, tubes with similar diameters can be compared: a (6,6) SWNT on Si has 0.21 eV compared with 0.12 eV for a (10,0) on InAs, (10,10) 0.42 eV on Si cf. (17,0) 0.21 eV on InAs, (18,0) 0.34 eV on Si cf. (18,0) 0.23 eV on InAs. At 4.284 Å the lattice parameter of InAs is well matched with 4.26-4.32 Å for nanotubes, compared with 5.65Å for GaAs. Considering that the poorer lattice matching may reduce binding, the upper bound for binding energies on GaAs is half that on silicon. This contradicts the observations above, which could result instead from interactions between GaAs and the SDS coating the nanotubes. Were this true it might also explain the tendency of dichloroethane suspensions to disperse as connected and overlapping nanotubes on GaAs, but isolated nanotubes on other substrates; with GaAs the relative attraction to the substrate is less and nanotubes are more likely to leave suspension to bind to another nanotube.

3.3.4 Manipulations on GaAs

This hypothesis was tested using an improved NanoScope IV controller and an SPM head with 'closed loop' feedback, increasing the accuracy of programmed tip move-

ments to a couple of nanometres, so that small nudges could be done reliably. Version 6 of the control software included the NanoMan module for lithography. The paths followed by the tip are all straight lines perpendicular to the nanotube.

Contact mode lines with feedback had no effect on the tube up to a deflection of 1 V, but did remove some surface debris. The tips used are designed for non-contact use and relatively stiff with a force constant of 42N/m, so the vertical force is around 5 μN. With feedback plus a -4.5 V bias applied to the tip the nanotube was cut neatly (the two new ends were not moved substantially) and all debris removed along the line (Figure 3.9a). Without feedback or bias but instead pushing the tip 100 nm towards the surface also cut the nanotube and cleared debris, but this time the new ends were pulled out (Figure 3.9a) showing that the tube was kinked before it snapped. This was with a small deflection of 500-800 mV (4 μN force) which is within the imaging regime (a minimum deflection of 200-500 mV is needed to maintain surface contact).

By moving the tip only a limited distance beyond the nanotube it could be kinked for 50 nm but by 100 nm had snapped (Figure 3.9b). To test the effect of SDS another chip was soaked for 24 hours in still, de-ionised (DI) water, which revealed a previously overlooked problem using SDS. The regions of GaAs exposed to SDS solutions were always found to be $\sim$5 nm lower than the surrounding substrate. After this long soak, the floor of the exposed squares had dropped 100 nm. Still DI water absorbs carbon dioxide, creating acid ions which can etch GaAs. However the areas outside the squares appear untouched, or if etched the rate is far slower. The SDS solutions have a moderately high pH and GaAs is selectively etched by bases [122], so it is probable that some etching occurs while the suspension droplets dry, and again due to the residual SDS coating when the chip is sat in still water. The rate is slow at around 5 nm/hour, indicating that it is a selective process removing only certain atomic species and may be the cause of the increased roughness. After the water soak no nanotubes remained but their shadows were visible in the etch profile.

The previous chip was rinsed in a DI weir for 20 minutes; there was no visible change and manipulations gave the same results. After annealing at 320℃ for 20 minutes the surface was cleaner and a nanotube could be pulled for 100 nm, not breaking until 200 nm. In comparison nanotubes from a dichloroethane suspension were kinked up to 400 nm, and did not snap even when pulled 500 nm (Figure 3.9c). This confirms that nanotube binding to GaAs is reasonably weak but enhanced by SDS.
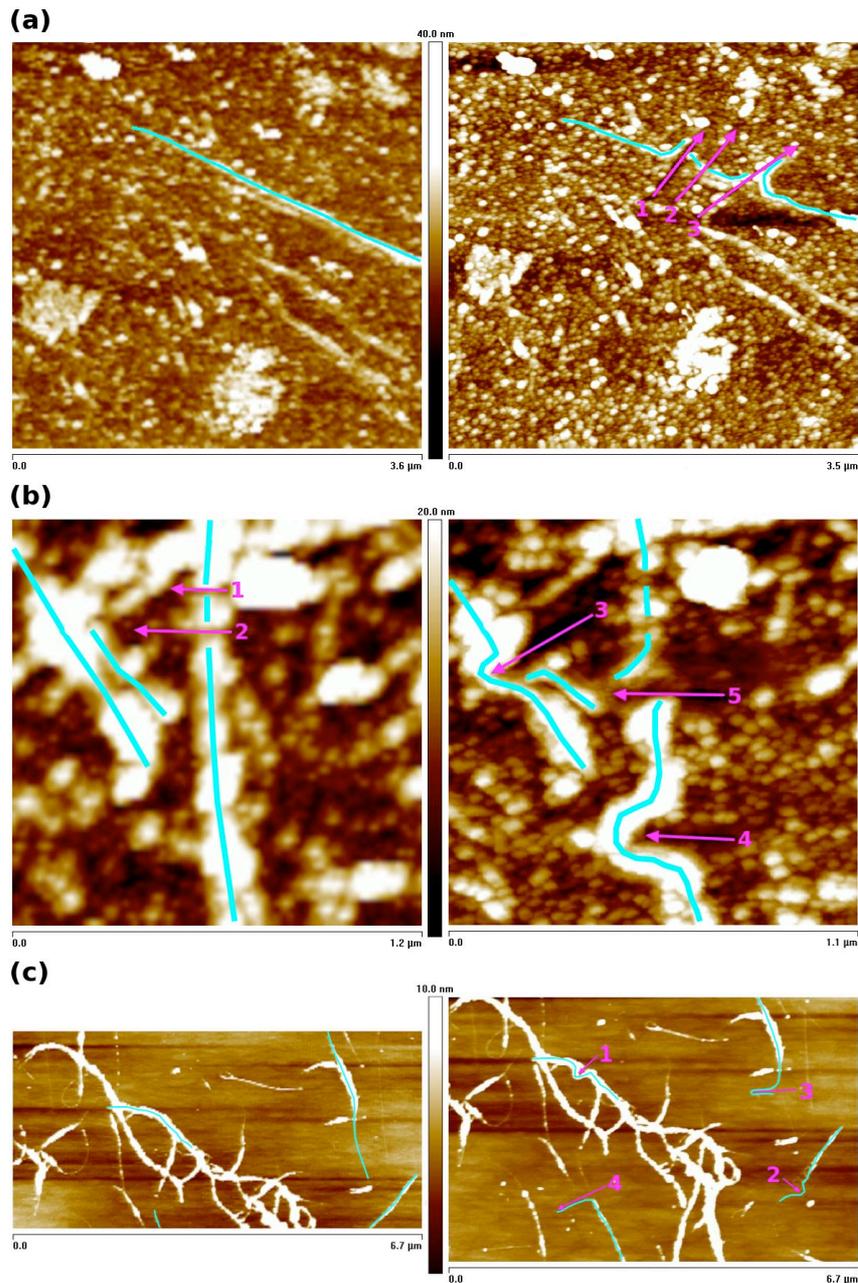
Figure 3.9: Nanotube manipulations on GaAs: AFM height images of nanotubes before (left) and after (right) linear manipulations by AFM tip. (a) Without AFM feedback an SDS-dispersed nanotube is dragged a short distance before it snaps (1 & 3). With a bias voltage applied to the tip the nanotube is cut neatly (2). (b) Bias voltages give neat cuts (1 & 2), a 50 nm nudge without bias kinks the nanotube (3 & 4) but a 100 nm push snaps it (5). (c) Successively longer pushes on dichloroethane-dispersed nanotubes do not snap them even up to 500 nm (1-4).

3.4 Device fabrication

3.4.1 Standard device fabrication

The Semiconductor Physics Group has well established processes and templates for GaAs device fabrication. Certain key features are present. The mesa is a limited region within which the 2DEG is confined, defined by etching through the dopant layer elsewhere on the chip surface leaving a raised plateau. Ohmic contacts are bond pads on the surface which make electrical contact to the 2DEG and surface gates are additional metal bond pads and finger wires which run over the surface of the mesa, either to gate the 2DEG or make contact to features such as nanotubes. The heterostructure wafer is chosen according to mobility and other requirements of the device being made, and the final chip is cut to fit and bonded into a standard package for mounting in the group's cryostats and fridges.

Features on the chip are defined using photo or electron beam lithography (Figure 3.10). In these processes a polymer resist is spun onto the chip surface, to give as even a thickness as possible, then baked to drive out the thinning solvent. The resist is exposed to the appropriate pattern, either to UV light through a shadow mask (photolithography) or to selective writing by an electron beam. Long chain polymer molecules in the resist degrade by irradiation, increasing solubility in exposed areas for removal by a chemical developer. This leaves the chip surface exposed for etching or the deposition of metal. Where appropriate the chip is cleaned to remove contaminant particles or residual resist using acetone, being a strong solvent, followed by IPA, which can be blown away before it dries, leaving no residues. If there are no fragile features such as sub-micron gates then washing is enhanced in an ultrasound bath.

Fabrication is carried out in a clean room, to limit dirt contaminants on the chip surface. Several devices are prepared on a single chip for consistency and efficiency. The chip is cut from a wafer by scratching the surface using a diamond-tipped scriber. A small force either side of the scratch then cleaves the chip.

For the mesa, Shipley Microposit 1813 resist is spun on at 5500 rpm for 30 seconds, and baked for 1 minute at 115℃. It is exposed to UV light for 9 seconds using a Karl Süss mask aligner, and developed in MF319 for 30 seconds. The developer is removed with de-ionised water. Prior to etching, the GaAs surface oxide is stripped with hydrochloric acid ($1HCl:9H_2O$) for 10 seconds. The chip is etched for around 15 seconds in a sulphuric acid/hydrogen peroxide mixture ($1H_2SO_4:8H_2O_2:91H_2O$) which etches at approximately 10 nm per second.
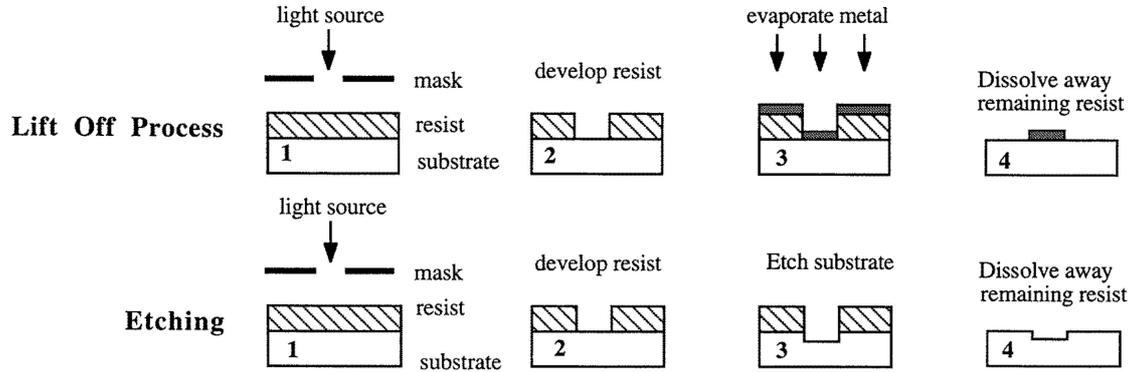
Figure 3.10: Basic device fabrication steps. (Top) Metal evaporation and lift off. (Bottom) Substrate etching (from [2]).

For adding the ohmic contacts the resist bake is at 90℃ and is followed by a 5 minute soak in chlorobenzene. This hardens the surface so as to give an overhang profile after developing, improving the reliability of metal lift-off and the metal profile. After oxide removal approximately 150 nm thickness of a gold/germanium/nickel alloy is evaporated in a low pressure chamber. The chip is soaked overnight in acetone to remove the remaining resist, and lift-off the unwanted metal with it. The ohmic pads are annealed at 220℃ for 15 seconds, 430℃ for 80 seconds, and a further 220℃ for 15 seconds in a nitrogen/hydrogen forming gas. Surface gates are made as per ohmic contacts but without the annealing step and usually titanium followed by gold is evaporated. The titanium wets the surface ensuring the contacts stick reliably and the gold provides good electrical properties and resistance to corrosion.

The chip is fixed into an LCC (leadless chip carrier) using G-varnish, and a Karl Süss thermosonic wire bonder is used to made electrical connection, with gold wires bonded from pads on the chip to pads on the LCC.

3.4.2 Integrating carbon nanotubes

Standard fabrication procedures were followed as far as possible to ensure reliability and compatibility. The Group's measurement rigs and probes all have 20 connections and a pre-existing mesa pattern with 20 each of ohmic contacts and gates was chosen, maximising the possible number of contacted nanotubes on each chip. The mesa was an open square area 100 μm wide, large enough to contain many contacts and gates and small enough to be scanned efficiently by AFM at high resolution. The mesa, ohmic contacts and gates with fingers which ran up onto the mesa, overlapping by around

10 μm, were all defined by optical lithography (see Figures A.3 and A.4). One hundred mesa devices were completed, spread over four large chips. The GaAs wafer used was grown at 560℃ with a moderately low As flux, and one of the flattest measured.

To give accurate coordinates for the nanotubes for subsequent contacting the mesas were patterned with e-beam written alignment marks. A new pattern was designed covering an 80 μm wide square (see Figure A.5). It is a non-repeating grid of 19 different cross-based marks. Each mark has two axes, to enable orientation and centering, and they are arranged such that any four in a square identify a unique location. If the orientation of the chip is known then in most cases only two marks are required. The pattern also includes a number unique to each mesa, printed once in a corner of the chip large enough to be identified with a magnifying glass, and once at the corner of the mesa, resolvable by the AFM optical camera. The cross marks are 1 μm wide and also visible in the AFM camera for aligning the tip with the scanning target. Though there is a 10-20 μm error between the motor alignment of the AFM tip and the actual area scanned a single scan is enough to identify the tip location, at a resolution high enough to also image the nanotubes.

The nanotubes were dispersed as late as possible in the fabrication process, reducing the risk of damage from cleaning, lift-off or e-beam exposure (see Section 3.4.4). Once the alignment marks were complete the chips were washed in Posistrip for 20 minutes at 70℃ followed by a DI water rinse and cut into smaller units of 4 mesas for dipping in the nanotube dispersions.

From this point onwards the chips were subjected to no sonication or aggressive cleaning, to prevent damage to the nanotubes. Once the nanotubes were contacted the chips were always handled with rubber tweezers to prevent an electrostatic discharge from the body blowing them up. For the same reason care was taken when bonding, with a constant de-ionising air flow, and the ball-forming spark applied manually and at some distance from the chip. Every effort was made to keep appropriate contacts grounded during electrical measurements.

3.4.3　Device design

Four mesas were mapped for tubes with 12 μm wide AFM scans including an alignment mark in each corner, at 1 Hz and 512×512 pixels. From 200 scans five to ten nanotubes were selected for devices on each mesa. Most were long straight tubes, easily identified as single tubes by narrow diameters. Some crossed and kinked tubes were also chosen to enable a variety of device functions, for a total of 29 devices.

The position of each nanotube was plotted relative to the alignment marks, and contacts designed around them. The e-beam writer aligns automatically to the alignment mark pattern, with an error of a few tens of nanometres. The centre of an alignment mark can be identified to within around 20 nm and a measured distance on a scan is accurate to within 40 nm, being a single pixel error at each end. There is a non-orthogonal scanning error of around 1% (100 nm), which can mostly be corrected for. Therefore the error in matching the contacts to the nanotubes is of order 100 nm in total, which is accommodated in the overlap (see Figure 3.22). The contacts also have a large overlap with the optically-defined gates.

The primary feature of each device is a gap between two nanotube contacts which might be used to define a 1D channel in the 2DEG. Twelve such devices had wide contacts, creating channels at least a micron long, and five were shorter. The channel widths (and hence contacted nanotube lengths) were 700 to 1900 nm. Two nanotubes were long enough to include side gates a short distance away, and three were divided into two sections by a third narrow contact. Two devices had crossed nanotubes, each contacted at both ends, and a further two had one tube just overlapping another, such that the first tube was only contacted at one end, with a side gate beside the junction. Two devices had kinked tubes, and one had two parallel tubes very close together and so sharing contacts, with a side gate to try to differentiate between transport in each. See Figure A.5 for a sample mesa, and Section A.2 for the specifications of all the devices.

The mesas were separated and bonded into packages individually. Since complete bonding required more than 20 nanotube and ohmic contacts on each mesa chip only 19 devices could be bonded initially. Two devices were lost when bonds lifted and removed the bond-pad metal.

3.4.4 E-beam and carbon nanotubes

There is increasing evidence that electron beam exposure damages or alters carbon nanotubes at surprisingly low energies [44, 123]. The threshold for removing carbon atoms by knock-on collision is high, at 86 kV [124], but Suzuki and Kobayashi [125] find that 20 kV, and Nygård [7] that 10 kV is sufficient to induce defects, reducing and eventually destroying conductivity. Even as low as 1 kV, 60 seconds irradiation at a dose of $5 \times 10^{14} \text{cm}^{-1}$ induced a metallic-semiconducting transition (the resulting conduction was ambipolar) [126]. Suzuki et al. [127] found that this could be reversed with annealing or treatment in UHV, and observed with Raman spectroscopy that

e-beam exposure extinguished the optical property which derives from the tube's 1D structure, an effect which peaked around 1 kV with doses equivalent to a few seconds of imaging.

1-2 kV is the most reliable region for imaging fine SWNTs, therefore even though SEM imaging is faster than AFM, it was employed here only after a device had been measured, and where access with an AFM tip was not possible due to bond wires or protruding gates. Also, an AFM is more reliable for imaging the finest tubes and gives information on their diameter, while the SEM deposits a layer of amorphous carbon, significant even within the time to take a single image. It is likely that during e-beam lithography the PMMA will protect the tubes from damage, and nanotubes were even imaged through plasma-hardened PMMA (Figure 3.20), which might be used as a protective layer. A comparatively low lithography dosage of 40kV was nevertheless used. The writer is a Leica VB6 UHR EWF Electron Beam Lithography System.

3.4.5 Metal contacts

The metal used to make contact to the nanotubes has implications for the device performance. Gold is the most frequently used material in early nanotube work [7, 62, 78, 91, 95] as it is inert, easy to evaporate and also reliable as pads for ball-bonding. A thin layer of titanium or chromium underneath wets the surface, ensuring good adhesion and reliable lift-off [78, 80, 89, 115, 128–130]. The gold top layer also protects titanium from oxidation. Contacts evaporated on top of the tube give lower contact resistances and more ohmic behaviour than tubes dispersed on top of contacts [49, 77]. Contamination occurring between the nanotube dispersion and contact evaporation processes reduces conductance [7], though Mann et al. [131] find that transport only occurs at the very edge of the metal contact, and the length of tube beneath the electrodes is electrically 'turned off'. This is another indicator that the e-beam lithography will not affect nanotube conductivity significantly.

Schottky barriers exist at the metal-nanotube interface [132]. When the work function of the source electrode metal is larger than that of the nanotube, the Fermi level of the electrode is closer to the nanotube's valence-band edge and the barrier height for hole transport is low. Therefore holes are readily injected and the nanotube exhibits p-type behaviour [133]. Conversely, if the metal has a small work function and a Fermi level closer to the tube's conduction band electrons will be injected, giving n-type behaviour (see Figure 3.11). A nanotube work function is around 4.8 eV [133], compared with 5.1 eV for gold and 4.3 eV for titanium. Titanium therefore gives
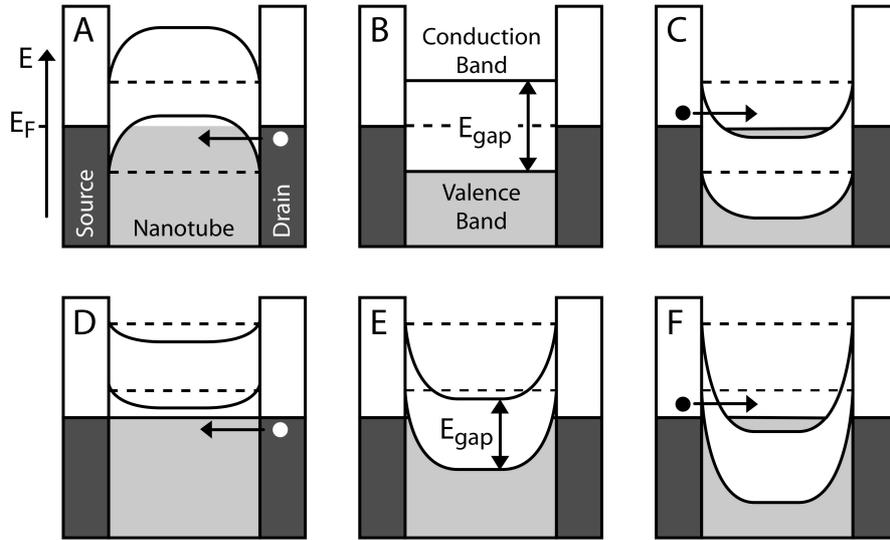
Figure 3.11: Band diagrams for nanotubes contacted by different metals.
(A-C) Metal with a similar work function to the carbon nanotube: conduction occurs
via Schottky barriers when gating the bulk of the nanotube p-type (A) or n-type
(C). (D-F) Metal with a higher work function: p-type conduction (D) is ohmic, but
n-type conduction (F) occurs via a large Schottky barrier. Movement of charge for
a small positive bias on the source is indicated (from [10]).

quite significant Schottky barriers [109] and large contact resistances, and gold less so.
Annealing the contacts under inert gas at temperatures from 180°C [92] to 600°C [78]
reduces the contact resistance and can achieve more ambipolar behaviour [129].

Palladium has a higher work function (5.1eV) and good wetting interactions [132].
It gives more reliable contacts than titanium, rarely [131] if ever [134] requiring post-
annealing to make them ohmic, and raising the channel conductance closer to the $4e^2/h$
limit for metallic tubes [132]. Platinum has an even higher work function of 5.7 eV
and is commonly used [50, 82], but gives lower p-channel conductance even than gold
[92], non-ohmic contact [131], and non-metallic behaviour for semiconducting tubes in
their 'ON' state [132]. This is thought to be due to poor sticking or wetting. Further
alternatives are bare titanium [106], nickel [132, 135] and aluminium [47].

Previous nanotube work in the group [10] had been hampered by low device yield,
therefore in the first instance palladium was chosen for this work to ensure the largest
number of working contacts. Since palladium is soft and not robust for wire bonding or
probing [132] it was only used for fine features on the mesa, and at the edge overlapped
with optically defined gold-on-titanium contact arms leading to large bond pads. The

alignment marks had been around 80 nm high, which gave too sharp a profile for the AFM tip and feedback leading to shadowing of the image and some blunting of the tip. The contacts were therefore evaporated to less than 45 nm. They were not annealed.

3.5 Transport measurements

The primary aim was to achieve mutual gating between the GaAs 2DEG and nanotubes; subsequent to that was measurement of gate-controlled transport in a nanotube quantum dot and then remote sensing of charge movement through the dot using a point contact. The Fermi energy of electrons in the 2DEG is around 6 meV, corresponding to a temperature of 70 K. The energy scale of nanotube quantum dots is $\delta E = h v_F / L$, giving a temperature equivalent of $T_{\delta E} \approx 40$ K for a 1 μm dot [49]. Experiments are carried out below these temperatures to prevent thermal smearing obscuring transport features, for which liquid helium temperature (4.2 K) is adequate.

For measuring, the chip was fixed in a Charntec sample holder on the end of a dipstick and submerged in a bath of liquid He^4 in a Dewar flask. At the other end a breakout box had BNC coaxial connectors for electrical measurement and input to the chip contacts. For most DC voltage source and measurement purposes Keithley 2400 Source Measurement Units (SMUs) were used. These received instructions from and returned data to an Acorn PC running the CryoMeas software (written by Dr C. J. B. Ford), and later a version adapted for LabView. The temperature was deduced by measuring current through the 2DEG as the dipstick was lowered; once the device was submerged in liquid helium the 2DEG resistance plateaued at 5–10 kΩ.

3.5.1 First measurements

The gate dependence of nanotubes at room temperature identifies their conduction type (no gate dependence indicates metallic) [136], but such assessment is not possible here because GaAs is moderately surface conducting at room temperature. When applying a voltage to the 2DEG at low temperatures the conduction electrons break through to the chip surface at around -1 V. This limits the gating range available from the 2DEG. All devices were tested at 4.2 K with back gate range -1 to $\sim$10 V and DC bias $\pm$100 mV, or until conduction was observed. SWNTs tolerate a maximum current of around 25 μA [132], which with perfect contact transmission would be driven by a 160 mV bias. All devices on two mesas were found to be short-circuited to the back gate, and on the other two only 3 out of 11 devices showed any conduction. With
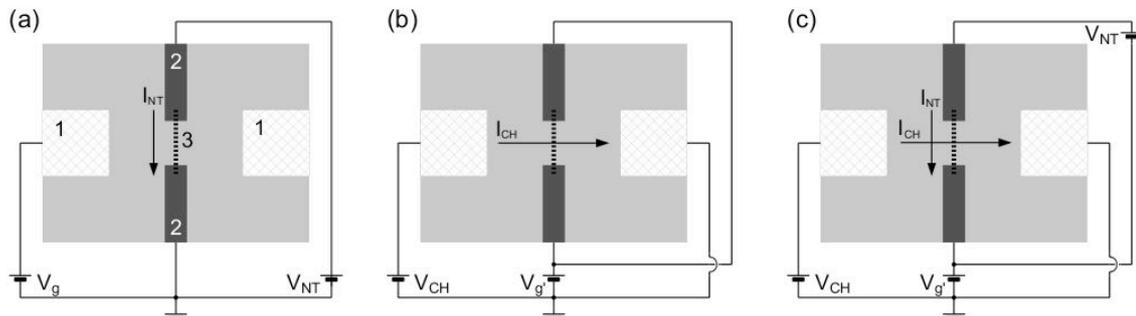
Figure 3.12: Nanotube measurement set-up. Circuit diagrams with ohmic contacts (1), nanotube contacts (2) and a nanotube (3). (a) To measure I_{NT} through the nanotube V_{NT} is applied to one nanotube contact with the other grounded, and V_g to one ohmic contact. (b) To measure I_{CH} through the channel V_{CH} is applied to one ohmic contact with another grounded, and $V_{g'}$ to both nanotube contacts. (c) To measure both currents simultaneously V_{CH} is applied to one ohmic contact with another grounded, and $V_{g'}$ to both nanotube contacts, with V_{NT} added on one.

hindsight much higher bias voltages with an appropriate current compliance might have revealed conduction in more devices (see Section 3.5.3).

To measure nanotube conduction a DC bias voltage V_{NT} was applied to one nanotube contact and the other contact grounded. A DC gating voltage V_g was applied to one ohmic contact with the others floating (see Figure 3.12a). All other nanotube contacts on the chip were also floating, since if they were grounded they would deplete the 2DEG when negatively gated beneath. To measure channel conduction through the 2DEG beneath a nanotube, V_{CH} is applied to one ohmic contact and another located on the other side of the nanotube is grounded. Negative $V_{g'}$ applied to both nanotube contacts depletes the 2DEG beneath and defines the channel beneath the nanotube (see Figure 3.12b). A poor initial understanding of the equipment and expected transport behaviour meant that much of the data was poorly sampled and averaged, but all three devices stopped working prematurely, indicating damage to or loss of the nanotube, and could not be re-measured.

All three devices show some common features. Biasing behaviour is non-linear, without a clear semiconductor band gap (Figure 3.13a). The highest conductance observed is $0.15e^2/h$, indicating either very poor contact transparency, which is surprising for Pd contacts, or additional scattering centres along the nanotubes. In back gate sweeps the steep rise in current beginning around $V_g = -1$ V is the 2DEG breaking through to the surface, and all three show an enhancement of conduction just before

the 2DEG channel pinches off at $V_g \sim 0.8$ V, followed by a broadening and reduction in conduction peaks (compare Figures 3.13b & c).

Gate sweeps for device E2t8 show peaks and spikes on top of a slowly varying background. Slight smoothing reveals these peaks to be approximately periodic at $\Delta V_g = 23 \pm 3$ mV, which is a good indicator of Coulomb blockade. This implies a quantum dot with a gate capacitance, $C_g = e^2/\Delta V_g$, of 7 aF (see Figure 3.13c). At high bias the peaks are smoothed out as expected. Coulomb blockade peaks would have appeared sharper and more defined at lower temperatures but when the chip was transferred to a 1.5 K cryostat the device failed. Approximately every 7th peak is missing or strongly suppressed. E2t13 is divided into two sections, each showing repeatable, well-defined conduction peaks, but with less recorded evidence of periodicity. Both sections measured together in series show an asymmetry between positive and negative bias; some features are absent or even inverted, indicating a degree of diode behaviour . E7t4-3 does not show any explicit peaks.

The channel depletion profiles themselves are surprising. With conventional GaAs point contacts the regions under the gates deplete when the surface contacts are between 0.3 and 0.5 V lower than the 2DEG, defining the channel and giving a step in conductance. Pinch off is proportional to the channel width, and occurs after 3 V for a width of 700 nm [15, 137]. However E2t8 (1.5 μm wide) pinches-off at 0.8 V and E7t4-3 (0.7 μm) at 1.1 V. Both show only limited definition, as if this is washed out by the rapid pinch-off. Although one section of E2t13 (0.7 μm and 1.4 μm) is twice the width of the other, it pinches off only 0.05 V later at 0.45 V (Figure 3.14). There is a mere hint of definition just before pinch-off. The most likely cause of these early and apparently arbitrary pinch-offs is that nanotubes are acting as charged extensions of the metal gates, and are themselves gating the 2DEG. The voltage separation between definition and total depletion is a measure of this effect. Other devices which showed no nanotube transport did also show this early pinch-off, such as at 0.35 V with a 0.8 μm gap for E7t12. This could be because the nanotubes are present but broken, and so still extend the metal gates.

When using a 1D channel for charge sensing the channel is normally brought to the point of steepest conduction gradient just before pinch-off. This is the most sensitive region, and electrons moving in and out of a nearby quantum dot can cause a saw-tooth profile on top of it. However differential biasing of the contacts of one of the devices above showed no effect in the channel, beyond the expected change in width.
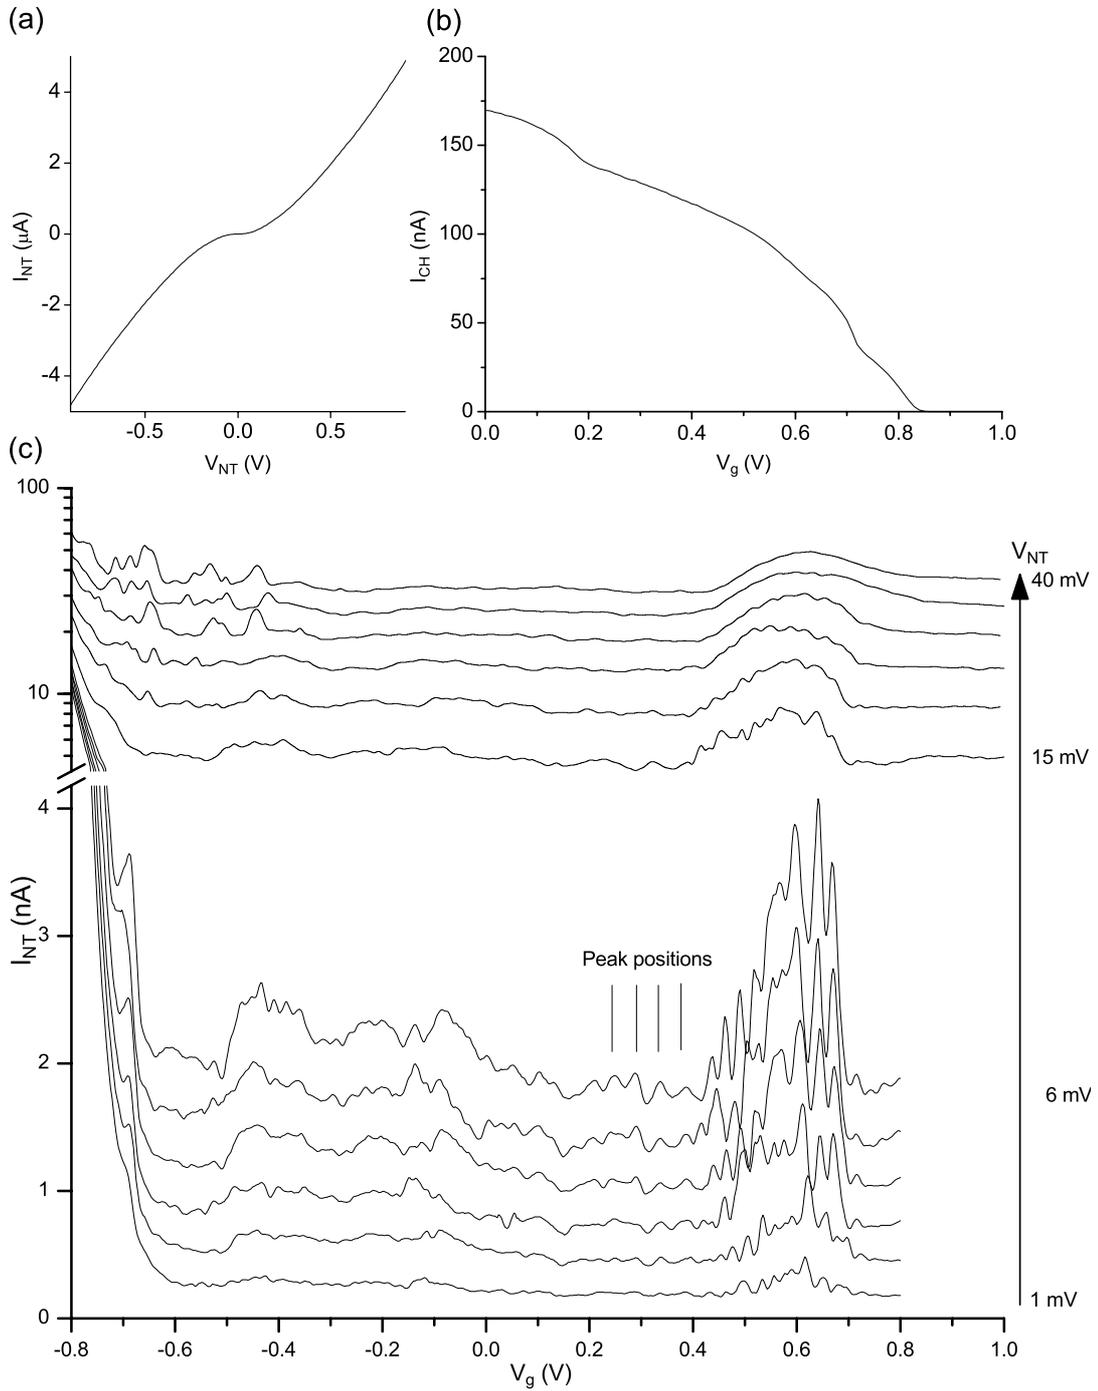
Figure 3.13: Coulomb oscillations from a nanotube. (a) Bias voltage sweep for device E7t4-3. (b) Channel current depletion profile for E2t8. (c) Gate voltage sweeps at different nanotube bias voltages for E2t8. The data shows periodic oscillations on top of a smooth background, which broaden after pinch-off.
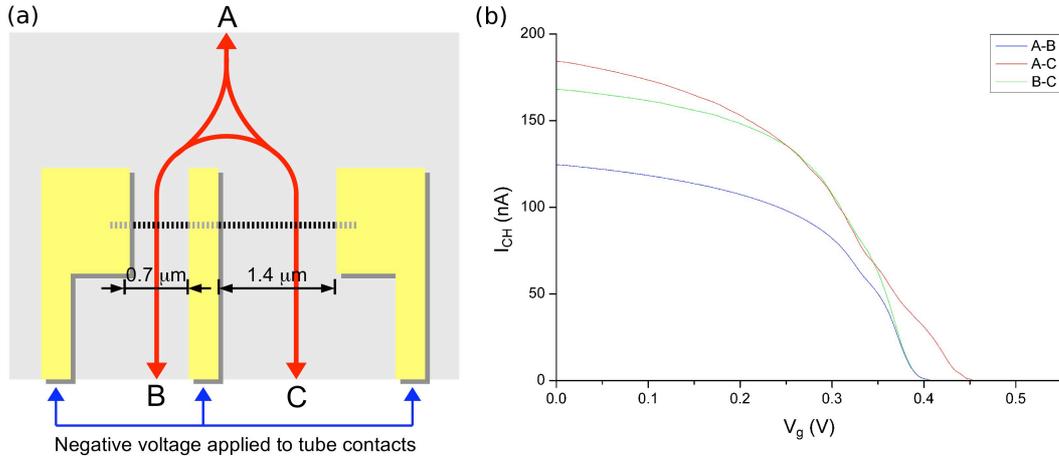
Figure 3.14: Channel pinch-off beneath nanotubes. (a) Diagram showing current paths through the 2DEG beneath device E2t13 (compare with Figure A.5). (b) Channel current depletion profiles for each path; the right-hand channel is twice as wide, so A–C would be expected to pinch-off at twice the gate voltage of A–B.

3.5.2 Device recovery and re-measurement

When examined optically contacts to 12 of the 19 bonded devices were lifted from the surface or broken, with adjacent damage to the GaAs. Non-bonded devices were unharmed so the damage was the result of an electrostatic discharge (ESD) event. Since it occurred with all four mesa chips, despite careful handling during measurement, the most likely cause is their storage; there were insufficient shielded chip cases so the devices had to be stored in polythene pots which are susceptible to charging. Under SEM the two chips which had been wholly short-circuiting had crystalline lumps of an insulating material clinging to metal features (Figure 3.15a). Over several months this spread across the chip surface and cracked apart (Figure 3.15c). EDX identified only Ga, As and O, and the nature of the material remains unknown. Also seen were substantial mats of nanotubes connecting ohmic and contact pads as well as short-circuiting devices (Figure 3.15b).

The bond wires were removed and all four chips demounted and washed in Shipley 1165 resist stripper at 70℃ for 30 minutes (found to be more effective than posistrip). This removed residual PMMA from the final lithography stage so that nanotubes could be resolved again by AFM. The nanotubes in seven of the surviving devices were either missing altogether or broken. E7t12 had a narrow break in the nanotube, supporting the conclusion that it is gating by nanotube that causes pinch-off soon after definition.
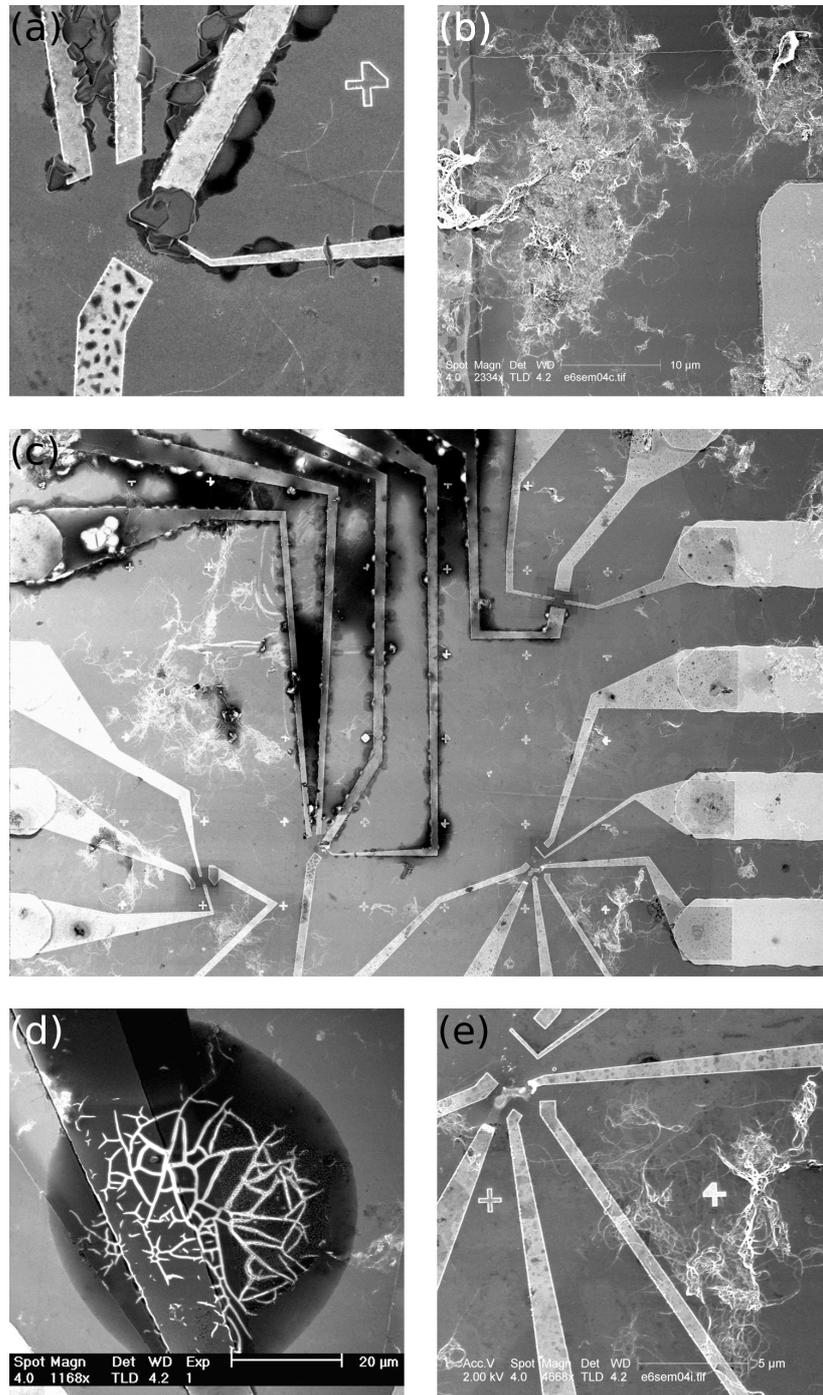
Figure 3.15: Contamination and nanotube mats. (a) Crystalline material clinging to contacts. (b) Nanotube mat connecting an ohmic contact (L) and a gate (R). (c) After several months the material has spread and some devices blown up. (d) Close-up of cracked material. (e) Close-up of destroyed and shorted device.

3. CARBON NANOTUBE DEVICES ON GALLIUM ARSENIDE

To recover more devices for measurement the chips were coated in photoresist and developed with a pattern leaving resist covering only the mesa, protecting the devices. The chips were exposed to an RF oxygen plasma etch at 100 W for four minutes [138], to burn away exposed nanotubes. Bonding was attempted to all devices with surviving contacts and bond pads, including those with tubes missing or damaged; four failed.

Almost all optically visible nanotube material was now gone and only two devices electrically connected to the back gate. However conduction was seen for only one device (E2t1-2). Chip E1 was plasma etched again for six minutes to remove all resist and nanotubes, after which the pinch-off for E1t5-2 (0.7 μm wide channel) shifted from 0.9 V before etching to 1.5 V after, and for E1t11-2 (1.2 μm) from 0.7 V to 3.2 V. The new values are roughly proportional to the channel widths, therefore the single broken nanotube observed across each gap can gate the 2DEG. That a nanotube can do so effectively suggests that detecting single electron charging in a nanotube will be possible with a GaAs point contact.

For E2t1-2 the nanotube and channel currents were measured simultaneously using the arrangement as for measuring channel currents with a second DC bias added onto one of the nanotube contacts (Figure 3.12c). The gate voltage is applied to the nanotube contacts therefore breakthrough occurs at positive and pinch-off at negative voltages, the inverse of the data already presented. There are a large number of conduction peaks across the gating range, which are less frequent after pinch-off at -0.4 V (Figure 3.16). They are repeatable, irregular and broaden to higher bias but some regions show periods of 27 mV and 85 mV (Figure 3.17). The current-voltage characteristics are non-linear and with no conduction gap, and subtracting a smooth background from the current (values taken at $V_{g'} = -0.32$ V - Figure 3.17a-inset) reveals irregular Coulomb diamonds (Figure 3.17b).

Zooming to a narrower region reveals that gate-dependent conduction features do not shift with changing bias to the same degree as the channel conduction. Pinch-off moves relative to $V_{g'}$ by a value equal to V_{NT} but nanotube conduction features move by only 80% of V_{NT} (Figure 3.18). However, certain features were found in the depletion profile which are neither symmetric in bias nor aligned with the changing nanotube bias. When compared with the nanotube conduction these are seen to align with bias-asymmetric features there. This is slight but promising evidence that conduction correlations between the nanotube and the 2DEG channel are feasible and may be observed (Figure 3.19). At this stage the pinch-off started shifting dramatically between sweeps, down to nearly -0.1 V. This could not be reversed by illuminating or warming the device, and made comparing further data impossible.
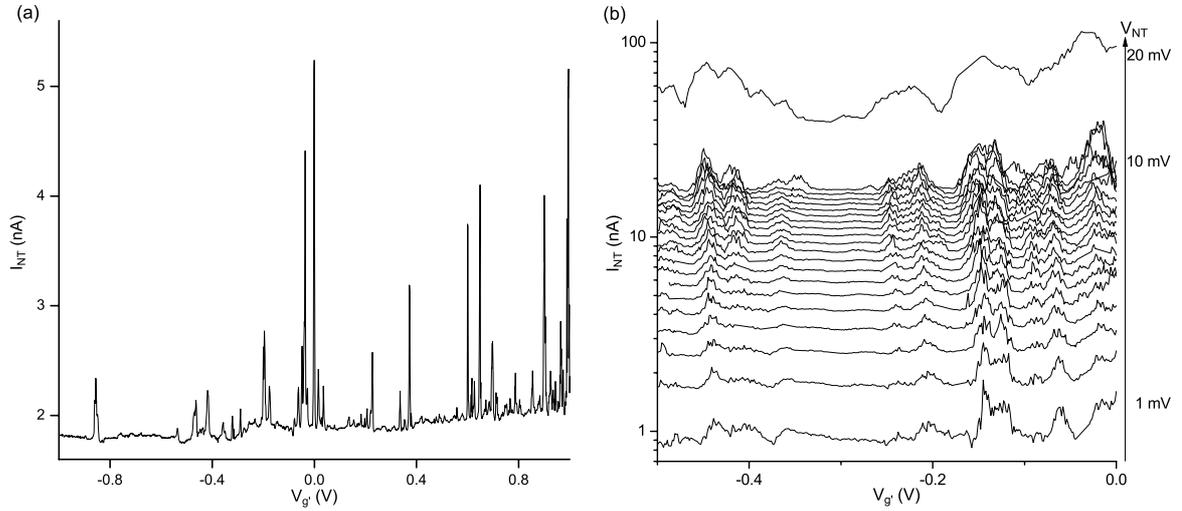
Figure 3.16: Device E2t1-2: Coulomb oscillations. (a) Broad gate sweep with $V_{NT} = 1$ mV. (b) Narrower gate sweeps at a range of bias voltages.
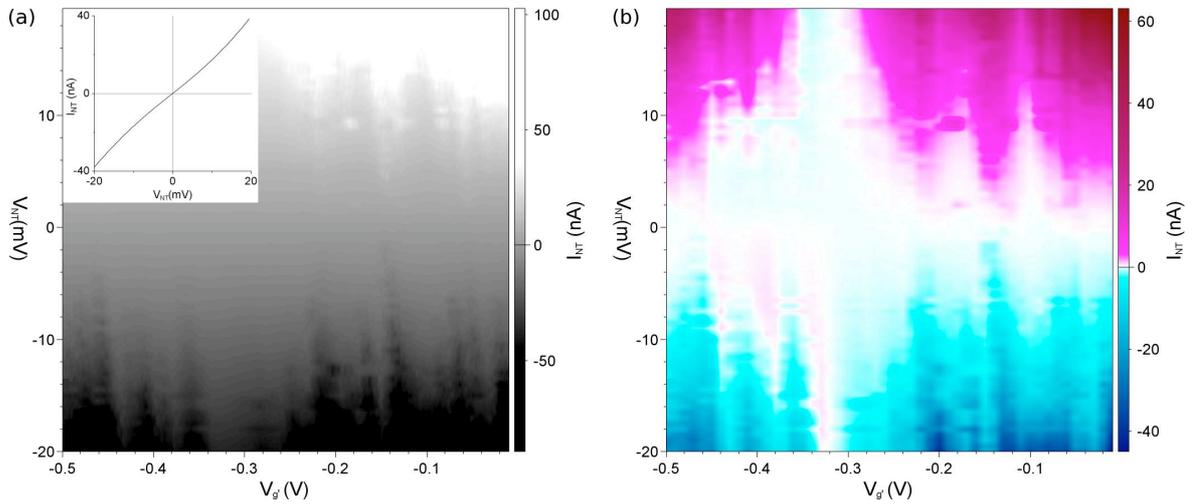
Figure 3.17: Device E2t1-2: Coulomb diamonds. (a) Surface plot of current against V_{NT} and $V_{g'}$ - there is limited periodicity on top of a smooth background (inset). (b) The data with the background subtracted, revealing Coulomb diamonds.
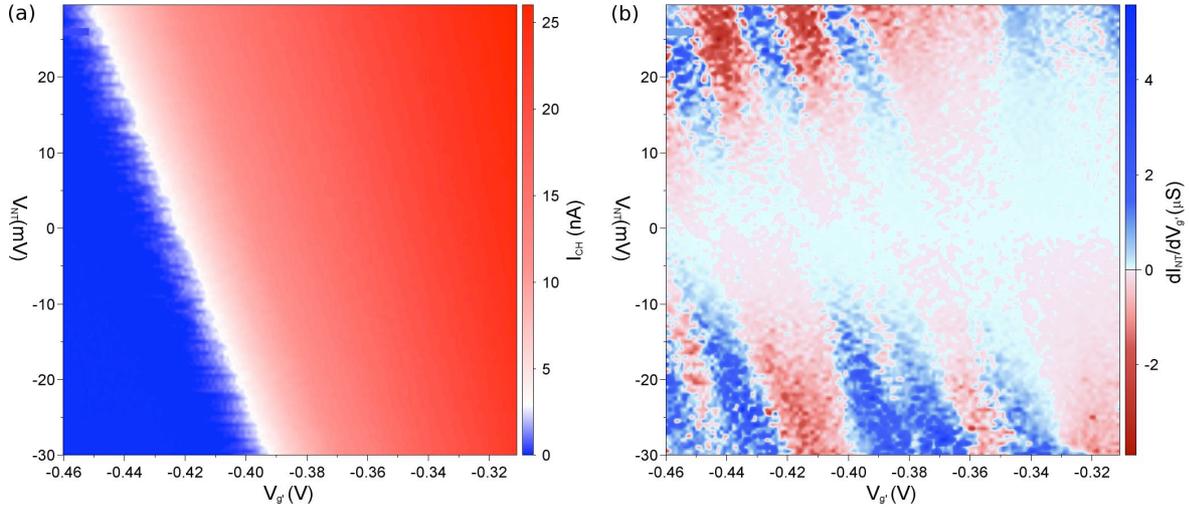
Figure 3.18: Device E2t1-2: Localised gating. (a) Surface plot of channel current profile (colour scale chosen to highlight pinch-off). (b) Surface plot of nanotube differential conduction (the differential plot highlights the Coulomb peaks).
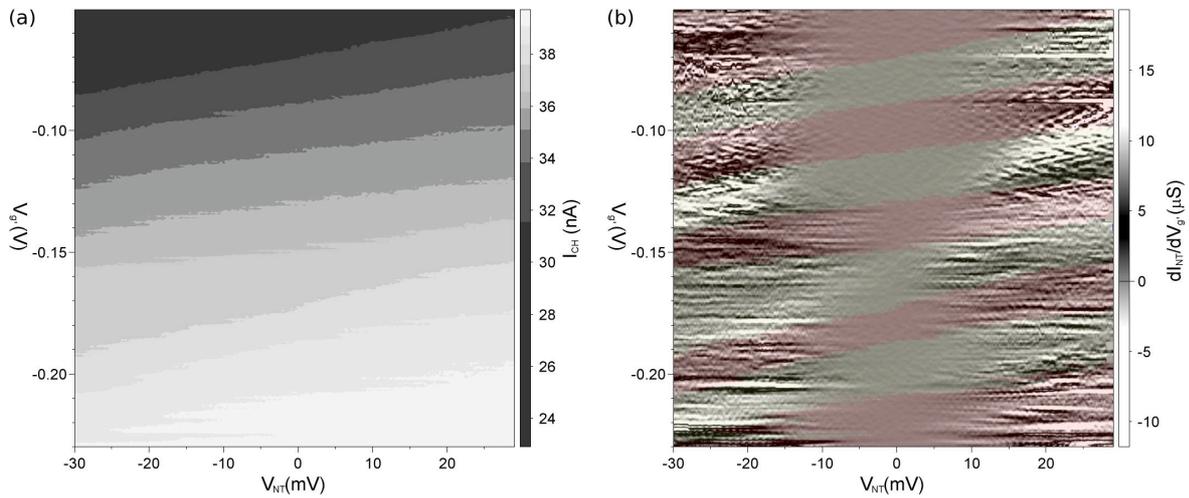
Figure 3.19: Device E2t1-2: Conduction correlations. (a) Surface plot of channel current profile with a banded colour scale - note the broadened region at negative V_{NT} below $V_{g'} = -0.15$ V. (b) When overlaid on the nanotube differential conduction this region matches up with broadening in the nanotube transport.

3.5.3 Discussion

Of 29 devices only 4 showed conduction, of which three were destroyed by ESD events before measurement was complete. Six were lost at the bonding stage and of the remaining 18, seven nanotubes were broken or removed during contacting and nine were destroyed by ESD events before their failure could be explained.

The ranges and sampling of the transport measurements described above were not chosen with consideration to predicted scales of the expected transport properties but rather, given the limited gating range available, to identify and focus on any features that emerged. Semiconducting tubes have a band gap approximated by $E_{gap} = 2\gamma_0 a_{C-C}/d$ for a tube of diameter d, where the C-C tight-binding overlap energy $\gamma_0 = 2.7$ eV, and the nearest neighbour C-C distance $a_{C-C} = 0.142$ nm [42]. The nanotubes used here are all less than 2 nm wide, giving a minimum band gap of ~ 380 meV. At $V_g = 0$ V this would be the width of a zero-conductance region in the low temperature current-voltage (I-V_{NT}) characteristics. At high gate values conduction is restored at lower bias voltages in this region, but remains immeasurably small for low bias at all V_g for $T < 20$ K [139] (resistances are above 10 GΩ even at 77 K [109]). Furthermore, since the nanotubes are made p-type by the Pd contacts conduction is more likely to be observed at negative gate voltages, but the gating range is foreshortened by breakthrough. It is therefore unlikely that these measurements would have revealed any transport in semiconducting nanotubes at 4 K, which may account for the apparent failure of some unbroken nanotubes. It also suggests that the four conducting tubes were metallic (4 tubes out of 16 measured tubes is roughly consistent with the incidence rate for metallic and small band gap tubes). The 4 K current-voltage behaviour of metallic nanotubes is also non-linear at low bias due to non-ohmic tunnel junctions at the metal-nanotube interface, and without an insulating gap [106, 140] (in agreement with Figures 3.13a and 3.17a-inset).

The Coulomb oscillation back-gating period $\Delta V_g = e/C_g$, where C_g is the capacitance between the nanotube and gate. Assuming all nanotubes have a similar width which is small compared to the distance to the gate d, then $C_g \propto \epsilon_r L/d$ [97]. GaAs has a relative permittivity ϵ_r of 12.1, compared with 3.9 for SiO$_2$, increasing the gate coupling with the nanotube. Published data from 200 to 300 nm long dots on silicon [50, 89, 91, 141] implies $\Delta V_g \sim 10$ to 20 mV for similar dots on GaAs with a 70 nm deep 2DEG. If the dots on GaAs stretched all the way between the nanotube contacts (here up to 1.4 μm) then the gating period might be as little as 2 mV. The highest resolution data taken here has only 3 data points per mV, while the sweeps for E2t8, with the

clearest oscillation period, span 3 mV per data point so a quantum dot spanning the whole inter-contact length would not have been detected. However, barrier-free regions much longer than 300 nm are not expected [132]. By the above comparisons a 23 mV period for E2t8 implies a 120 nm long dot, and for E2t1-2 27 mV gives 100 nm. The reported extent of non-conducting diamonds in V_{NT} is around ± 15 mV for a 300 nm long dot [91] and ± 45 mV for a 100 nm dot [139]. This agrees with the observed Coulomb diamonds in E2t1-2, while the other devices were measured over too narrow a bias range to see them.

At 4 K a single quantum dot in a metallic nanotube gives strongly periodic conduction peaks with zero conductance in between, in contrast to the irregular peaks and smooth background seen here. The former are a sign of disorder in quantum wires [142] and arise with multiple series tunnel junctions [143]. Disorder gives rise to significant backscattering which localises electronic states, dividing the nanotubes into a series of shorter dots [139] with different Coulomb blockade periods superimposed.

The irregular Coulomb behaviour observed is unlikely to be through nanotube ropes, whose transport is dominated by one nanotube, without loss of clarity in Coulomb diamonds [77, 81, 97]; the need to conserve crystal momentum along the nanotube axis limits inter-tube transport between parallel tubes of different chiralities [144]. An alternative transport regime to Coulomb blockade with varying peak periodicity is Kondo resonance, but this arises with high contact transparency where conduction is in excess of e^2/h [49, 76, 91]; the largest measured here was $0.15e^2/h$. Nanotubes can be doped by certain substrates, giving rise to inhomogeneous charge distribution in the nanotubes [145]. Modelling on polar III-V substrates [120] shows that surface dipole potentials strongly affect the nanotube band alignment resulting in some overlap with the substrate, hence autodoping the nanotube. However the absence of actual charged dopant atoms means there is little increase in scattering. Also the top layer of the heterostructure used here is undoped and the surface topography was chosen to be as smooth as possible to avoid sharp kinks.

Because the devices were designed around only the target nanotubes, the high dispersion density meant that some devices were contacted by multiple nanotubes, though some still showed no conduction. E2t1-2 was the only conducting device which survived to be imaged after measurement and has a substantial network of overlapping tubes and ropes connecting the contacts at a short distance from the intended device tube (they were too close to the device to be removed by plasma etching, and are imaged through the protective photoresist - Figure 3.20). Conduction through nanotube networks a few microns wide is smoothly varying [140, 146], with transport
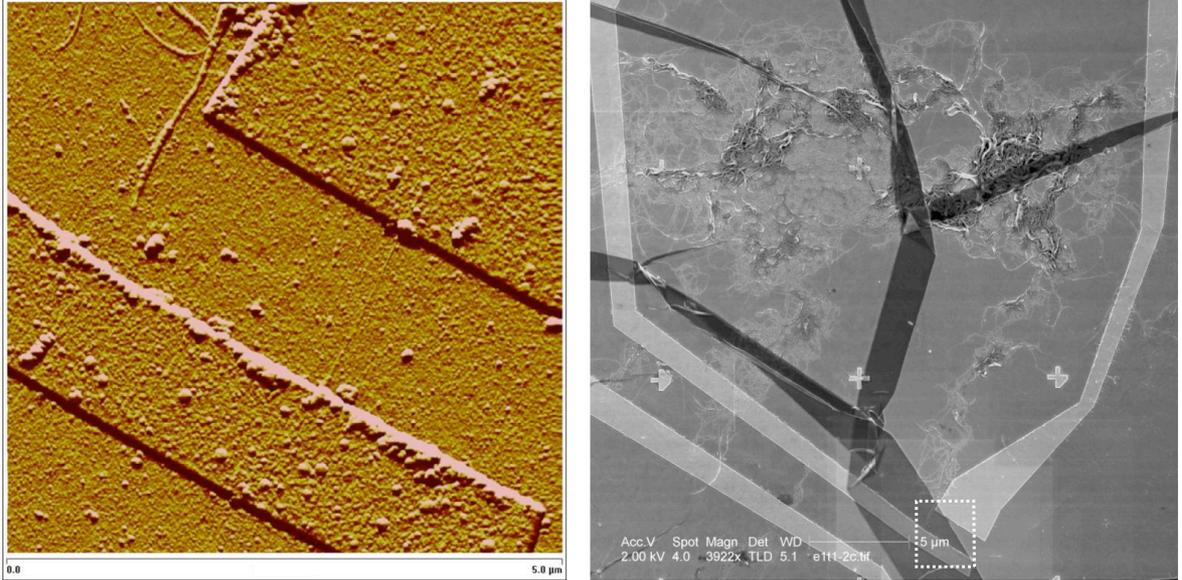
Figure 3.20: Nanotubes short-circuiting E2t1-2. (left) The target nanotube shown by AFM and (right) an SEM scan showing a large mat of nanotubes connecting the contacts at a distance from the target nanotube (at the bottom of the image in the box). The dark regions are folded PMMA.

dominated by hopping between nanotubes [136]. Crossed metallic nanotubes show linear junction conductance of order $0.1e^2/h$ (250 kΩ), similar to that observed for E2t1-2. Metallic-semiconducting junctions have a substantially reduced conduction of $10^{-3}e^2/h$ (30 MΩ), with a rectifying Schottky barrier [7, 85], and limit conduction in low density networks where all-metallic percolation paths are not possible [140]. The weak Coulomb blockade data for E2t1-2 is qualitatively similar to that for a single metallic-metallic junction (Figure 3.21).

The pinch-off data falls into three groups. Devices spanned by a large number of nanotubes, such as E2t1-2, pinch-off at depletion; the surface gate is in effect continuous and the connected nanotubes are gating even if they are not conducting - they shine in SEM images because they are electrically connected. Devices with only a single tube spanning the channel (only confirmed data with broken tubes is available) pinch off a little after depletion, at between 0.7 and 1.0 V. Once all nanotubes were removed pinch-off is at a few volts. That limited correlation was seen between channel and nanotube conduction is not surprising in retrospect. Both quantum dot and 1D features would be better defined at low temperatures, and the inter-dependence between the channel and the nanotube for gating and bias makes this device design an unsuitable system
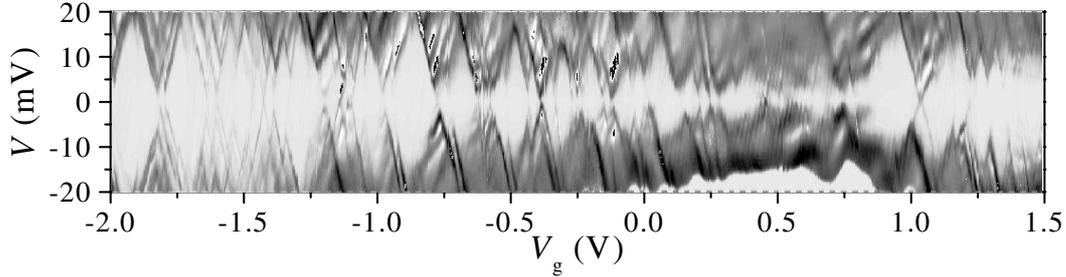
Figure 3.21: Metallic nanotube junction. Conductance dI/dV plot for crossed metallic nanotubes at 1.5 K. Note the weak Coulomb blockade around $V_g = 0.5$ V (from [7]).

for detecting single electron transport.

If a quantum dot spanned the whole length between the contacts, tunnel barriers at the metal-nanotube interface might be influenced by additional depletion in the gate below due to the bias difference. This would bend the nanotube bands close to the contacts in such way as to broaden Coulomb blockade peaks, but shorter dots in the middle of the tube would not be similarly affected. The degree of slant in the conduction features of E2t1-2 relative to $V_{g'}$ and V_{NT} is indicative of the location of the tunnel barriers - were they near the drain contact then no slant would be expected, and near the source contact it would be in line with the channel conduction data. The quantum dot(s) being measured is apparently closer to the source. As the 2DEG depletes, the gating experienced by the nanotube no longer varies linearly with the applied gate voltage, hence the broadening of peaks. Once the gate has depleted altogether there is no effective change, and the nanotube conduction plateaus.

In conclusion, carbon nanotubes and a shallow GaAs 2DEG have been successfully used for mutual gating. Four devices showed repeatable conductance peaks which dominate the transport, but no conductance gap. The most likely scenario is that all are metallic nanotubes, separated into multiple Coulomb islands either by structural defects or substrate interactions, with GaAs inducing greater disorder than silicon substrates. The lack of zero-conductance regions in Coulomb blockade data indicates parallel conduction paths through additional interconnected nanotubes. Coulomb oscillations for E2t8 suggest a single quantum dot, and weak Coulomb diamonds for E2t1-2 could arise from a single metallic-metallic nanotube junction.

3.6 Alternative device processing strategies

3.6.1 Modified device design

The first generation of devices successfully demonstrated proof of principle for charge transport detection between carbon nanotubes and GaAs but suffered from two substantial flaws. The poor nanotube dispersion introduced short circuits and made the nature of the measured devices uncertain, and the device design gave too little independent gating control of the nanotube and channel, relying instead on a coincidence of transport properties.

A second generation focused on improving the device design to address these problems. All the new devices were identical and included a split gate without a nanotube, allowing comparisons between nanotubes, and between crossed and uncrossed channels. The contacted nanotube length was reduced to 300 nm, to make single quantum dots more likely, and the additional split gate completed a cross formation with the contacts (see Figure 3.22). These were 750 nm apart, differently shaped and not centred on the nanotube. This asymmetry allowed them to act simultaneously as both a split gate and side gates for the nanotube.

The nanotube dispersions were examined more thoroughly, including SEM imaging of the bond pad area (but not over the mesa, so as not to damage device nanotubes). Where before only the coordinates of device candidate tubes had been plotted for the overall chip design, this time the entire mesa was imaged by AFM at high resolution and the contact design overlaid on this map to avoid all other nanotubes. For additional security most of the contact length was laid on top of a thin layer of insulating cross-linked resist. This involved a three stage e-beam process. First short sections of contacts were laid to pin the nanotubes in place and prevent their loss during subsequent stages. Next the cross-linked resist was fixed using a high e-beam dose over the desired regions. Finally metal was laid to connect the devices to the optically defined contact arms. Some bond pads were shared between multiple devices to maximise the number of contacted nanotubes. When bonded the chip was fixed on top of another piece of wafer to raise it within the package and allow AFM access without demounting the chip. Seventeen devices were created over 4 mesas (see Figure A.6 for a sample mesa and Section A.2).

At 4 K none of the devices showed nanotube conduction, and channel conduction was erratic. Under SEM 11 devices were found to have suffered ESD events, again probably a result of inappropriate storage (Figure 3.23). Channel conductance mea-
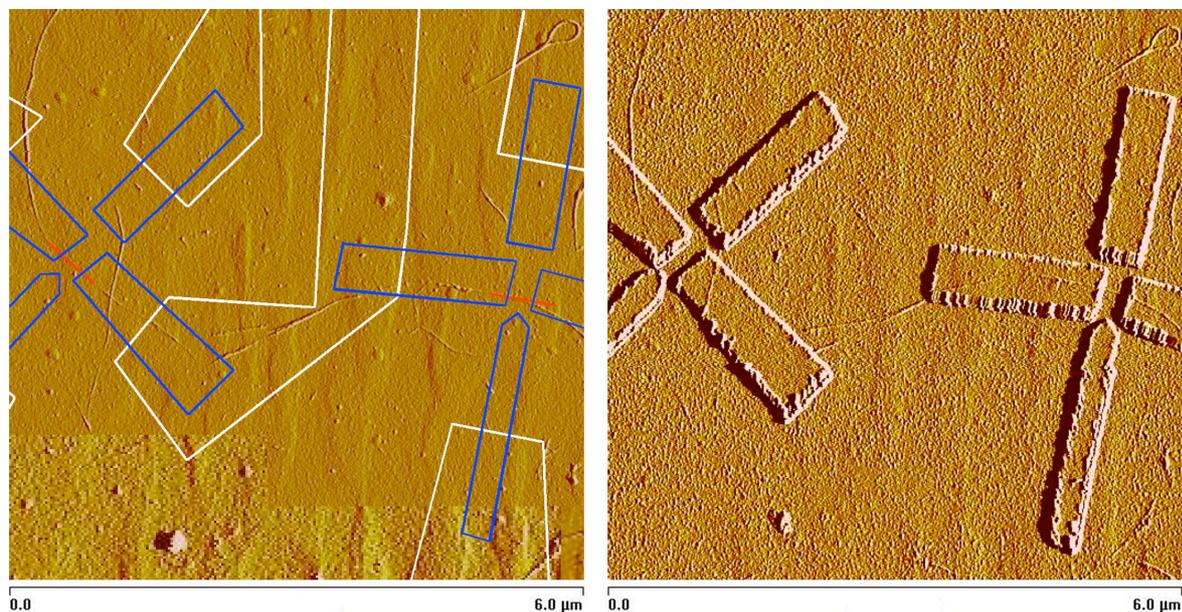
Figure 3.22: Alignment in e-beam lithography. The left-hand image shows the device design overlaid on the original AFM map, and the right-hand image is the device after the first stage of lithography. The alignment error is less than 200 nm.
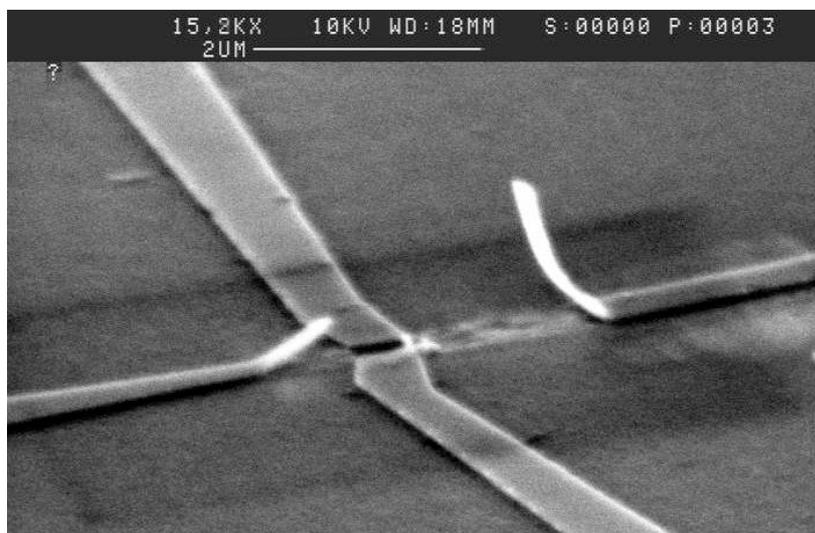
Figure 3.23: Electrostatic discharge damage. The split gate ends have been blown off the surface, and some GaAs beneath the device expelled. The darker surface region is due to carbon deposits from SEM imaging.

surements of the surviving devices showed split gate pinch-off at $V_{g'} = -3$ V, consistent with values from the first devices in the absence of nanotubes.

3.6.2 SEM nanomanipulator

The fabrication time for these devices was over nine weeks, if beginning with nanotubes already dispersed on complete mesa devices. With the high failure rate, and also several months planned downtime for the e-beam lithography system, it was not practical to embark upon a third set using the same processes. The Advanced Technology Institute at Surrey University have successfully positioned MWNTs using a nanorobotic manipulator housed within an electron microscope [147, 148], by using the focused electron beam to weld with carbonaceous deposits [149]. An attempt was made using their system to position SWNTs across the surviving contacts from the second generation devices, thus saving substantial fabrication effort. Initial work was promising but was not continued by the Surrey group.

3.6.3 Ink-jet printing

The apparatus available the Group for nanotube dispersion was inadequate, and gave rise to many of the subsequent fabrication and measurement problems. Superior facilities and experience were available at the Cambridge Centre for Advanced Photonics and Electronics (CAPE), who offered to attempt dispersions on GaAs. They ink-jet print thin film transistors using an Autodrop piezoelectric dispensing system [150]. HiPCO SWNTs are suspended in N-methyl-2-pyrrolidone (NMP) [98] with one hour in a 20k Hz Bioruptor sonication bath at 200 W output. After 30 minutes ultracentrifugation at 30 krpm they are further fractioned and filtered though a 0.7 μm filter paper to remove aggregates. Once on the substrate the solvent is allowed to evaporate slowly in ambient conditions, then annealed at 140℃ to remove residual NMP.

Lower concentrations were used for tests on GaAs and initial results gave very clean surfaces and a suitable nanotube density, but all connected in loose networks. CAPE did not process sufficient further substrates to permit deduction of whether this was a problem with the suspensions or a function of the substrate, and no isolated nanotubes for device use were obtained.

3.6.4   In situ CVD growth

Since this project was begun there had been some developments in low temperature CVD growth of carbon nanotubes, including at CAPE. They achieved vertically aligned nanotube material at temperatures as low as 120℃ using plasma enhanced vapour deposition and a nickel catalyst, which lowers the growth activation energy from 1.2 eV to 0.23 eV. The lowest temperatures give rise to mostly short carbon fibres; SWNTs are produced from 350℃ but show a strong Raman signal for defects [151, 152], and from 400℃ nanotubes produced are suitable for devices. Later improvements used iron and Al/Fe/Al multilayer catalyst films patterned with magnetron sputtering. The feedstock is acetylene and growth time is five minutes [153].

These results are approaching temperatures compatible with CVD growth on GaAs, which would solve many of the problems involved with dispersing nanotubes. The sample is heated from the graphite stage on which it is mounted, with thermocouples on a reference silicon substrate used to gauge the temperature. Since the calibration is uncertain the effect of the growth routine on GaAs was tested by processing a set of Hall bars. They annealed four chips from the same wafer used for the nanotube devices, in vacuum at estimated temperatures of 350, 400, 450 and 500℃. Along with a control sample these were then patterned with a standard Hall bar (Figure A.1), bonded and measured in a cryostat rig dedicated to wafer assessment at 1.5 K by Dr I. Farrer.

Heating to 500℃ in a vacuum was expected to cause As desorption, creating a conducting layer at the wafer surface such that Shubnikov-de Haas oscillations no longer tend to zero at zero field. Additional Si diffusion into the 2DEG would cause an increase in carrier concentration but decreased mobility, and the opposite if away from the 2DEG. The mobility and carrier concentrations measured, even for the chip annealed at the highest temperature, were comparable to the measurements made immediately after the wafer was grown (see Figure A.2). This indicates that either the annealing temperature is lower than estimated, or the time spent at temperature is too short to detrimentally affect the heterostructure.

Iron atoms migrate into GaAs, however the concentration is only 10^4 atoms per cm^3 at 500℃ (compared with dopant atoms at 10^{18} atoms per cm^3), and less than 1 at 350℃ [154]. As a precaution the GaAs was first patterned with 10 nm thick aluminium oxide, and then a sub-nanometre iron film. CVD was carried out at 500℃ and CAPE recorded significant nanotube material by SEM (Figure 3.24), but none was found by AFM on the chip returned to us nor did CAPE carry out any further growth.
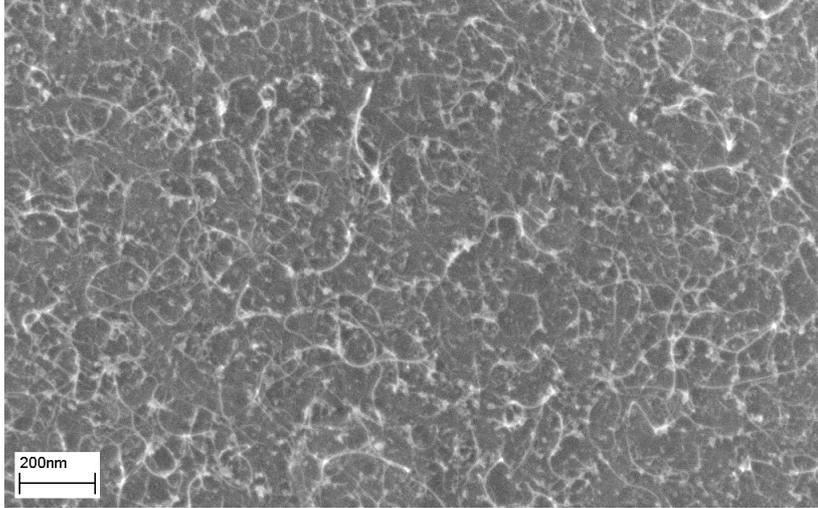
Figure 3.24: CVD-grown carbon nanotubes on GaAs. The nanotubes are too dirty and dense for device use, and were not present on chips received from CAPE.

3.7   Subsequent developments in the field

Since this project was ended some progress in aspects of the work has been reported elsewhere. Jensen et al. [155–157] successfully spun dichloroethane suspended HiPCO tubes onto amorphous As topped GaAs. The amorphous As is later desorbed from the surface and the tubes are contacted with gold and MBE grown epitaxial (Ga,Mn)As. Stobbe et al. [158] did similar work on undoped (100) GaAs. Meyer et al. [159] suspended tubes in 1% SDS solution for deposition on GaAs, but mostly achieved small bundles while Liang and Roth [160] did achieve isolated tubes for devices but on Si-GaAs interfaces.

The most complete work has been done by Tsukamoto et al. [141], who fabricated a device with a Ti-contacted nanotube on a GaAs heterostructure mesa. They used the 2DEG as a gate and measured Coulomb oscillations in a quantum dot between the nanotube contacts 200 nm apart. They also use a nearby quantum point contact as a switch to control the gate, but not for charge sensing. They too found that the resistance of nanotubes on GaAs was consistently around ten times that on silicon, but have not pursued the project further [161]. This agrees with my low measured conduction and implies that interactions with the GaAs substrate are inducing substantial disorder. The only other transport measurements are reported by Wang et al. [96]

67

who fabricated a device with a Y-junction MWNT lying across a micron wide channel etched in a GaAs/AlGaAs heterostructure. They also report pinch-off at around -1.1 V in the channel, consistent with my data from single nanotubes.

A complementary experiment has been done by Gustavsson et al. [162], performing the same measurements as attempted here, with an InAs nanowire on a GaAs substrate. They successfully used a quantum dot defined in the nanowire to probe current fluctuations in a quantum point contact within the 2DEG below. This is a useful demonstration of charge transport sensing between two different materials and between the wafer surface and a buried 2DEG.

The only in situ growth on GaAs has been reported by Engel-Herbert et al. [163], who use a fast sample loader to minimise the time the wafer is exposed to the growth temperature (typically 10 minutes). They use a CrNi catalyst but require 700℃ to achieve a reasonable yield of SWNTs, and still find visible damage to the GaAs. A new 'hot zone' technique by Dittmer et al. [164] only heats a small area of the substrate. A 400 nm wide tungsten electrode on the chip surface is covered in catalyst. This resistive bridge is heated to 800℃ with a high current, and SWNTs grow from its edge. An additional DC electric field parallel to the substrate surface aligns growth away from the bridge. On silicon the substrate temperature decreases by $100℃/\mu$m, and the bulk substrate reaches only 60℃. Nanotubes produced are up to 8 μm long, therefore a significant length of nanotube lies above unheated substrate, making this process useful for CMOS devices. However the thermal conductivity of SiO_2 is 1.5 W/m/K, compared with 50 W/m/K GaAs. The temperature gradient on GaAs would be around $3℃/\mu$m, so the entire mesa would be heated to over 600℃.

Ruppalt et al. [165–167] have developed a dry contact process for dispersing nanotubes which gives results qualitatively independent of the substrate. HiPCO powder is first impregnated onto a fibreglass sheath which is placed into direct contact with the substrate surface. Van der Waals forces between the substrate and tubes strip individual SWNTs (and infrequently small bundles) from ropes on the sheath. Comparable results are seen on GaAs, InAs and silicon, and isolated tubes only are observed to preferentially align to the [1$\bar{1}$0] lattice direction on GaAs, agreeing with predictions by Kim et al. [120] on InAs. They also observe carrier transfer between the two systems, with n-GaAs doping the nanotubes slightly n-type so that their conduction gap aligns with the GaAs gap. This is promising for balancing the p-type shift caused by the contacts and may also provide some explanation for increased disorder in the nanotubes [168]. Another substrate independent process was found by Gonzalez et al. [169]. They observed that nanotube bundles formed in the gas phase from a catalyst aerosol were

charged, while isolated tubes were electrostatically neutral. The latter were filtered out and deposited on a substrate at ambient temperature by an electrostatic precipitator.

3.8 Future directions

Interactions between nanotubes and the GaAs surface are a critical factor in these experiments, for dispersions and potentially transport too. One of the substrate-independent dispersion techniques described above might address the former, and etching a narrow trench beneath the nanotubes suspends them and reduces substrate induced disorder [130]. This was in fact attempted early on during this work, in relation to another research project, but the chip shattered before measurement. A solution to both issues could be the addition of a barrier layer. A few nanometres of SiO_2 on top of the mesa would present a silicon surface during nanotube dispersion, which is better studied, would prevent autodoping and might smooth the GaAs surface topography, without shielding the nanotube significantly from the 2DEG.

Because breakthrough occurs at such small negative gate voltages a shift towards n-type conduction would increase the gating range available for measuring semiconducting nanotubes. Doping with potassium achieves some n-type conduction [136, 170], and more recently others have reported using calcium [133] and magnesium [135] electrodes. Both have low work functions (Ca 2.8 eV, Mg 3.7 eV) and give a low Schottky barrier height for electron transport. Measuring nanotubes on a calcium substrate Okada and Oshiyama [145] observed that the metal Fermi level crosses several nanotube π-bands, facilitating electron transfer to the tube. However, the calcium electrodes gave high contact resistances, and at a given positive gate voltage less than 1% of the current seen with palladium electrodes with an equivalent negative voltage. This is thought to be due to a dipole layer at the metal-organic interface causing an abrupt potential drop [133]. Nosho et al. [133] discovered an alternative where amine or toluene treatment gives a temporary shift to n-type conduction, which disappears after a few days.

Useful tools for assessing devices would have been electrostatic force microscopy (EFM) [86, 94, 117, 171], Kelvin probe microscopy (KPM) [172] or scanning gate microscopy (SGM) [86, 173]. These non-contact AFM modes are described in detail in Section 5.2. Scanned images reveal scattering centres such as defects, kinks and breaks not detectable by topographic AFM [173]. Under EFM charged regions appear in high contrast to the insulating substrate, so that conduction paths can be followed. These techniques would reveal the locations of tunnel barriers bracketing quantum

dots, percolation through nanotube networks, and explain why some devices were insulating. The Group has a fridge with an AFM appropriate for such measurements, as described in Chapter 5, but time was not available on the facility for this experiment.

3.9   Conclusions

The eventual goal of this work had been to use a GaAs heterostructure to detect charge transport in a nanotube. Such a result requires a substantial degree of control over device fabrication and measurement, but the broadest conclusion of this work is that fabrication involving nanotubes is both complex and very specific to the materials and processes used. The additional challenges and difficulties that arose made that original goal impracticable in this instance, but ultimately a number of nanotube devices were fabricated and measured on GaAs. The dispersing process was found to be highly substrate dependent, so that established recipes were of little use, and the available apparatus was inadequate in the circumstances. Nevertheless, reasonable dispersions were achieved using dichloroethane and the first manipulations of carbon nanotubes on GaAs provided evidence for a weak surface binding energy. Anionic surfactants were shown to be inappropriate for use on GaAs.

The weaknesses of the dispersions were revealed in transport measurements, where the high density of nanotubes and low device yield made it impossible to accurately determine the nature of the devices measured. Despite some poor data sampling, highly localised gating of nanotubes by the GaAs 2DEG was demonstrated, through the measurement of Coulomb blockade and quantum dots, as well as gating of the 2DEG by nanotubes. There was also limited evidence of correlation between transport in the two systems and disorder in nanotubes induced by the GaAs substrate. Many lessons were learnt and implemented for the improvement of device design, and steps made towards improved dispersion and CVD techniques for use on GaAs.

The added complexities of working on GaAs make this system less immediately attractive than well studied nanotube-on-silicon devices, but since this work was begun no fabrication strategies have been demonstrated which exceed the potential of GaAs heterostructure-based devices for functional complexity in relatively simple structures. The successful fabrication of working devices in this project is a sign that most of the experimental difficulties have been identified and largely overcome, and further improvements have been suggested, therefore such a system remains a valuable avenue of practical future research.

Chapter 4

AFM Lithography of Graphene

4.1 Introduction

Recently discovered graphene has substantial potential for use in microelectronic devices and circuitry, on account of its physical dimensions and exceptional electronic properties. A significant challenge towards this goal is patterning graphene flakes. In particular, the ability to define graphene ribbons sufficiently narrow to introduce an energy gap would give transistor functionality. However, the length scales needed are at the limit of state of the art electron beam lithography.

This chapter describes the first work using a charged AFM tip to pattern graphene, including single layers, achieving feature sizes as small as 5 nm. Combining experimental observations with a detailed study of the literature, a comprehensive molecular-level mechanism for the lithography process is presented. Additionally the role of water on and beneath the flakes is explored, and an explanation found for the extra measured depth of thin flakes.

The first successful cutting was performed in December 2006, with the bulk of the experimental work carried out during May to August 2007, and revisited following repair of the AFM from May to July 2009. The introductory sections describe the state of the field at the time the work was begun and later sections reference reported developments as it progressed.

4.1.1 Motivation

Being a zero-gap semiconductor, bulk graphene cannot be depleted to define conduction paths using only surface gates, nor used for switching applications such as field-effect

transistors; physical barriers or breaks are necessary to confine the charge carriers [174]. Analogous to the derivation of carbon nanotubes' semiconducting behaviour, graphene ribbons of finite width will have quantised k_y, and their two-dimensional energy dispersions will be split into a number of one-dimensional modes. If these 1D modes do not pass through a conduction and valence band intersection then such quasi-1D ribbons are semiconductors, with a finite energy gap [175], and the carrier concentration can be depleted with conventional gates.

The energy gap is inversely proportional to ribbon width W: $\Delta E_C \approx 2\pi\hbar v_F/3W$ [46]. Barone and Scuseria [176] show theoretically that to achieve band gaps comparable to Si or GaAs (1.1-1.4 eV) the ribbon would need to be 1-2 nm wide. Ribbons with 'armchair'-orientated edges of 8 nm width would yield a maximum band gap of 0.3 eV, and by 80 nm the maximum band gap is 0.05 eV, which is of a similar order to thermal energy at room temperature. Son et al. [177] model more orientations, and the results from Barone and Scuseria [176] match their highest bounds, while Liang et al. [178] find the band gaps to be even smaller.

Reduced flake dimensions may also permit observation of 1D confinement, visible as steps in a gate sweep of the conduction, when the energy level spacing is greater than thermal broadening. The band-structure of graphene nanoribbons (GNRs) is complex, due to edge effects. Ouyang et al. [179] predict lowest subband gaps of 0.56 eV, 1.18 eV and 2.12 eV for a 2 nm wide channel and Brey and Fertig [180] also find the level spacing of similar order to the band-gap. For 0D structure, the level spacing for massless carriers in a box of size d is $\delta E \approx v_F h/2d$, so a quantum dot of dimension around 10 nm is needed in order to detect level spacing in graphene at room temperature [181].

A comment here on graphene versus carbon nanotubes. In a 1D quantum box of width W, quantisation requires $k_\perp = n\pi/W$ so the allowed energy states are spaced π/W. For a cylindrical box of circumference C, $k_\perp = 2n\pi/C$ so the levels are spaced $2\pi/C$. Therefore CNTs have a larger gap for the same confinement dimension [179], plus the electronic states in a GNR are not degenerate, but in a CNT are doubly degenerate because, according to boundary conditions, the CNT wavefunction is periodic around its circumference, where the GNR wavefunction must vanish at its edges [46]. This means that, in the face of difficulty achieving narrow GNRs, CNTs appear to be a more practical solution for devices. However, the problem remains with CNTs of positioning and characterising more than a few for complex devices, where a large graphene sheet could be patterned into many identical and orientated ribbons, and even complete circuits.

4.1.2   Principles of scanning probes

The work in this thesis makes extensive use of scanning probes, for micron-scale imaging, as a mechanical tool, as a lithographic tool, and for electrical measurement. The fundamental elements of a scanning probe are a flexible cantilever with a tip hanging from its end, sharpened at its base to the width of a few atoms, and some means of measuring bending of the cantilever. The substitution of different cantilevers and tips allows different functions such as those listed above to be derived from the same microscope apparatus. The first scanning probe microscope capable of imaging on the atomic scale was the scanning tunneling microscope (STM) [182]. An STM tip is metallic, and when scanned over a conducting surface at a small distance the tip-to-sample tunneling current is monitored. A negative feedback loop keeps this constant by controlling the height of the tip above the surface. A surface plot of the resulting height adjustments is a topographic map of the local density of states of the sample, which is closely related to its surface morphology.

The atomic force microscope (AFM), developed soon after [183], instead employs an insulating tip on a flexible cantilever (Figure 4.1). When the tip is pushed very close towards the sample surface it is deflected by repulsive electrostatic forces. The cantilever bends with a deflection proportional to the force between the tip and the surface, and so a scan at constant deflection (again maintained via a negative feedback loop) gives an image of surface topology. This is called contact mode operation. Alternatively a stiff cantilever is oscillated close to resonance for tapping mode imaging [184]. In this case, as the tip approaches the sample, interaction forces cause the amplitude of oscillation to decrease. The surface profile is given by the probe height adjustments (through feedback) needed to keep the amplitude constant. Because the tip is not dragging across the surface there are no lateral forces and contact with the surface is intermittent [185]. This technique is therefore more suitable for delicate or soft samples, whose surface morphology would be changed or damaged by contact mode scanning. Unlike STM, AFM can be used to image insulating surfaces.

Methods to measure the cantilever deflection or oscillation amplitude include tunneling [183], laser interferometry [184] and piezoresistivity [186], and the lateral scanning motion and height adjustments are facilitated by piezoelectric positioners, acting on either the cantilever or the sample plate. The vertical measurement resolution is < 1 Å but the lateral resolution is limited by the tip profile, since all fine surface features appear as a convolution image with the tip shape. Measurements such as nanotube diameters are therefore taken from the feature height, not width, and the
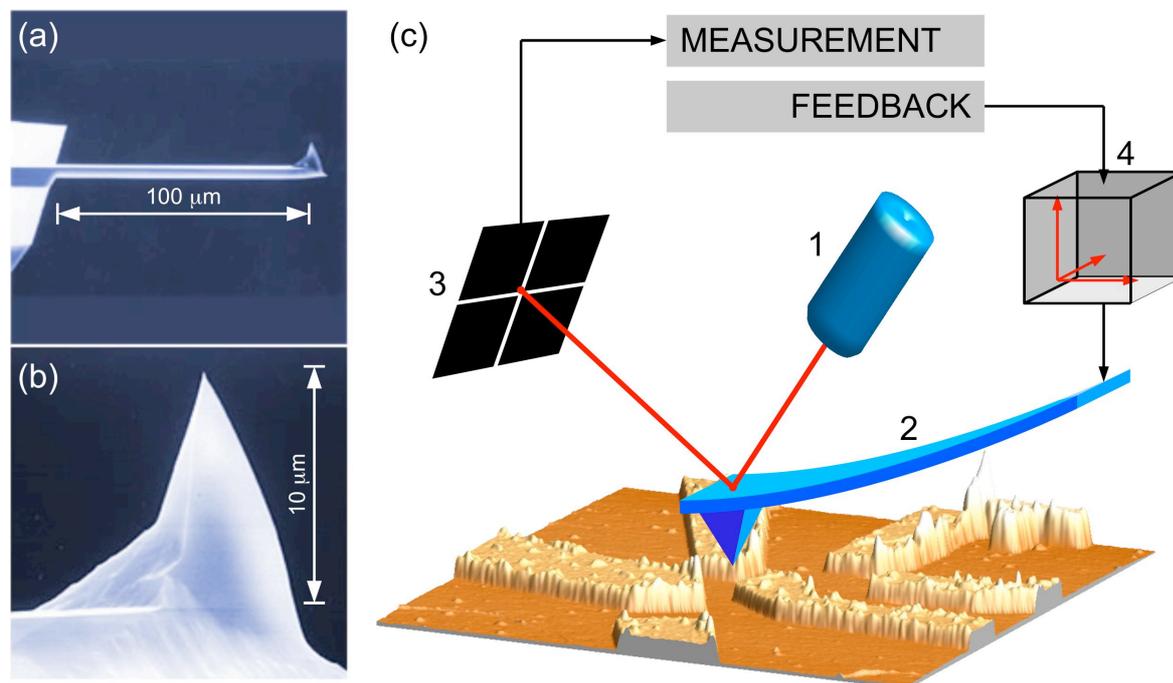
Figure 4.1: AFM with laser interferometry detection. (a) A probe cantilever and (b) tip. (c) A laser spot (1) is reflected off the deflected cantilever (2) towards a four-quadrant split photodiode detector (3). A feedback signal directs adjustment of the piezo-electric positioners (4) acting on the cantilever.

tip's ability to detect narrow indentations is limited by its sharpness (the typical apex width is around 10–30 nm). Appropriate scanning speeds are determined by the speed of the feedback response and the sample roughness - high or sharp features will be broadened if the scanning speed is too fast for accurate adjustments to the probe height. With the Veeco AFM used for most of this work the maximum tapping mode scanning rate at high resolution is around 10 μm^2/min.

Descriptions of additional scanning probe techniques used to make local electrical measurements are given in Section 5.2.

4.1.3 Lithography techniques

Optical lithography is the industry standard for patterning semiconductor devices, with electron-beam lithography giving increased resolution for use in specialised research applications (see Section 3.4.1). Combining these techniques with subsequent chemical or plasma etching, the smallest reported feature size achieved in graphene prior to this

work was 80 nm [187] (subsequent advances are described below). This was not the best resolution of state of the art lithography systems, but graphene's extreme thinness requires exceptional control of the etch speed, and it is vulnerable to undercutting of the resist and substrate. An order of magnitude improvement is needed to create devices operational at room temperature.

Scanning probe microscopes have been extensively used in research for patterning the surfaces of materials [185], employing their sub-nanometre positional accuracy and sharp tip profiles, and may be able to achieve the necessary feature sizes. Cutting and indentation on semiconductor surfaces has been demonstrated using a range of AFM techniques and substrates: GaSb cut in hard contact mode with SiN tips [188], GaAs scratched using doped silicon tips [189] and diamond tips [190], and holes punched in GaAs with an oscillating tip [191]. The first of these achieved lines 20 nm wide. An alternative procedure is local anodic oxidation (LAO), where the tip of the SPM probe is used to locally oxidise the surface beneath it, chemically or electrically altering the material, or removing it altogether. Murphy [192] and Day and Allee [193] have used an AFM, and Dagata et al. [194] an STM to pattern silicon and GaAs; a tip held at negative voltage with respect to the substrate draws lines of oxide, which have a higher surface profile. The mechanism deduced is that the electric field dissociates water, and O^- or OH^- ions are driven into the material surface [192]. Irmer et al. [195] and Gordon et al. [196] have achieved similar results on thin titanium films, for defining split gates, and Majumdar et al. [197] used an AFM to pattern PMMA, for subsequent etching.

Of particular relevance to this study, LAO has been widely reported on HOPG (Highly Ordered Pyrolytic Graphite) using STM [198–205], and once with contact mode AFM [90] (see Figure 4.2). Amongst these reports there is no universal agreement regarding the necessary conditions for lithography, nor the mechanism itself (see Section 4.5). However, the common threads are a useful guide to the basic conditions for this experiment, and the disagreements a good indicator of likely difficulties.

All authors agree that the success and extent of cutting is determined principally by the applied potential difference between the SPM tip and the sample, not the current flowing. Almost all find that the tip must be at negative voltage with respect to the sample, with the exception of Venema [116], Penner et al. [204] and Mizutani et al. [203], who achieve ambipolar or positive bias cutting. In all cases lithography is only successful above a repeatable threshold voltage value, determined by conditions. It is most commonly found at around 4V [116, 198–200, 203, 204], and occasionally at 2.5V [201, 202]. The only recorded AFM cutting had a threshold of 6V [90]. Larger
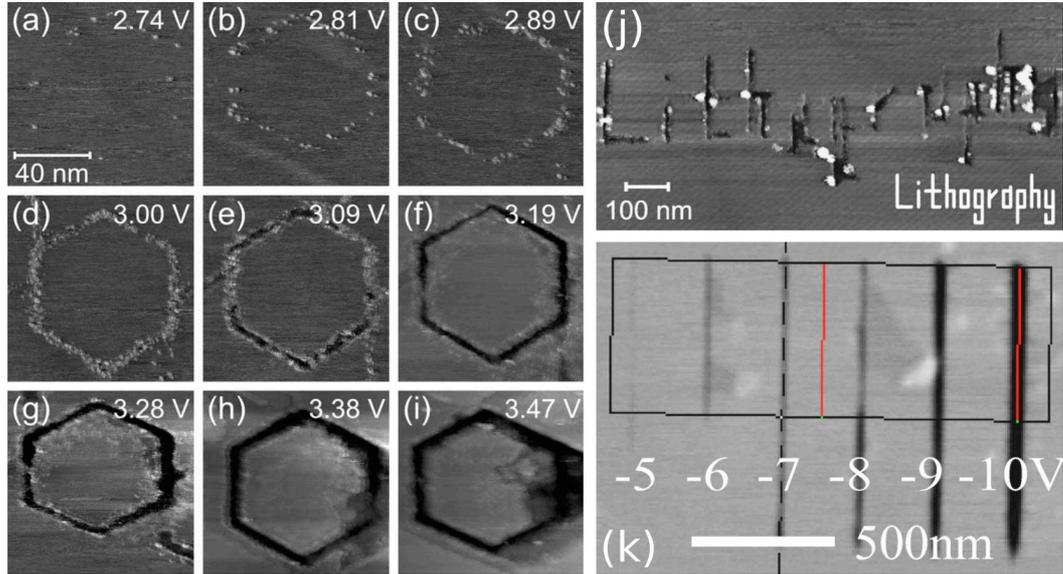
Figure 4.2: Prior STM and AFM LAO on graphite. (a)–(j) A pattern drawn by STM at increasing sample bias, transitioning from humps to holes which become deeper and wider, and an attempt at writing letters (from [200]). (k) Lines cut by AFM with changing tip bias, which become deeper and wider above a threshold bias value (from [90]).

voltages give wider and deeper holes, and slower tip velocities (or longer stationary voltage pulses) enhance the lithography and lower the threshold [90]. Hiura [200] finds lithography to be effective at tip velocities from 0.1 nm/s to 1 mm/s (with corresponding thresholds from 1.95V to 3.68V), and that 100 nm/s gives the best compromise between uniform carbon removal and speed. The magnitude of the tip-sample current during lithography has no effect on feature size, and even very large currents cannot induce lithography below the threshold voltage [198]. At voltages just below threshold some sources observe deposits in place of holes, the nature of which is unknown [204, 205]. All except Venema [116] and Kondo et al. [201] use a humid atmosphere (40% humidity and above [90]), and Albrecht et al. [199], Tang et al. [205] and Penner et al. [204] observe that the presence of water is obligatory. Results reported include holes reliably one atomic layer deep [200, 203], lines 30 nm wide [200] and even smaller dot features [199]. The most common proposed mechanism is that under the applied electric field carbon is oxidised by water on the graphite surface [200, 201, 203], but tunneling current induced sublimation of surface atoms [201] and field emission [90] are also suggested.

4.1.4 Experimental aims

The main aim and challenge of this work was to controllably cut thin lines in single-layer graphene that could be used to define devices for low-dimensional transport measurements, in particular a point contact. The target channel width was 20 nm or narrower, which was at the resolution limit of the most advanced electron beam lithography equipment. As well as offering smaller potential feature sizes, SPM LAO does not leave any chemical residues, eliminates the risk of e-beam damage to the graphene (see Section 3.4.4), is a single-step process and is less vulnerable to misalignment and writing errors. The voltage is applied only locally and the bulk graphene is left pristine. Using an AFM rather than an STM means that a small graphene flake can easily be found and imaged on a non-conducting substrate, a UHV environment and passivated surfaces are not required [192], and a device can be measured in situ during lithography. Since the resolution-limiting feature is the tip profile, a relatively basic AFM is sufficient (the hardware used here was 10 years old), requiring only minor modifications for applying a voltage to the tip.

The uncertainty and variation in the literature implied a substantial parameter space. The primary goal was to explore the most likely parameters, establish a repeatable process and achieve a fuller understanding of the mechanisms involved, before applying this to fabricating devices for measurement. Also of interest were mechanical manipulations, using the AFM to fold, peel or cut single layers from thicker flakes.

4.1.5 Isolating graphene

Single-layer graphene was first isolated on a non-graphitic substrate by Novoselov et al. [27] in 2004. Their exfoliation technique involved etching the surface of a platelet of HOPG, a material with large, pure, ordered graphite crystals, to leave small mesas. The graphite surface was pressed against wet photoresist, which was then baked. The mesas, now fixed to the photoresist, could be cleaved from the HOPG. Scotch tape was used to repeatedly peal layers from the mesas until suitably thin flakes were left. These were released in acetone, into which the silicon substrate wafer was dipped, and to which some of the flakes stuck. Ultrasonic cleaning in propanol removed thicker flakes, with the few-layered flakes (FLG) adhering strongly to the substrate due to van der Waals forces.

Very little of the graphitic material deposited by this process is single layer graphene (SLG), and films thinner than 50 nm are partially transparent to light [27], a single

layer absorbing only $\pi\alpha = 2.3\%$ of white light (α is the fine structure constant) [206]. However they can still be seen using an optical microscope due to an interference effect. The substrate wafers have a surface oxide layer, and the extra optical path for light passing through the graphene and then reflecting off the oxide-silicon interface introduces a slight colour shift compared to the path through the oxide alone. The oxide thickness and wavelength of incident light are chosen to give this shift a noticeable contrast; under white light a 300 nm thick oxide is blue-violet but thin graphene on its surface shows as blue [27] (Figure 4.3). A 90 nm oxide gives even better contrast but devices are more susceptible to gate leakage [207]. Even a 5% thickness error substantially reduces contrast [28]. For monochromatic light a wavelength can be chosen to view the graphene on any oxide thickness [207]. The very palest flakes visible are those of one or two layers.

AFM is not reliable for confirming the exact number of layers. The interlayer distance in graphite is $\sim$3.35Å [27] but tapping mode AFM measures the height of SLG on silicon oxide at between 1 and 1.6 nm. The extra spacing and variation may arise from van der Waals interactions, contaminants on and beneath the graphene, and the surface morphology of both flake and substrate. The most reliable technique for establishing the thickness is Raman spectroscopy, as single and double layer graphene have markedly different spectra [208]. At 514 nm excitation SLG has a much smaller G peak (1580 cm^{-1}) at slightly higher shift, and a sharper G' peak (2700 cm^{-1}) at slightly lower shift and with no trace of the secondary peak seen in bulk graphite. In bilayer graphene the G' peak is broader and comparable in height to the G peak. In FLG and graphite it is the G peak which is tallest and sharpest. This enables almost unambiguous differentiation of single, bi- and few layer graphene (Figure 4.4a). An alternative technique has since been found by Hibino et al. [209, 210], where the resonance of incident electrons with the quantised conduction band states in FLG causes dips in the reflectivity, as viewed by low-energy electron microscopy. Adjacent regions of different thickness therefore appear in high contrast, and the incident electrons can be tuned to highlight regions with a specified number of layers (Figure 4.4b).

Novoselov et al. [212] subsequently found that a good yield of SLG could be achieved simply by rubbing HOPG against the substrate wafer, and SLG can be transferred from silicon to other substrates using a PMMA transfer process [211]. Exfoliation currently remains the quickest and most reliable route to the largest, structurally coherent flakes, but is not practical for industrial fabrication, and even for research purposes the inability to specify size (presently up to 7000 μm^2 [213]) and location presents limitations, though improvements will no doubt arise. Alternatives being investigated
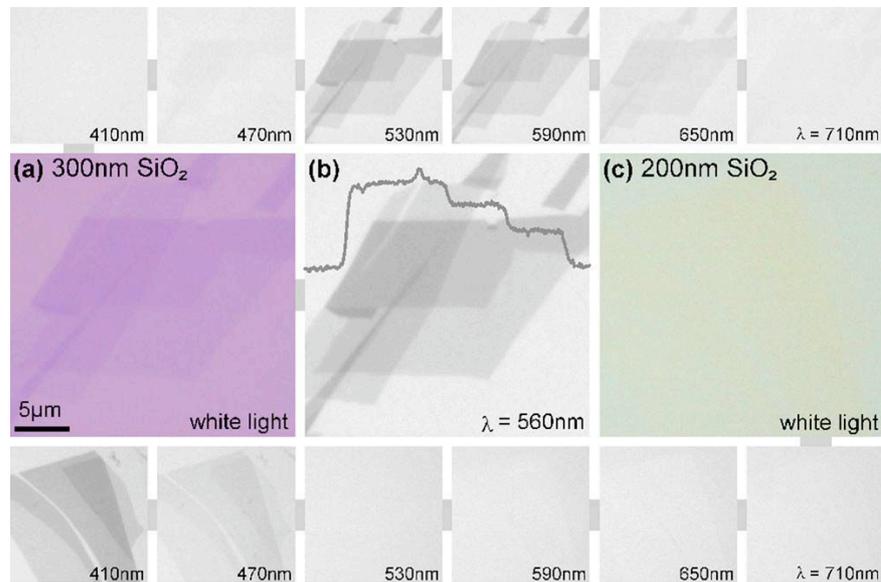
Figure 4.3: Optical identification of graphene. (a) Graphene on 300 nm thick SiO$_2$ imaged with white light. The top panels show the same flake under a range of monochromatic light wavelengths, of which 560 nm (b) gives the best contrast. (c) A similar flake on 200 nm SiO$_2$ is nearly invisible, but can be seen clearly under monochromatic light at 410 nm (bottom panels) (from [207]).
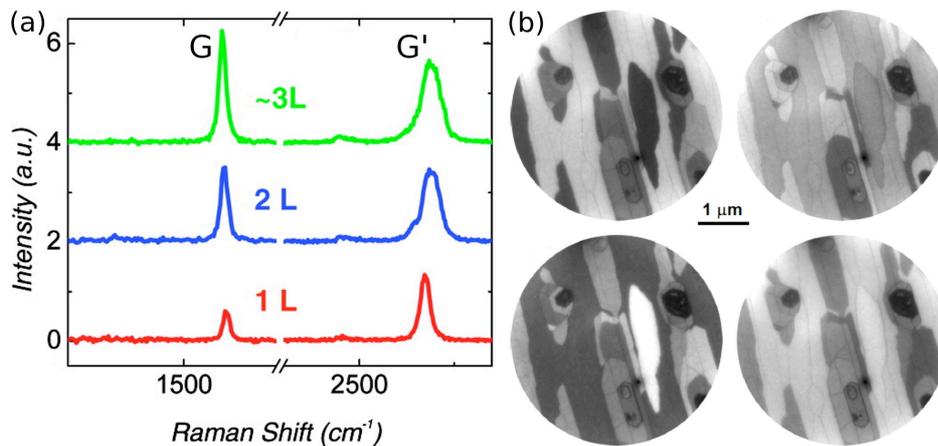
Figure 4.4: Counting layers with Raman spectroscopy and electron microscopy. (a) By comparing the G and G' peaks Raman spectroscopy at 514 nm can reliably distinguish between single, bi- and few layer graphene (from [211]). (b) Low energy electron microscopy shows regions with different numbers of layers at varying contrast, in images taken at different beam energies (from [210]).

include chemical vapour deposition (CVD) [211], reduction of silicon carbide [187, 214], epitaxial growth [215, 216], and transfer printing [217].

Like carbon nanotube dispersing, effective graphene exfoliation is a mechanically and physically complex process; success is a function as much of practice as of understanding. During the latter stages of this project a Masters student assisting with the work, David Tyndall, investigated the exfoliation process [218]. He used Indian natural graphite [219] and substrates cleaned with tetrafluoromethane and oxygen plasma etches to ensure good adhesion. Several tapes were compared, and the best chosen for the physical properties of the glue layer, which must be homogeneous on a scale smaller than the size of flake desired (some tapes have lumpy and sparsely distributed glue) and sufficiently adhesive to easily fracture the graphite. He found Minitron 1008R dicing tape to be the most effective. Repeated exfoliations of the same graphite fragments were seen to reduce the lateral size of flakes more than their thickness, so he employed only five tape-to-tape transfers, giving a reasonable dispersion of flakes. The transfer was done with the mechanical action of a rubber clamp, to give repeatability. The pads could be heated and the applied pressure controlled, and the size and density of thin flakes was found to increase up to the available limit of these variables (70 ℃ and 700 kPa). Higher temperature and pressure left more glue residue, which was removed with a 5 minute anneal at 400℃ in hydrogen/nitrogen forming gas.

Graphene production is labour intensive with a low throughput, therefore every flake found must be considered for study. The cost of a graphene flake purchased for this work [213], at 50 pence per square micron, equates to around £750 trillion per gram, which is several times the sum of human wealth for an area of graphene the size of an Olympic swimming pool.

4.2 Experimental set-up

4.2.1 Apparatus

The AFM was a Digital Instruments 3100 Nanoscope IV, featuring a scanner with 'closed loop' feedback to monitor and correct probe drift. This feature was critical both for mapping surfaces and plotting patterns, reducing positional error to less than 3% of the scanning field size (it was broken for an extended period, and the substitute 'open loop' scanner had a 100% uncertainty in travel distance and no directional fidelity during lithography). The AFM was shielded from external vibrations and electrical noise by a damped table and retractable hood. A selection of different AFM tips

were used. Ordinary imaging, and ultimately most lithography, was done with conventional antimony-doped silicon tips from Veeco (NCH-W/TESPW). These give an acceptable sharpness for imaging with an apex diameter of 30–40 nm (features atomic scale in height could be imaged with widths at around 10 nm if the tip had a sharp convex point). They pick up contaminant particles from the surface and gradually blunt, giving a typical life-span of around 50 scans. When changing tips would cause additional problems it is sometimes possible to sharpen a blunt tip, or remove dirt and secondary tips, by disabling feedback and hitting the tip rapidly against a high surface feature. The tips are designed solely for non-contact (tapping) mode use, so are rigid and brittle, and liable to damage a delicate sample if used in hard contact mode.

For LAO the tips must be conductive, and previously reported AFM lithography employed metal coated tips, such as Pt-Ir [117]. Some Ti-Pt coated tips from Mikro-Masch (NSC14) were used here. They were a little sharper but because conduction was only through the metal coating they became insulating after wear or damage. This effect was hard to identify above other factors that could be responsible for failure of the lithography. A third type used were tungsten carbide coated (apex diameter 20–30 nm) from NT-MDT (NSG 03/2C), as also used by Kim et al. [90]. These were more robust and had more flexible cantilevers, giving them greater compatibility with contact mode imaging and lithography, but too expensive for use throughout this work, given the short tip lifetime. All images here show cuts by doped silicon tips.

Usually the bias voltage was supplied to the tip from the AFM's internal source, with a range of ± 12 V. Lithography was controlled by the Nanoman component of the Veeco Nanoscope 6 control software. The humidity within the AFM hood could be raised by bubbling nitrogen gas through a glass cylinder of water and then a tube into the hood. The cylinder rested on a hot-plate, to warm the water and enhance the humidifying, which was measured with an R500 Digitron Instrumentation hygrometer. Simply placing a beaker of water inside the hood did not have a perceivable effect, possibly because the cavity was not sealed. Device packages appeared to make good ground contact to the metal sample plate, but a clamp was sometimes used for added certainty. The plate could be isolated from ground via software control, and was later additionally connected to the inner ring of a coaxial cable. The current path from the device could then be made direct to an external ground, or via an ammeter or alternative voltage source (a Keithley 2400 SMU was used for both tasks). During the later stages of the work a package holder designed to accommodate an LCC package was added by David Tyndall, allowing measurements of bonded devices to be made in situ and during lithography.

4.2.2 Graphite and graphene devices

For lithography tests on graphite, an Indian natural graphite flake was cleaved with a scalpel blade, exposing large areas of clean and continuous graphite lattice. It was then fixed with silver glue onto a small metal plate, which was set on the AFM sample plate. Most of the FLG and SLG used was provided by Andrea Ferrari's group at CAPE. We provided pre-cut silicon substrates, topped with a 300 nm thick oxide and metal alignment marks and bond pads, onto which Cinzia Casiraghi and later Antonio Lombardo exfoliated graphene and performed Raman analysis to locate and identify SLG. An average yield was two flakes on a 4 mm by 4 mm chip. A change to a new, denser alignment mark pattern, needed for more reliable device patterning, reduced the effectiveness of their exfoliation process, so the flakes used for in situ measurement were exfoliated by David Tyndall.

The device design process was simpler than that for carbon nanotubes, but shared many features (see Section 3.4.2). Apart from laying the bond pads and alignment marks, no further preparation of the cleaved wafer was required. The alignment marks were visible by optical microscopy, giving an approximate location co-ordinate for flakes. Flakes and their nearest alignment marks were imaged by AFM, often requiring several scans to be stitched together due to the larger size of the flakes and alignment mark spacing, into maps spanning up to 100 μm. On this scale the error between the scan tube X and Y axes becomes significant, so images required geometric corrections. Contacts to the flakes were plotted onto these maps in a CAD file containing the existing metal features, to which the e-beam writer could align; the final alignment error in the laid contacts was typically less than 100 nm. Because of the low density of graphite material there is very little risk of short-circuiting, except close to the flake, but the largest, optically visible, graphite fragments are high enough to disrupt thin metal laid across them, and must be avoided.

Initially each flake had a custom-designed pattern, to make optimum use of the scarce material (we received only 10 flakes in 18 months). The first devices consisted of multiple interlocking fingers overlapping the flake completely, to guarantee good contact. When the alignment proved to be better than expected later devices had contacts just overlapping the flake edges, leaving a large area of clear flake for lithography experiments (see Figure 4.5). A large number of contacts gave redundancy for a reliable ground connection and helped pin the flake; in instances with few contacts, or where alignment marks were added after exfoliation, flakes tended to fold or roll up during wet processing stages. Where possible, contact metal evaporation was the

only processing stage carried out subsequent to exfoliation. The yield and size of flakes produced in-house towards the end of the project was greater and Dr Malcolm Connolly designed a set of standard contact patterns in an optical lithography mask, greatly speeding up the fabrication process by eliminating the design stage and e-beam lithography queue (between a fortnight and 3 months).

Because of the shortage of material, robust and reliable contacts were preferable to good electrical conduction, therefore gold over titanium was used instead of palladium (see Section 3.4.5). All devices survived lift-off complete, and none were damaged by ESD events under normal use. The metal had to be thin (less than 40 nm total) to reduce the region of shadow at its edges in AFM imaging, and limit high image contrast which makes shallow features such as cuts hard to resolve. The 2-terminal resistance of contacted SLG devices was 1–2 kΩ. Chips were fixed with G-varnish into an LCC package on top of a GaAs riser, lifting the silicon surface to be level with the top of the package rim, and bond pads were used on only two sides of the chip, to leave access for the AFM tip unobstructed by bond wires. To protect the delicate graphene efforts were made to prevent static discharge during bonding, as with nanotube devices. The silicon back gate and selected contacts were connected to a ground contact using silver glue over the edge of the chip. When the package was resting on the AFM plate all bonded contacts were grounded, and when in the package holder the ground path could be chosen.

A major problem was the reliability of the PCs controlling the AFM. Performing certain sequences of actions in the lithography software, or queuing certain tasks, caused the software to crash. If the microscope was scanning at the same time it would therefore lose feedback control while in contact with the surface and, in the time taken to restart the computer, inevitably broke the tip and usually damaged the flake and contacts. This occurred several times before it was possible to isolate the cause. In addition a motherboard and a fan on successive PCs developed faults, also causing unpredictable crashes and freezing. Five out of eight devices used were damaged or destroyed in this manner (see Section A.4). In some cases the graphene itself was unharmed and the contacts simply dragged across its surface but, when this broke the ground connection, cutting was made impossible.
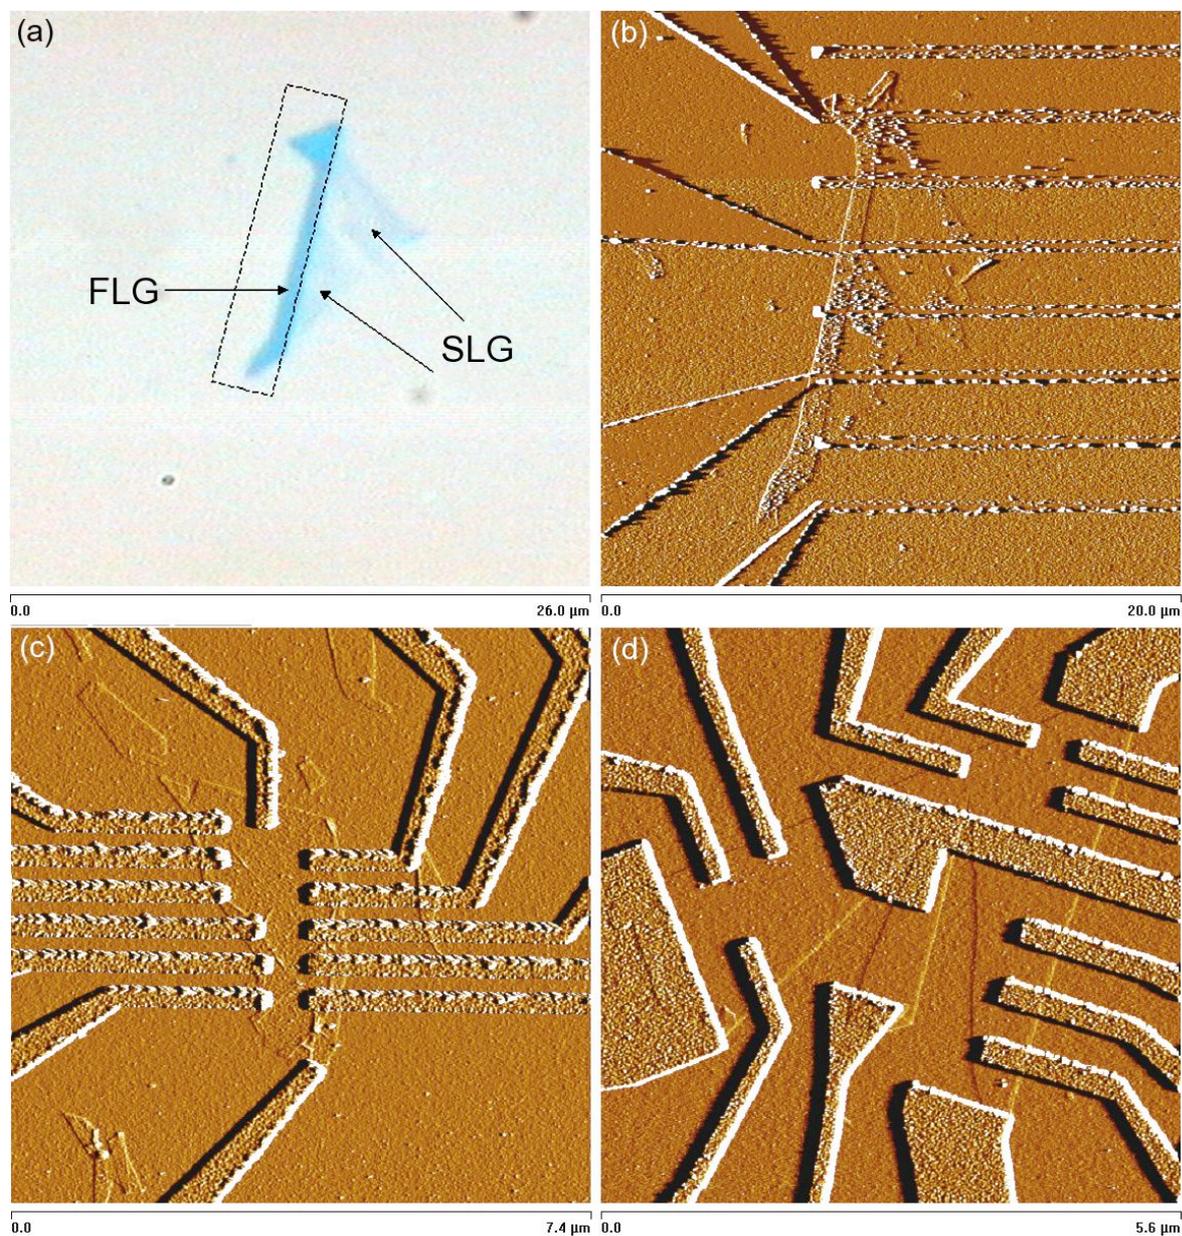
Figure 4.5: Graphene devices for lithography. (a) Optical image of a flake and (b) the same flake with finger contacts (GPE04). (c) & (d) SLG flakes with edge contacts (G04a & G07 - device patterns and fabrication by Mark Buitelaar).

4.3 Cutting graphene

4.3.1 Experimental outline

When carrying out lithography a scan is first made of a small region of flake, usually in tapping mode. The Nanoman software allows paths to be drawn over the scan image. When directed to proceed the tip moves to the beginning of a path at a safe distance above the sample, switches to the chosen lithography settings and follows the path, then returns to imaging settings and re-scans the region. For parametric testing paths were usually straight lines, around 300 nm long and aligned in rows with a single changing variable. Lines were drawn at a small angle to the imaging scan direction, since this presents a wider trench profile to the tip and because extended features parallel to the scan direction may be filtered out by plane-fitting between scan lines.

The variables to be tested were the tip-surface bias voltage, humidity, tip velocity and the tip-surface force. Voltage variation is concerned with the threshold for cutting and the voltage-dependence of the cut profile, and the tip velocity gives a measure of the interaction time for a given surface area. There are three lithography modes, with associated scanning engage parameters which determine the force: contact mode with feedback, where the path is followed at a fixed tip deflection (*deflection setpoint*), tapping mode, with a fixed oscillation amplitude (*amplitude setpoint*), and either contact or tapping mode without feedback, where the tip is pushed vertically into the surface a fixed distance (*Z move*). In tapping mode a smaller amplitude makes the average tip height closer to the surface, and the scanning action tends toward contact mode as the setpoint tends to zero. Also to be confirmed were the need for a ground path, voltage polarity and current independence.

The force of the tip on the surface can be calculated only approximately. Each model of tip is rated with a spring constant, but the standard deviation is up to 70% of the mean, and it is not practical to measure individual cantilevers. The deflection sensitivity is calculated by pushing the tip a known distance towards the surface and measuring the deflection (see Figure 4.8e). Deflection is reported as the output voltage difference between a vertical pair of photodiodes and, since the laser spot must be centred manually each time a tip is mounted and subsequently drifts, the measured deflection when the tip is fully retracted is rarely at zero; the actual deflection is the difference between the measured value and the retracted value.

The noise in the spot position is around 0.2 V, and it drifts by a similar amount within a single image scan. This variation is 20–50% of the typical deflection used here

during lithography, hence the actual force applied has a substantial uncertainty. The doped silicon tips have a mean spring constant $k = 42$ N/m and a mean sensitivity $S = 80$ nm/V, giving an applied vertical force $F_z = 3.4$ μN/V. The W$_2$C-coated tips have $k = 1.1$ N/m and $S = 110$ nm/V, giving $F_z = 0.1$ μN/V. Because of the large error in this conversion, deflection is reported here at its actual value in volts. Unless otherwise stated, all contact mode imaging and lithography is done with feedback and a deflection of around 0.5 V, being the smallest value to reliably achieve and maintain surface contact. This corresponds to a 40 nm deflection and force of around 1.7 μN with silicon tips, and 55 nm deflection and 0.06 μN force with W$_2$C-coated tips.

4.3.2 Lithography on FLG

Preliminary testing raised more questions than it answered, partly because the very first attempts at cutting were immediately successful. These used a doped silicon tip in contact mode on an 8 nm thick flake ($\sim$15 layers) with a single large contact (device G11). The flake surface was covered in small lumps of unknown material. A line with the bias at -4 V scraped away the lumps along its path, -5 V made a shallow 10 nm wide cut and -6 V made a 30 nm wide cut as deep as the flake (Figure 4.6a). However further increases in bias did not achieve any cuts. The first cut was not revealed until a sharper Pt-Ti coated tip was used for imaging. Lines with this tip did not have any effect until the bias reached -10 V. At -11 V a 65 nm wide split gate was made by drawing two 40 nm wide lines end to end (Figure 4.6b). By -15 V the cut was several hundred nanometres wide and no longer had straight edges, instead having many sideways intrusions into the flake as if the oxidation process flared irregularly.

The dirt on the flake surface accumulates over time, forming a continuous layer 1–2 nm thick in a week. It cannot therefore be solely residual PMMA from the metal lithography, and is not observed to increase on silicon, metal or GaAs surfaces, so is hypothesised to be mostly hydrocarbon molecules accrued from the air. If left it makes the identification of successful cuts harder, since shallow cuts in the graphene cannot easily be distinguished from lines cleared only in dirt of unknown depth. Lumps are often dropped at the end of lines, so the tip does not just brush it aside but also picks it up. Large lumps in the way of a path can be picked up by applying a large voltage (-8 to -10 V) locally, and dropped elsewhere by applying a similar bias.

Before cutting, a square area was henceforth swept clean by scanning in soft contact mode, which piled up debris around the edge (Figure 4.7a). Scanning is possible, albeit with poor resolution, with zero and even slightly negative deflection, and therefore
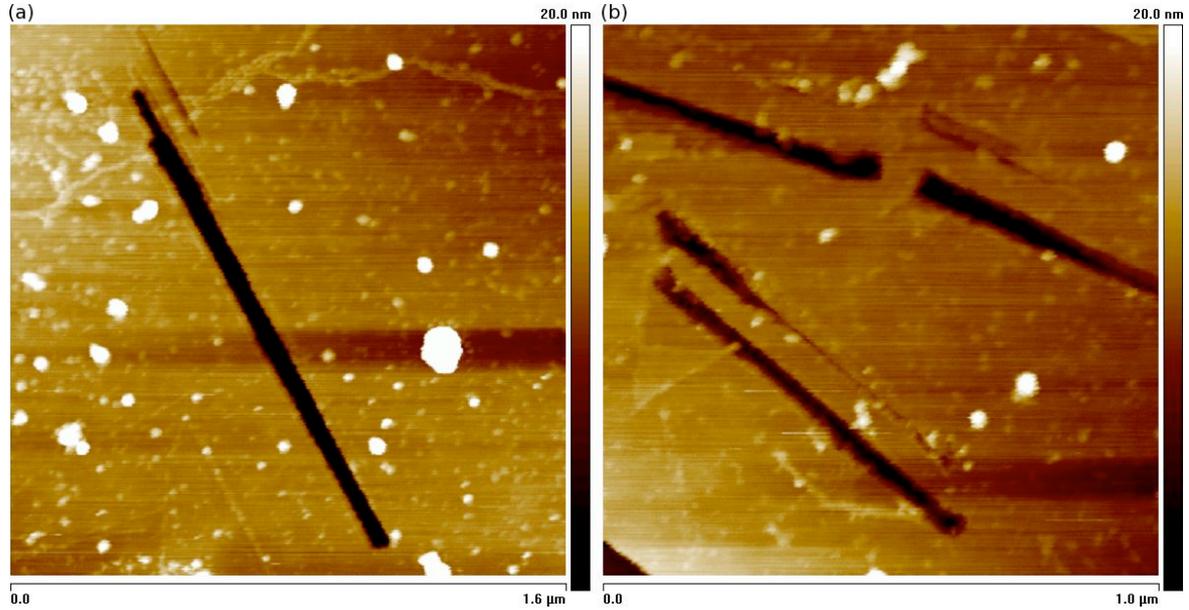
Figure 4.6: First AFM lithography on FLG. (a) The first attempt at cutting FLG - a bias of -6 V penetrated the whole flake thickness. (b) A 65 nm wide split gate cut at -11 V and 10 and 30 nm wide cuts below it.

minimal force, ensuring that there is no damage to the surface during cleaning. Almost all surfaces are covered in a thin layer of adsorbed water (see Section 4.5) and when an approaching tip touches this, attractive capillary forces drag it downwards and it snaps into contact with the surface beneath (see Figure 4.8f). As the tip descends further it is deflected by the repulsive electrostatic forces and the cantilever bends to accommodate this. During retraction the attractive forces keep the tip extended for longer before it snaps out of contact. To scan in the attractive force regime the tip must first be brought into contact with a small deflection and then retracted until it is extended. On graphite a pristine surface was achieved by scanning in hard contact where, with a deflection of up to 3 V and high scan speeds, the tip dug into the surface and pulled several layers of graphene away (Figure 4.7b). A fresh tip must be used for lithography after both cleaning methods.

Another change in surface texture was observed after FLG flakes were connected with bond wires (as opposed to others which were bonded with silver glue) with the appearance of well defined blobs hundreds of nanometres wide and 1–2 nm high (Figure 4.7c–f). These were later identified as pockets of water trapped beneath the flake (see Section 4.4).
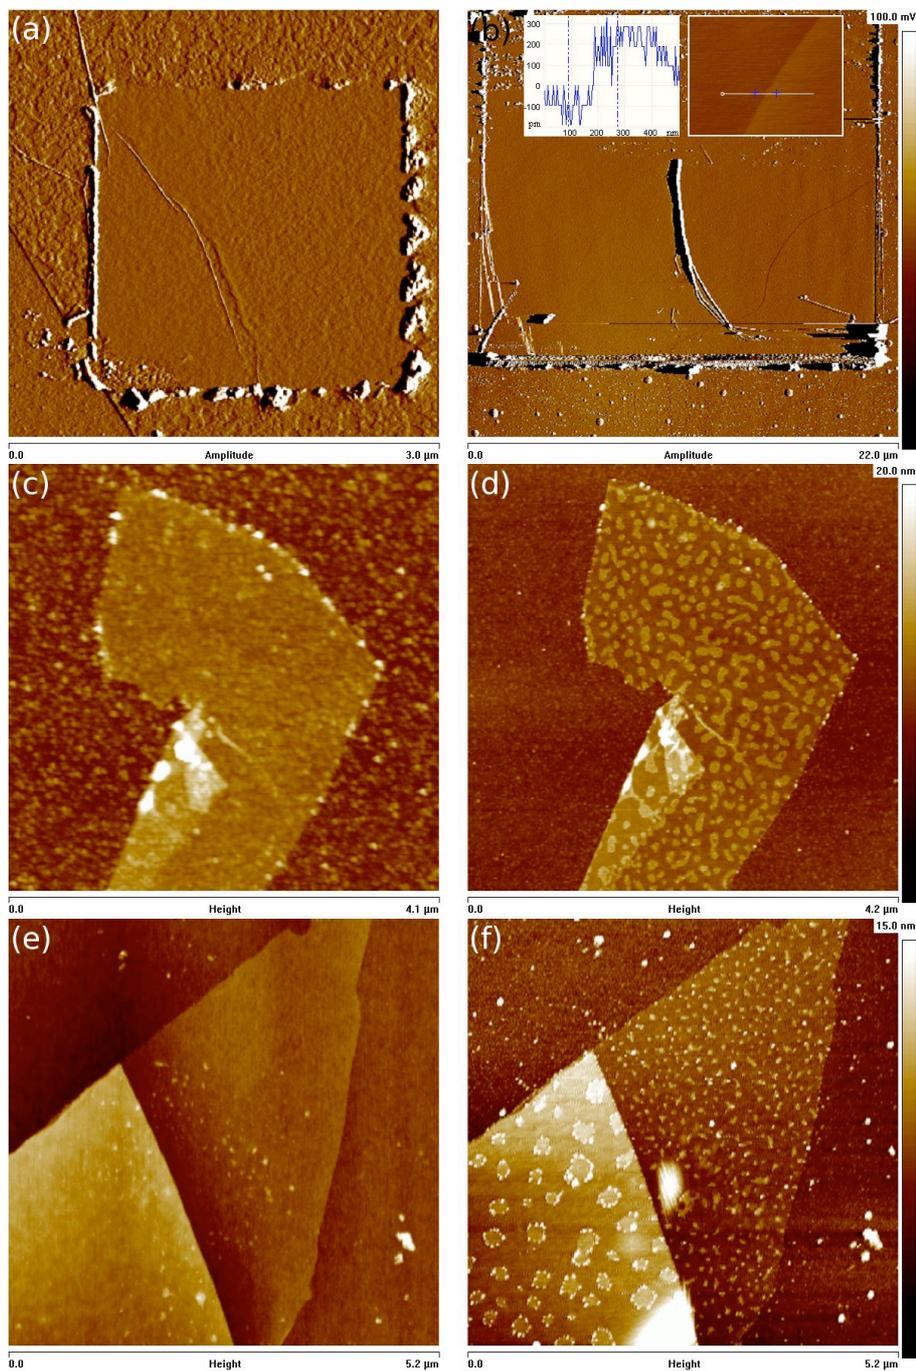
Figure 4.7: Dirt on graphene and AFM cleaning. (a) A square swept clean on SLG, with the dirt piled up around. (b) Several layers of graphite peeled back to reveal a clean surface. (Inset) Enlargement and height profile of a 0.34 nm step edge, buried several layers beneath the surface. (c)–(d) and (e)–(f) Surface changes after wire bonding. Raised patches appear on FLG and smaller lumps on SLG.

A further new surface feature was often seen after a few successful cuts, when tapping mode images were disrupted by what appeared to be puddles scattered across the surface, which were almost certainly water droplets (Figure 4.8) [220]. They represented a real feature on the surface, since they were continuous between scan lines, though with an error of a few nanometres, and could not be eliminated by changing feedback and scanning parameters or the tip. However they did shift between scans (Figure 4.8c–d), and were moved but could not be imaged by contact mode scanning. They were measured as between 1.3 and 2.3 nm high, 50–500 nm wide and very flat, clung to the edges of cuts and disappeared completely after heating the chip. Imaging is sharper outside of the puddles than within, indicating that they are distinct from the thin surface layer, and may be substantially thicker than their measured apparent height; the tip still touches the flake surface but reports a higher position because of a substantially changed mechanical impedance to the tip oscillation. As the tip enters and exits a droplet it pushes and pulls the edge, giving the horizontal streaks in the images. Their incidence did not seem to be related to the ambient humidity; the images in Figure 4.8 were all taken at relative humidities of 40–50%, while in the absence of cutting no puddles were seen even up to 95% humidity (at which Figure 4.7d was taken). Since the puddles appear after the application of a voltage they may be the result of an electrowetting effect, with the balling of existing surface water into discreet droplets or a change in the crystalline phase at the water-gaphite interface.

When stepping the bias voltage between consecutive paths, ridges were almost always drawn at lower voltages than clear trenches, which themselves had raised edges. After the time taken for a couple of scans of the same region ($\sim$10 minutes) the ridges and raised edges almost completely disappeared (see Figure 4.9), so they were formed either of a very volatile substance, or were sufficiently weakly bonded to the surface to be moved even by a tapping mode tip pass. Applying a large bias nearby often cleared them instantaneously. Most cleared ridges revealed a cut beneath, except sometimes those drawn at the lowest voltages. Imaged in contact mode the ridges and raised cut edges were 0.3 nm high, but in tapping mode were measured as high as 3 nm.

A number of problems that led to the failure of cutting were difficult to isolate. Often after a few successful lines a tip would cease to cut. In some cases no effect would be seen until the bias was raised by 4 to 5 volts when it would make a substantial hole in the flake, tens of nanometres deep, with even larger lumps left on the surface and specks and changes to the surface texture around the hole (see Figure 4.11a). This may be because the tip, even on a cleaned surface, had been collecting dirt which did not disrupt scanning but interfered with the cutting. At a sufficiently high voltage a
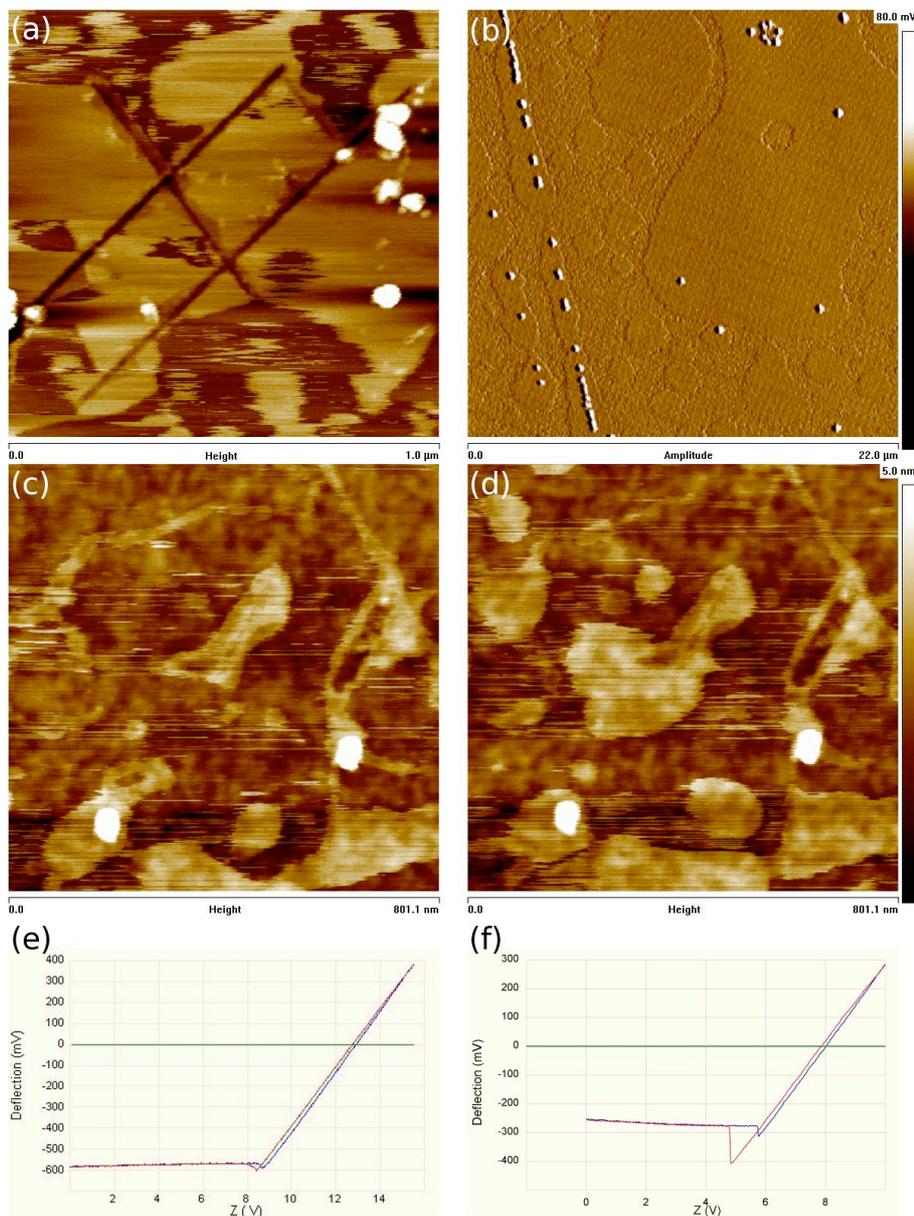
Figure 4.8: Water puddles on the graphene surface. (a) The same cuts as in Figure 4.9 after the appearance of ~ 1.5 nm high patches, which interfered with tapping mode scanning. (b) They are inferred to be puddles of water, seen hear with small droplets. (c)–(d) The puddles move and change shape between scans. (e) Deflection plot for a tip lowered towards the surface and retracted; on normal graphite the thin water layer gives weak capillary forces, and very little 'snap-to-contact'. (f) The presence of the puddles greatly increases the 'snap-to-contact', pulling the tip down early, and pinning it to the surface longer during the retraction.

90
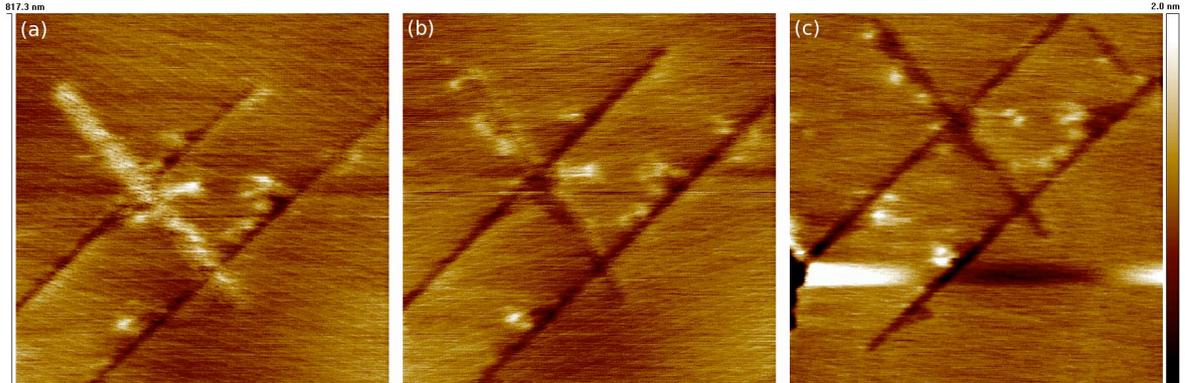
Figure 4.9: Transitory ridges formed on cuts at low bias. (a)–(c) A 3Å high ridge seen immediately after the cross cut is written (-4 V) gradually fades over successive scans, revealing a trench beneath. The short cut in the top-right of (c) has a 25 nm wide point contact at either end. The unusual image noise is a result of weak contact mode imaging using a tapping mode tip; the tip has caught and torn the cut edges in several places.

spark can cross the dirt, both removing it and burning the hole in the carbon [192]. These tips sometimes then resumed effective cutting as before.

In some cases cuts appeared at substantial offsets from the drawn line, outside the scan region. If the cutting mechanism is a function of electric field strength (discussed in Section 4.5) then this is explained by a double-ended tip, with one end blunt but longer, which gives the image, and another slightly shorter and sharper, which gives a greater field density for cutting. Several brand new tips never achieved cuts, indicating that they may have been too blunt from the start.

Additionally several devices suffered from broken ground connections, which was later confirmed to be critical for lithography, and at one stage the voltage source failed, but with little progress yet made in characterising the lithography conditions the fault may have occurred several weeks before it was identified. Because all these problems were frequent and could not readily be distinguished it was impossible at this stage to identify clearly the function and effect of parameters being studied, such as the bias voltage and humidity, in the success and extent of cutting.

The lowest bias at which cutting was detected was -3V (with a W_2C-coated tip), and in most cases the threshold bias V_{th} was around -4V. As the bias was raised beyond this the width and depth of cuts increased non-linearly. The smallest unit of cut depth is a single atomic layer, therefore to identify a functional relationship between depth
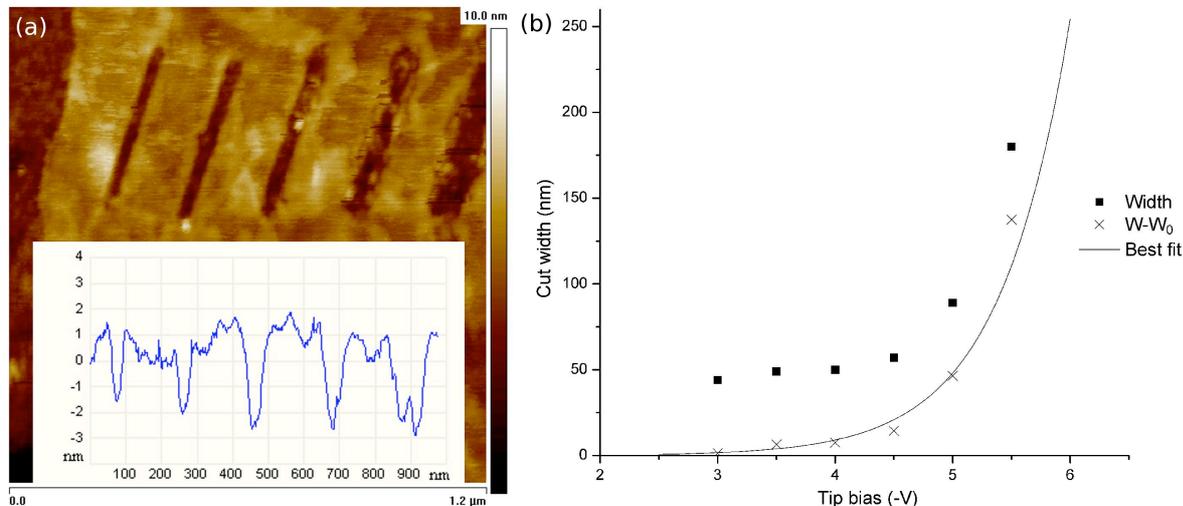
Figure 4.10: Variation of cut width with tip bias. (a) Lines cut with a tip bias of (L–R) -3, -3.5, -4, -4.5, -5 V. (b) Plots of the original cut width W (averaged over 100 scan lines), the remainder after subtracting the tip width $W_0 = 42.5$ nm, and an exponential fit $W = W_0 + \exp(1.7V - 4.5)$ against bias V.

and bias the cuts need to be deep relative to the layer thickness. Most testing was done close to V_{th} and on FLG, therefore the cut depths were shallow and inadequate to derive such a relation - the cuts in Figure 4.10 are all 3–5 layers deep. The cut width W appears to rise exponentially, but tends to a non-zero minimum width. After subtracting this minimum W_0 the remainder is a very close exponential fit. An AFM tip can detect any feature of sufficient height, no matter what its width, but will not detect surface depressions that are narrower than the tip apex, being the width over which the tip is approximately atomically flat. Whatever the cutting mechanism, the whole apex area will act equally on the graphene, and a cut grows outward from a minimum width W_0, not zero as anticipated. V_{th} is therefore the voltage at which a change in depth is first detectable, with the removal of perhaps half the atoms in the area of the first carbon layer beneath the tip. In occasional instances the cutting apex width is narrower than the imaging width, giving an apparent raised V_{th} since the cuts at lowest bias are not detected by imaging.

Where cuts were observed from lower threshold voltages they also tended to be narrower. This was initially thought to be because V_{th} was roughly constant, and narrower tips could image narrower lines and therefore detect cuts made at lower voltages than a broader tip would have found. However, if the first lines cut are the
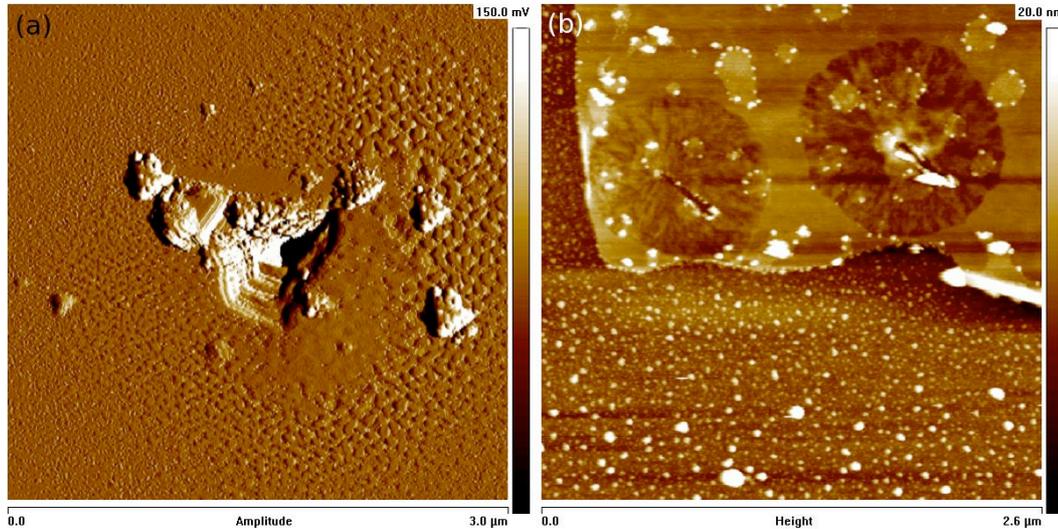
Figure 4.11: Sparking events and high bias cuts on FLG. (a) A hole in graphite resulting from a spark through dirt accumulated on the tip at a bias of -7 V. Note the change in surface texture around the hole. (b) On FLG high bias (-5 V) cuts cause extensive damage to the upper layers, but the lowest 1–2 layers remain intact, with the cut itself penetrating the whole flake.

same width as the minimum imaged feature size, then the observation of thinner cuts from lower thresholds suggests that V_{th} is itself dependent on the tip width. This is discussed in Section 4.5.

At high bias (2–3 V beyond V_{th}) on thick flakes the cut width becomes significant relative to the line length (usually 300 nm in testing), giving a rounded hole in comparison to the straight edged, thinner cuts; over longer cuts straight edges are apparent at higher voltages. However, on FLG the high bias cuts have a very different appearance once the depth has penetrated the whole flake. The width is not graduated between layers, and the upper layers are cut irregularly to much greater widths; bonding to the substrate may stabilise the lowest layers and accelerate the reaction in the layers above (Figure 4.11b).

No correlation between humidity and the success of lithography could be deduced at this stage, and the lowest relative humidity at which cutting was definitely achieved was 34%. Cutting was also attempted in tapping mode, but such a small amplitude was needed (1 mV equivalent, which is around 0.02% of the normal scanning amplitude) that this was qualitatively little different from contact mode. It was also less reliable and did not give any noticeable increase in tip lifetime.

4.3.3   Lithography on graphite

Because of the very limited supply of FLG and SLG further testing was done on bulk graphite, giving an effectively unlimited experimental surface. Even if necessary conditions were not identical, the physics of LAO lithography on graphite and SLG was expected to be the same, allowing detailed analysis of experimental parameters before returning to SLG for final tests. The AFM was now fitted with an external ammeter and ground connection. This permitted confirmation with each path drawn that the tip bias was functioning, the tip had made contact with the flake, and the flake was connected to ground through the chip package, eliminating several sources of uncertainty and increasing the efficiency of the experiment.

The ground connection was shown to be obligatory; when it was disconnected the tip could only sweep a line in the surface dirt, even at high bias, but when reconnected cutting was immediately restored. After drawing a line with a high bias (-12 to -15V) and no ground connection subsequent tapping mode images had a fogged appearance due to the flake charging - they initially showed a smooth, flat surface, through which the highest features, and ultimately the whole true surface, emerged progressively over around 5 minutes.

A current signal was produced each time the tip with an applied bias made contact with the surface. The success and quality of a cut could be established from the current profile during the path's progress. Re-imaging after each path was no longer necessary for monitoring the process; after a set of cuts was complete the whole row was then imaged for measurement of the cut profiles. Successful, clean cuts gave only a single sharp spike of 30–100 pA at the start of the cut, which was a capacitative current induced in the flake (Figure 4.12a). Any further features were buried in the signal noise (of order 10 pA). The bias voltage was ramped after the tip makes contact with the surface, normally within 10 ms. When the ramp time was extended to a few seconds the current spike decayed exponentially over the ramp period.

Certain tips gave a mean current of up to 10 nA for the duration of the cut, overlaid with spikes 10 to 100 times as large (Figure 4.12b–c). The base current is thought to be field emission from the tip (see below), and the spikes are instances of more ohmic contact with the surface. Such current spikes and high plateaux of up to 1 μA are seen more frequently when the tip is shedding lumps of dirt or piling substantial amounts of matter around the edges of cuts (distinct from the shallow raised edges normally seen), and are also associated with short gaps in a cut; the dirt changes the tip profile, weakening the field at the flake surface (Figure 4.12d). Unsuccessful cuts usually gave
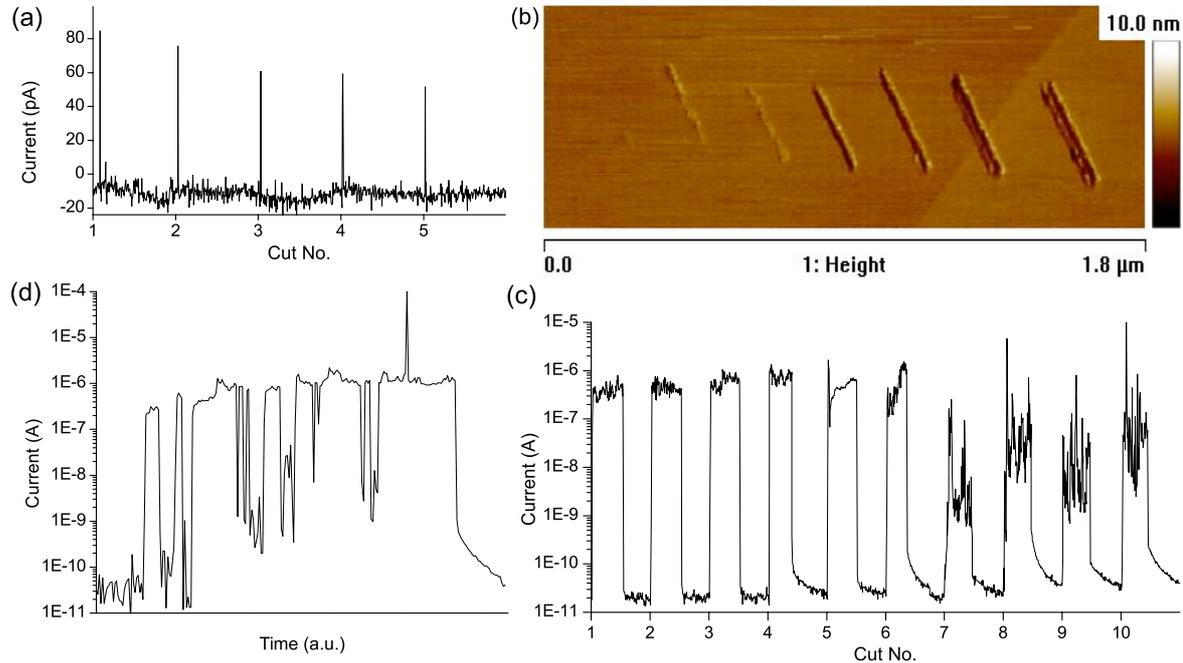
Figure 4.12: Tip-surface currents during lithography. (a) Successful, clean-edged cuts give only sharp pA spikes at the start of the cut. (b) & (c) A row of cuts and the associated currents, showing the transition from failed cuts with large currents to successful cuts, as the bias is stepped in 0.2 V intervals from -3 to -4.8 V. The non-zero current during cutting is associated with the high ridges. (d) Current for a cut where much dirt was deposited from the tip, including a large piece coincident with the mA spike.

a regular current of 1–10 μA since the tip is constantly in the more ohmic regime (Figure 4.12b–c). When a hole is blown out by a sparking event the current spike is 10 μA or more. It was important to have any imposed current compliance set to mA values, otherwise it acted to reduce the tip-to-surface potential drop during current spikes and aborted successful cuts.

When the bias was ramped with a tip static on the surface to measure the I-V characteristics, the current continually jumped between levels an order of magnitude apart. The higher value indicated more ohmic, electronic contact, and the lower regime was observed more clearly with another tip when the current was instead averaged over long periods held at static bias. The resulting data was a close fit to a Fowler-Nordheim curve for electron field emission, with the form $I = \alpha V^2 \exp(-\beta/V)$ (Figure 4.13) [221]. The current is very sensitive to changes to the tip; the final data point coincided with
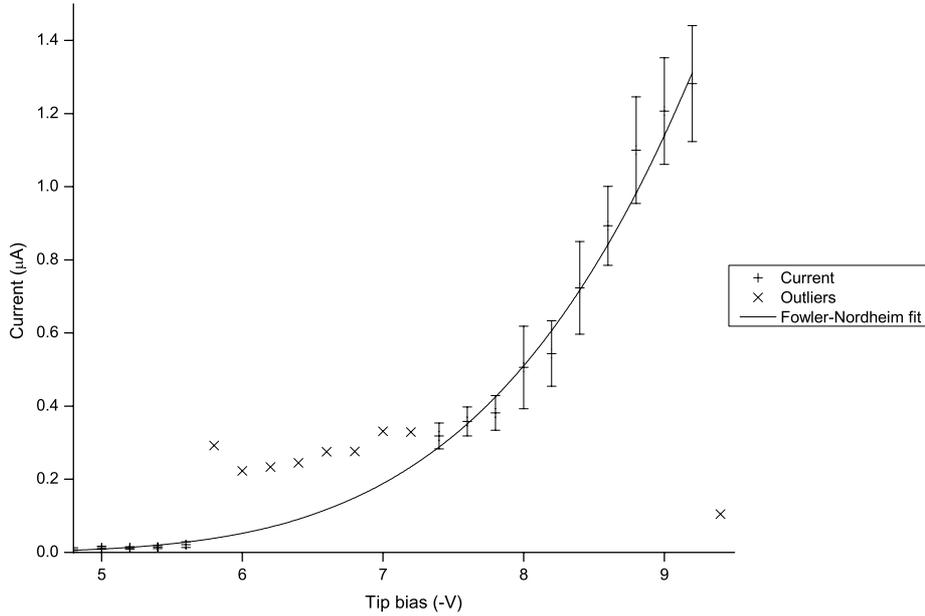
Figure 4.13: Field emission current in the absence of cutting. Current data with the tip held at a fixed bias on the surface. The fitted curve is a Fowler-Nordheim type relation: $I = 1.34V^2 \exp(-41.03/V)$. Each data point is averaged over a 15 s measurement, and there are several systematic outliers where the tip may have picked up surface dirt. The final data point created a hole in the flake.

a hole blown in the flake (up to this point there had been no effect to the surface), and several data points in the centre depart from the curve, presumably also due to some change in the dirt attached to the tip. With changing probe height the tip-surface current varies linearly with cantilever deflection throughout the regime where deflection is linear with height and the tip is in contact with the surface, therefore an increased force gives better electronic contact.

The improved reliability and diagnostic ability of this set-up permitted detailed assessment of the experimental parameters. The transition from no cutting, to ridges, to single layer cuts and then to deeper cuts was examined in detail (see Figure 4.12b); cuts were reliably producible with uniform single layer depth (Figure 4.14a). No surface change was seen applying positive biases up to 12 V. V_{th} rose after only a few microns of cutting, to the order of a volt, and repeated cuts at the same bias and humidity became narrower and eventually failed. A lower V_{th} was sometimes recovered after large lumps were dropped from the tip, or holes burnt in a sparking event. As well as changing V_{th}, a change in tip sharpness can give very different measures of cut widths
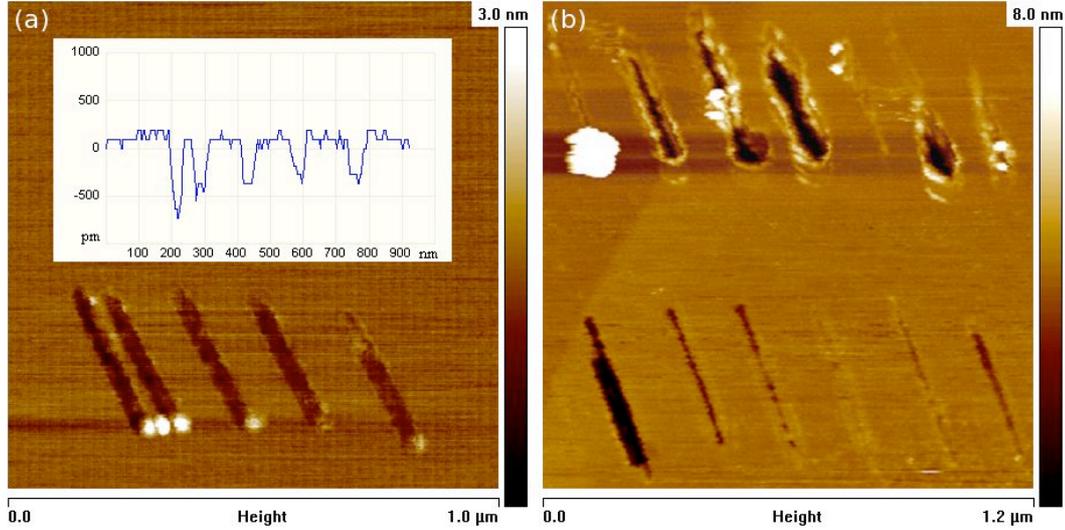
Figure 4.14: Cutting with different bias, force and AFM mode. (a) Lines drawn in contact mode with feedback, and bias -4, -3.9, -3.8, -3.7 and -3.6 V with (Inset) cut profiles. The middle three are all uniformly 1 atomic layer deep and the last is at V_{th}. (b) Top: lines drawn in tapping mode without feedback, and bias -12, -11, -10, -9.5, -9, -8.5, -8 V. The cuts are irregular and dirty, showing no clear trend with changing bias. (b) Bottom: lines drawn in contact mode with -5.5 V bias, and 'Z move' 120, 100, 80, 60, 40, 20 nm. Apart from the cut at greatest force there is no clear trend with changing tip height.

and depths, with implications for the detection of the lowest bias cuts.

The sparking events are a feature of drawing paths with feedback [192]. The scanning head maintains constant cantilever deflection, so as material is moved from beneath the tip the cantilever may be pushed into the hole if the feedback response is fast enough. This accelerates the downward growth of the hole. Without feedback the tip is moved at a constant height beneath its starting position (the 'Z move' distance) and relative to the atomically flat flake surface. Cuts without feedback had cleaner edges with less debris, and showed no current features other than the initial capacitive spike, since the tip no longer jumps into electric contact. Increasing the 'Z move' pushes the tip into the flake, with $Z = 0$ being the unperturbed surface. Cuts are weak and of roughly constant depth until $Z = 100$ nm, when there is a marked change (Figure 4.14b). This is equivalent to a deflection reading of 1 V. Continuing beyond this every additional volt of deflection adds only 1 atomic layer of depth, so within the available range of force the effect is much less than changing the bias. The

doped silicon tips apply 30 times as much force as the W_2C-coated tips for a given deflection, yet the latter typically gave smaller V_{th} by about 0.5 V, due to their sharper profile. This suggests that the main effect of the contact force is maintaining a short tip-sample distance.

Tip wear was better with non-coated tips, and the tip lifetime improved moderately as scanning parameters were better balanced to minimise the force. Further attempts at tapping mode cutting had the opposite effect. The margin between a narrow cut and a micron-scale hole was as little as 1 V bias, or 10 nm height change, and there was little correlation between bias and cut size. The best results were achieved without feedback (Figure 4.14b), giving $V_{th} = -8$ V. The effect of tip velocity was not studied in detail, and most cuts were done at 20 nm/s.

Keeping a constant tip bias and varying the humidity reveals a complementary effect. There is a threshold humidity beneath which no surface change is seen, then ridges appear, then cuts which gradually increase in size (Figure 4.15a). On a dirty surface with no tip bias lines at higher humidity clear proportionally deeper trenches in the dirt (Figure 4.15b). Comparing the deflection profile of a tip approaching graphite at low humidity and the same surface in the presence of water puddles, the puddles give a much greater snap-to-contact effect, due to the attractive force of the water meniscus (see Figure 4.8e–f). This effect is the cause of the deeper dirt lines, since the tip is pulled deeper into the soft dirt without changing its deflection. However, the sensitivity (the change in deflection with changing height once in contact) of both curves is almost identical, therefore an increased meniscus force alone is not responsible for the influence of changing humidity on cutting.

Ignoring tips which were known to be substantially worn or damaged, a plot of the voltage threshold for cutting with doped silicon tips against the relative humidity gives a reasonable correlation for increased V_{th} with decreased RH. Substituting the threshold for ridges in those cases where cuts were not found between all ridges increases the correlation by a further 30% (plotted in Figure 4.16). Since the cut and ridge thresholds were generally consecutive data points, the actual separation of values may be much less than the typical 1 V resolution of bias data points, and failure to detect the narrowest cuts is explained in Section 4.3.2. Therefore the formation of ridges probably indicates the start of the cutting process. Weakness in the correlation can be partly accounted for by the low bias resolution, as well as variations in tip width and velocity which are not accounted for here.

Amongst the features produced in these experiments were several 10 nm wide, single layer deep and 300 nm long cuts (Figure 4.17a), and a 5 nm wide and 100 nm
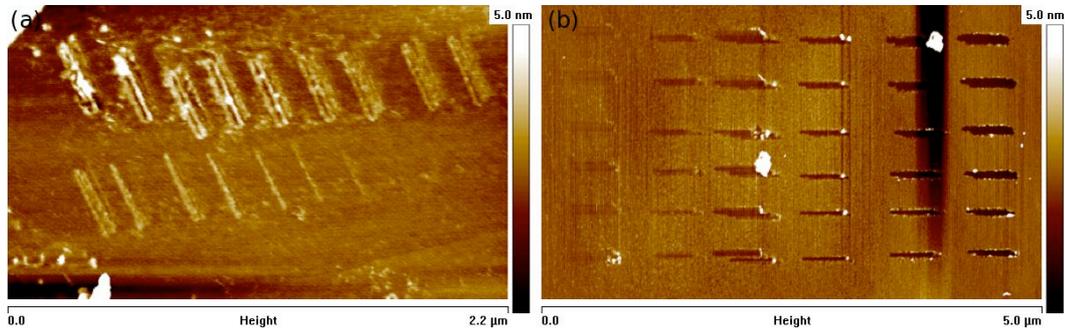
Figure 4.15: Effect of changing humidity on lithography. (a) Lines drawn with -4.6 V bias and relative humidity changing from 95% (top left) to 35% (bottom right). The lines transition through ridges to cuts, and deepen, in the same manner as when changing the tip bias. (b) Lines drawn on a dirty surface, which have not cut the graphite, at humidities (L–R) 18, 30, 41, 53, 66, 78%. The trenches deepen linearly with humidity.
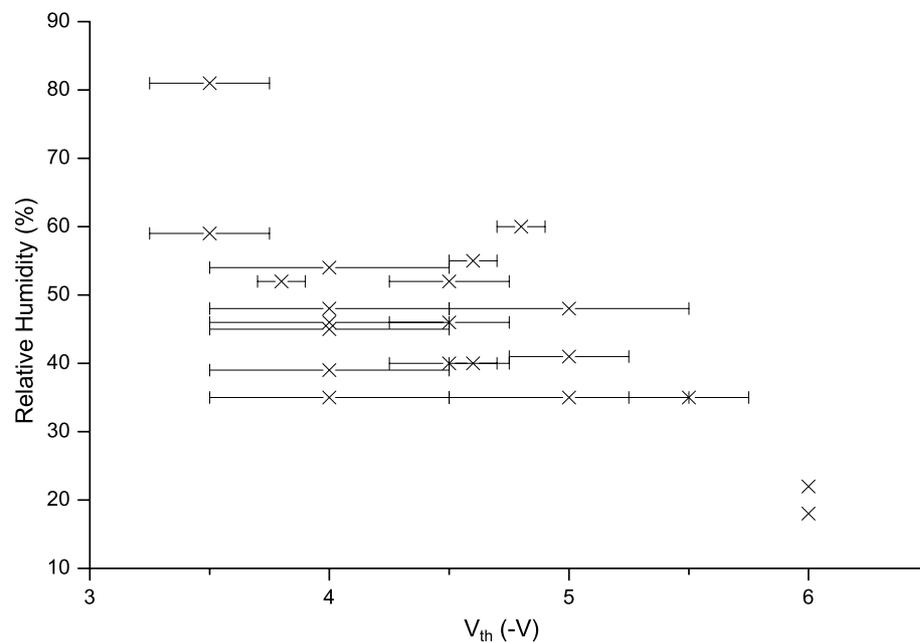
Figure 4.16: Correlation between humidity and the cutting threshold bias. The horizontal error bars represent the separation between sample bias measurement points. There is a clear trend towards lower V_{th} with higher humidity.
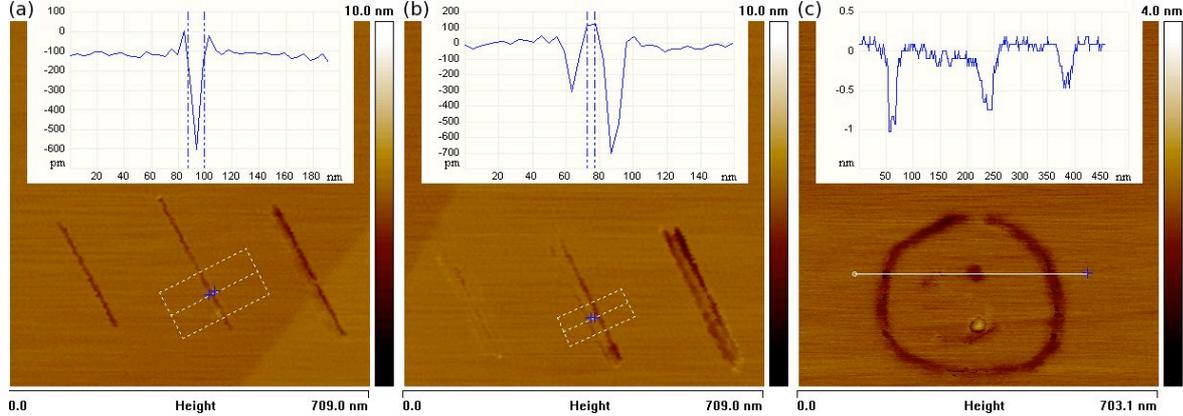
Figure 4.17: Fine features cut by AFM LAO. (a) 10 nm wide, single layer deep and 300 nm long cut (-4.4 V). (b) 5 nm wide and 100 nm long ribbon, between two cuts made by a double-ended tip (-4.8 V). (c) Circular path with a 12 nm wide gap at the top, and cuts 20 nm wide and 3, 2 and 1 layers deep in section (-5 V).

long ribbon, between the two cuts made by a double-ended tip (Figure 4.17b). The circle pattern in Figure 4.17c has a 12 nm wide gap at the top, and the cuts are 20 nm wide and 1–3 layers deep. After first finding V_{th} for a new tip, with a few short lines, the pattern was cut successfully at the first attempt. V_{th} for subsequent paths had dropped by around 0.5 V over 1 μm of cutting.

4.3.4 Lithography on SLG

Switching to SLG introduced an unexpected problem. Figure 4.18 shows a set of neat lines with bias stepping from -2 to -7 V. Unusually the trenches are all the same width despite the exceptional spread of voltages but they are also only 0.3 nm deep, and when the flake was imaged some weeks later there was no trace of any lines. This same effect was later seen with the other devices, and the surfaces involved were clean, so these were not lines in dirt. A later experiment annealing graphene (see Section 4.4) revealed that they were probably indentations, where there is a thin layer of water between the graphene and the substrate (which gives the extra measured thickness of SLG) and the tip pushed the flake into direct contact with the substrate, where it stuck for some period. These depressions were often hard to distinguish from cuts, and together with the ridges and puddles further obscuring lines, meant that the confidence level for successful lithography needed to be very high. The first set of SLG devices all developed broken ground connections, so very limited parametric analysis
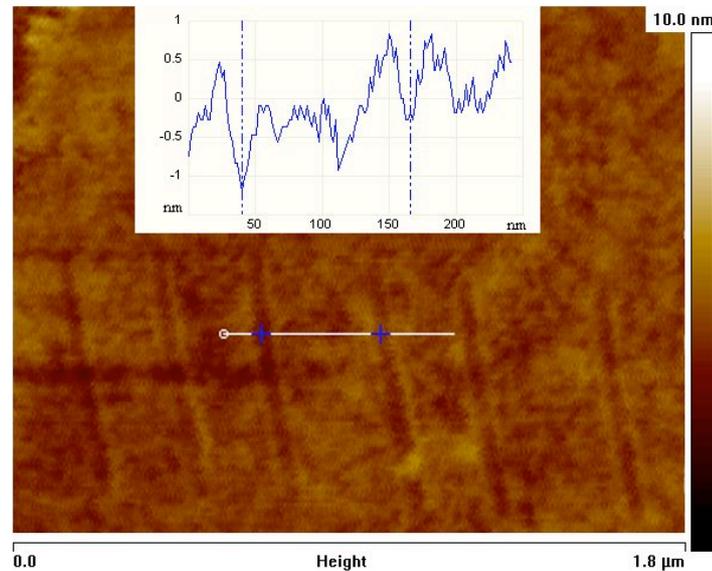
Figure 4.18: False cuts on SLG. The tip bias was (L–R) -2 to -7 V but the ground connection was broken. What appear to be cuts are depressions in the graphene surface, where the tip has pushed the flake into direct contact with the substrate, displacing the water layer beneath (device GPE04).

was achieved.

The first successful lines cut in SLG were only identified when, during contact scanning to clear some ridges, the tip caught in a hidden cut and tore the flake. A ridge appeared at -4 V, and when cleared revealed a cut less than 15 nm wide (Figure 4.19a). In further testing a ridge was sometimes seen at -3 V, but never a cut beneath it, and -4 V consistently gave a cut. Cut widths were more stable on SLG over a higher range of voltages - no blast-like holes were seen.

Figure 4.19b again shows depressions into the water layer and also an instance where the tip has cut only at the edges of a wide line, leaving the graphene between intact. Further to the right a similar piece can be seen rolled up at the end of the cut. That some carbon is not converted to gaseous oxide may provide an explanation for dirt collected on the tip even with cleaned surfaces, and for the material sometimes piled up around cuts, whose volume is generally comparable to that of the cut. This is in comparison to lumps dropped from the tip, which are often much larger than the hole volume.

Given the uncertainty of the success of a cut from imaging alone, reliable confirmation is given by measuring the change in resistance of a cut flake. Figure 4.20 shows
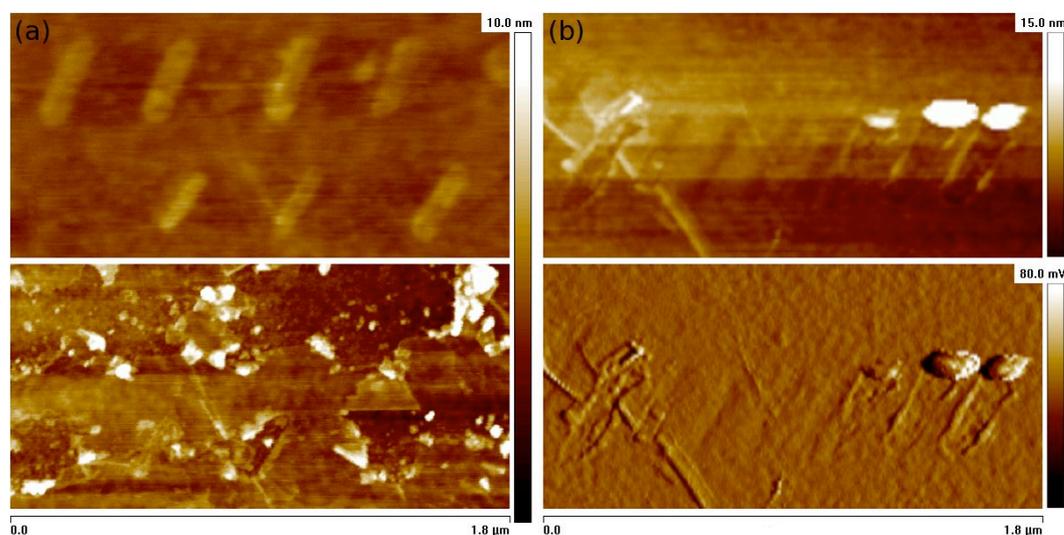
Figure 4.19: AFM LAO lithography on SLG. (a) Top: lines drawn at -6 V (top row) and -4, -5 and -6 V (bottom row) gave ridges. Bottom: after contact mode scanning the flake tore, revealing cuts beneath all the ridges (device G01b). (b) Lines drawn at -5.5, -4, -4.2, -4.4, -4.6, -4.8, -5, -5.2 V. V_{th} is -4.8 V - beneath this the tip has pushed the flake onto the substrate. At -4 V a short section has been cut only at the line edges, leaving the graphene between intact, and suggesting that the lumps at the ends of the successful cuts are rolled up graphene (G08).
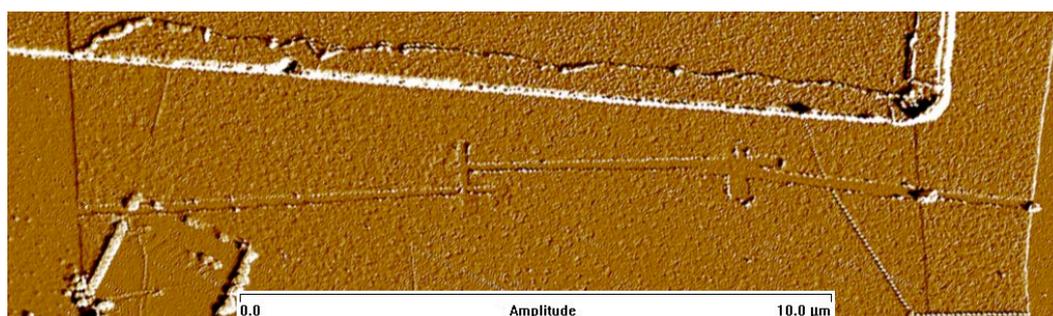
Figure 4.20: Large SLG flake divided by a cut. A series of 100 nm wide intersecting cuts (at -7 V) divide a SLG 16 μm wide flake in half. The flake is very dirty, giving many lumps along the path's course and around the area cleared for the cuts shown in Figure 4.19b. At the right the flake is rolled into a section 3–4 layers thick, and at the top is a metal electrode. Other line features are creases in the flake (device G08).
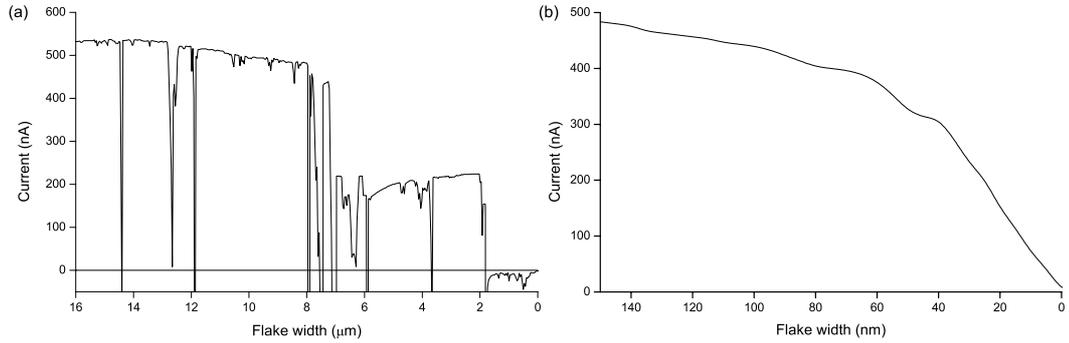
Figure 4.21: SLG flake transport during cutting. (a) Current between two electrodes on a previously annealed flake, while it is cut in half. The current drops irregularly, and the tip is making intermittent electrical contact with the flake, overlaying a profile of reverse current spikes and noise. (b) The same experiment on a pristine flake. Once the remaining flake width becomes comparable to the cut width the current drops linearly. In this instance there is no current flow through the tip.

a 100 nm wide (-7 V bias) cut across the whole 16 μm width of an SLG flake (G08), which broke the electrical connection between metal electrodes on either side. The chip package holder built for the AFM by David Tyndall allows current through the flake to be measured during cutting. A repeat of the previous cut between another pair of electrodes gave a very irregular drop in current between them, changing in large jumps rather than smoothly with the reduction in flake width (Figure 4.21a); the flake became completely insulating when a few hundred nanometres remained uncut. This flake had been annealed at 400℃, so may have already suffered damage to the lattice. Additionally the tip was making intermittent electrical contact with the flake, such that current spikes and noise driven by the negative tip bias were overlaid on the flake transport. The same measurement on a pristine flake (G07) gave a smooth decrease in current, which dropped linearly once the remaining flake width w became comparable to the cut width l at around 40 nm. Fitting a curve of the form $R = \rho l/w + R_0$ gives a resistivity ρ of order 1 Ωm.

4.3.5 Physical manipulation of graphene

An AFM tip can be used for other manipulations of single and multi-layer graphene, such a peeling, lifting and folding. Others have reported shifting multilayer sections by hitting a tip laterally against a graphite island [222], and arbitrary folding of single
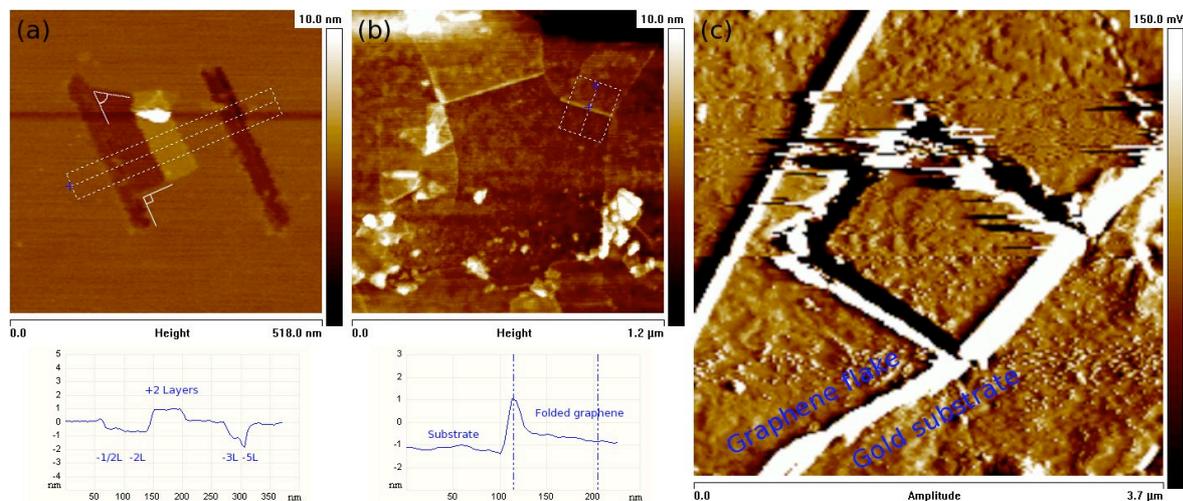
Figure 4.22: AFM manipulations of graphene. (a) A fold resulting from accidental tearing while scanning a cut edge. Also shown are 60 and 90° angles where the fold coincided with the lattice axes, and a cross section with the number of layers folded and cut. (b) A deliberate fold of SLG; sections of flake have torn, peeled off the substrate (central region) and folded. The curvature of the fold is greater than the bilayer thickness. (c) A square of FLG (on a gold substrate) was cut out, but could not be peeled or moved using the AFM tip.

layers by close vibration of an STM tip [223]. Such folds are stable because the interplane energy is larger than the repulsive folding energy, while tearing is energetically favoured along symmetry directions [224]. Stably folded graphene would be a route to fabricating a bilayer device, and selective peeling could give multilayers of chosen thickness and perhaps enable removal of single layers for transfer elsewhere. Controlled folding and peeling were briefly attempted here.

Small folds are often induced when scanning in contact mode after cutting, where the tip catches against the edge of the cut with enough force to make a small tear, and the fold is carried over with the tip as it continues (see Figure 4.9c); in one instance the angle of the cut aligned closely with a symmetry axis in the lattice and the resulting fold had angles of 30, 60 and 90° (Figure 4.22a). Layers can also be peeled with the tip in hard contact, as used to achieve a clean surface (Section 4.3.2). To fold SLG into a bilayer with a more controlled shape two parallel cuts were made, and then scanned perpendicularly in contact mode with moderate force. Figure 4.22b shows the two folded sections which resulted, the larger around 0.4 μm by 0.4 μm. More regular shapes might be achieved by cutting along the symmetry directions, but these are

not identifiable with an AFM. Cutting lines using only physical force is not practical, because unless alignment with the lattice is perfect the flake will tear and fold.

In an attempt to peel or move layers, deep cuts were made at -12 V on a thick flake to outline a square at its edge (Figure 4.22c). The tip was then positioned above the innermost trench, brought into hard contact with the surface, and moved with no feedback towards the edge of the flake. This should have applied a large lateral force to the island but repeated attempts had no effect. Possibly the tip was riding up the diagonal sides of the trench and lifting out instead of catching the edge.

4.4 Annealing graphene

An experiment was carried out to test for changes to the structure and morphology of graphene after heating in air. Spanning the range of elevated temperatures that might be encountered during processing, flakes were observed by AFM and Raman spectroscopy to detect changes to their surface topology and lattice integrity. Changes in transport properties were studied with magnetoresistance measurements.

4.4.1 AFM and Raman spectroscopy

Two uncontacted flakes (G10a and G10b) were baked on a hot plate at 40℃ intervals for 15 minutes, to a maximum temperature of 440℃. After each bake the flakes were immediately imaged by AFM then sent to CAPE and their Raman spectra taken by Vittorio Scardaci.

The AFM images revealed a striking transition between 100℃ and 180℃. The mean height of an SLG (flake G10b) edge relative to the substrate was 1.53 nm, 1.52 nm and 1.47 nm after annealing at 20, 60 and 100℃ respectively, but 0.48, 0.45 and 0.43 nm at 180, 220 and 260℃. At 140℃ the flake was divided into patches at two levels, some at an average of 1.36 nm and the rest around 0.42 nm (see Figure 4.23a-c). Taking all data together, the mean height before the transition was 1.51 ± 0.22 nm and after was 0.45 ± 0.11 nm. The layer height in graphite is 0.34 nm, so the mean final separation between the graphene and the substrate was only $\sim$0.1 nm, while the measured RMS roughness of the SiO_2 was 0.2 nm, and of the graphene on top of it was 0.3 nm. This substantial overlap between the two surfaces means that the graphene must be conforming exceptionally tightly to the substrate topography, and the other effects that might separate them such as intrinsic ripples in the graphene, adsorbed molecules and van der Waals forces are essentially negligible.

4. AFM LITHOGRAPHY OF GRAPHENE

In contrast, a step edge where the flake folded over to give a bilayer region was on average 0.37±0.08 nm below 140℃ and 0.40±0.05 nm above, therefore it had remained unchanged. The second flake (G10a) showed a height change of 1.04 nm between the same temperature steps, consistent with the first. The final height of 0.89 nm revealed it to be a double layer flake, where the Raman result had suggested it was a single layer; the flake was smaller than the spot size such that the significant edge signal caused uncertainty regarding its thickness. A third, contacted flake (GPE04) on a different chip was annealed at 20℃ intervals from 60℃ to 160℃; the lower height regions appeared at 120℃, becoming more defined at 140℃ (see Figure 4.23d-e).

The coincidence of this transition with the boiling point of water suggests that the height change is due to the evaporation of a thin water layer. The water puddles on top of graphene seen previously (Section 4.3.2) had a wide range of heights, and a lower bound of 1.3 nm, but the height change at the transition when heated is always 1 nm, and the patch edges are consistent between scan lines. The water being removed is therefore underneath the flake, and it is this which gives the enhanced step height for SLG on silicon. Such a thick water layer on top of the flake would only be detected by tapping mode imaging, and would not have given an enhanced single layer height from contact mode scanning. All the lower height regions connect to the edge of the flake and there are no isolated islands at the lower height, therefore the water is escaping at the flake edges, and the flake drops onto the substrate as it leaves.

A SiO_2 surface has many silanol (Si-OH) groups [225], which are active centres for water adsorption such that the surface is hydrated under normal conditions. Most of this water can be removed by heating to 120℃, agreeing with the raised evaporation temperature seen here, but complete removal of the silanol groups requires temperatures of up to 1000℃ [225], hence the surface re-hydrates after moderate heating. The presence of water has been inferred beneath carbon nanotubes on silicon, by transport measurements and subsequent annealing [226], so it is reasonable to find it beneath mechanically exfoliated graphene.

Images showing blobs appearing on FLG subsequent to bonding were now re-examined (see Figure 4.7). The bulk flake had in fact dropped in height by around 1 nm, and the blobs were at the original flake height. For gold wire bonding the chip is heated to at least 120℃ to aid the adhesion of the bond balls. The water under the flake was therefore leaving during bonding, but the greater size and stiffness of the FLG meant that some remained trapped in pockets. Under a more flexible region of SLG the bubbles are smaller and more numerous (Figure 4.7f).

A second feature which appeared with the height transition at 140℃ was a large
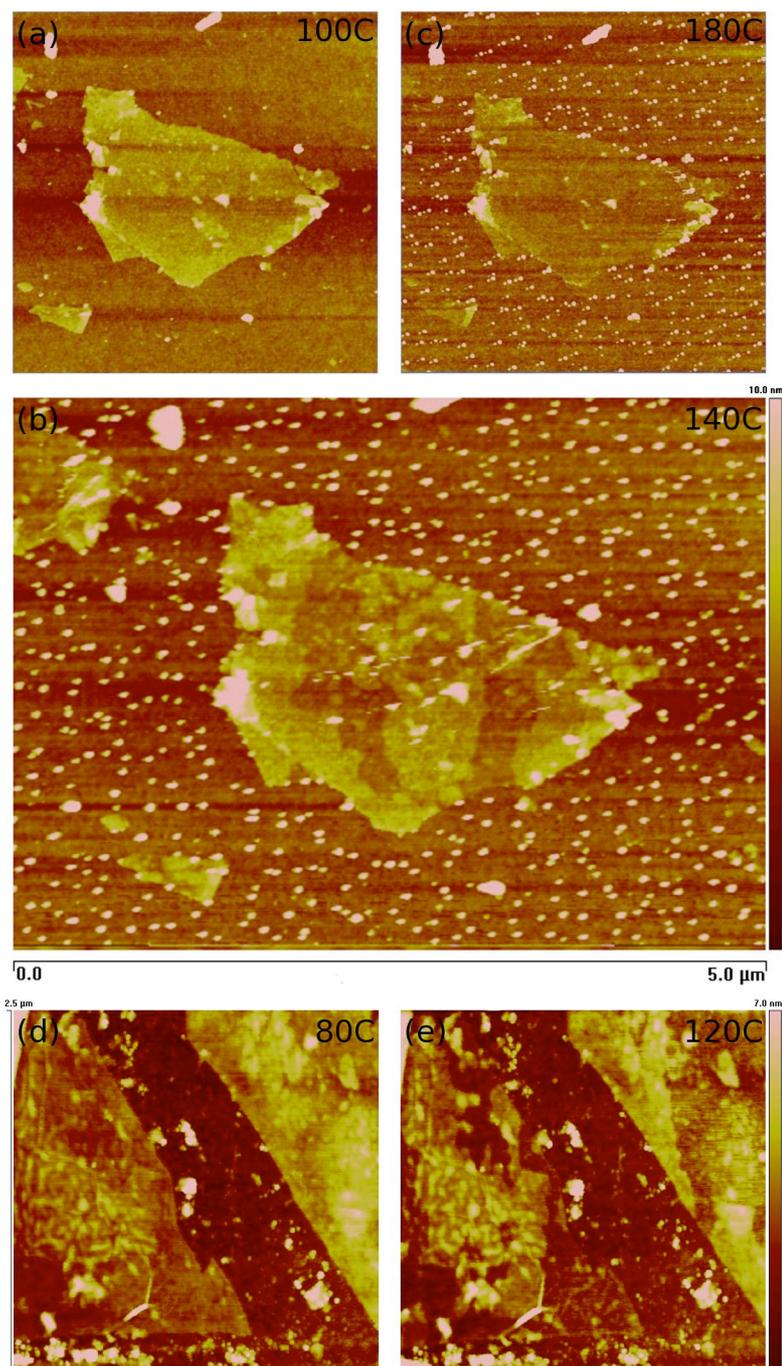
Figure 4.23: Height transition in flakes annealed above 100°C. Images after heating a flake to (a) 60°C, (b) 140°C and (c) 180°C - in the second image two height levels appear and in the third the whole flake is at the lower height. (d) Another flake heated to 80°C and (e) 120°C - the flake is in two parts with a central gap; in the second image many regions of the left-hand part have dropped to a lower height.

number of small round lumps, 5–10 nm high, scattered across the substrate but not the graphene. They were highly mobile, being pushed by the tip even in tapping mode and collected on the edge of the scan region or the flake edge, and even more mobile on the graphene, where they are evidenced only by thin streaks in the image as they are moved off the flake. They are consistent in size and behaviour with small water droplets observed by Luna et al. [220] on graphite (and may also have been recorded in Figure 4.8b). The latter were similarly mobile and clung to step edges, and were achieved by removing a macroscopic quantity of water from the surface. The source of the droplets here is more ambiguous.

Their density is around 40 per μm^2 which, assuming a molecular density similar to bulk water (30 nm^{-3}), gives an average substrate surface density of order 1 water molecule per nm^2. Under atmospheric conditions most surfaces are covered by a few molecular layers of water, but even a monolayer has a molecular surface density closer to 10 nm^{-2} (again neglecting surface tension, which makes both estimates too small). The droplets could therefore be either the remainder of the surface water, drawn into beads by surface tension after too much has evaporated for it to remain continuous, or it could be the first water to return to the surface after cooling. There are also water molecules and OH$^-$ groups present and diffusive within silicon dioxide, but the concentration in the oxide is 10^{-5}–10^{-2} nm^{-3} [227], so formation of the droplets would require evacuation of all water within the 300 nm thick oxide, which is improbable given the high activation energy and a diffusion rate of only 0.01 nm/s [228].

Raman spectra of the large flake (G10b) showed no change until 180℃, when there was very slight broadening of the G' peak. At 220℃ there was definite broadening and at 260℃ the peaks changed substantially to resemble the signal for the small flake G10a; the heating had introduced structural defects into the lattice such that it gave a strong signal for edges. There was no change associated with the height transition after 100℃. The Raman data was never returned from CAPE, so cannot be reproduced in full here.

4.4.2 Transport measurements

Flake G10b was then contacted with two Ti/Au electrodes, bonded into a chip package and measured in an Oxford Instruments cryostat with a top-loading variable insert in a magnet dewar. The cryostat can reach a temperature of 1.4 K by pumping on a small chamber of liquid helium-4 to reduce its vapour pressure; the thermal expansion of the helium as it evaporates cools the gas surrounding the sample. A needle valve bleeds

more helium into the sample space, which is surrounded by a liquid helium bath, a vacuum jacket and a liquid nitrogen bath. The superconducting magnet is capable of fields up to 10 T.

Immediately after wire-bonding (which heats the chip to 120℃) the conductance minimum was found at $V_g \sim 15$ V. This identifies the Dirac point and the transition from p-type to n-type carriers. Two weeks later it had shifted to 45 V, consistent with a gradual re-admittance of water, which dopes graphene p-type and shifts the Dirac point to positive gate values [27]. The transition is unusually broad and shallow, which may reflect damage done to the graphene though heating (see Figure 2.6 and Section 5.4 for comparison with untreated flakes).

Gated magnetoresistance data at 1.38K (Figure 4.24) shows aperiodic oscillations with respect to both changing magnetic field and back gate, and a drop-off in current near zero field. The latter is evidence of weak localisation [5]. This is a quantum interference effect whereby in zero magnetic field time-reversal invariance guarantees that trajectories forming a closed loop in a 2D system have equal probability amplitudes for clockwise and anticlockwise propagation. The resulting constructive interference enhances the backscattering probability, and gives a positive correction to the resistance. A weak perpendicular magnetic field breaks the time-reversal invariance by introducing a phase difference between the two trajectories, restoring conductance. Theoretically the chiral nature of electrons in graphene favours weak anti-localisation, giving a negative resistance correction. However, in a real graphene sheet scatterers and edges suppress this and restore conventional negative magnetoresistance; the long range propagation of electrons in the real, disordered material does not manifest their chirality [229].

The current oscillations are 'universal conductance fluctuations' [230]. Quantum interference also results in significant sample-to-sample fluctuations in metals if the sample size is not much larger than the phase coherence length. These are a function of changing the impurity configuration, which is equivalent to changes in the magnetic field or Fermi energy. They are always of order $\delta G \approx e^2/h$ at $T = 0$, when phase coherence is maintained over the entire sample, irrespective of the number of modes. At finite temperature the magnitude of fluctuations is reduced due to a finite phase coherence length and thermal averaging.

Taking a parabolic best-fit to zero field gate sweeps as the mean conductance $\langle G \rangle$ (Figure 4.25), the magnitude of the fluctuations is given by their variance $\mathrm{Var}(G) \equiv \langle (G - \langle G \rangle)^2 \rangle$ which is $0.18e^2/h$. This is an order of magnitude larger than a previously reported value measured by Horsell et al. [231] for a comparable flake geometry at
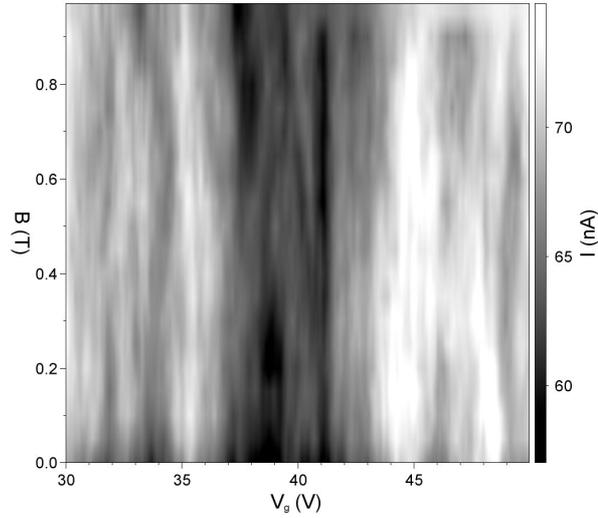
Figure 4.24: Magnetoresistance of annealed graphene. Two-terminal current measurement at 1.38 K through a graphene flake heated to 440°C, with a varying back gate and perpendicular magnetic field. The data shows oscillations relative to both swept variables, and a drop in current at low field.

the same temperature, and is for a two-terminal measurement, so would be larger still with the contact resistances excluded. A better fit applies the parabola over a narrower region, extended by straight lines, but the calculated magnitude of the fluctuations would not differ substantially. The fluctuations are smoothed out at higher temperatures. There is no evidence of Shubnikov-de Haas oscillations or quantum hall plateaux, normally visible above conductance fluctuations in pristine graphene [31, 187], so the 2D nature of the flake has evidently been disrupted by the heating. A large number of resulting structural defects make many elastic scattering centres, giving significant phase coherence across the flake and the large fluctuation amplitude.

After a further two weeks transition the Dirac point had moved to $V_g \sim 90$ V and the flake profile was higher, confirming that the water layer between the flake and the graphene had returned. The rate of return is slow but it is again unlikely that sufficient water is diffusing through the oxide, rather that it enters at the edges of the flake. Water on top of the graphene and substrate surfaces probably returns at a much faster rate, but it cannot be detected by AFM, nor can the effect of the two bodies of water be distinguished in this experiment. Heating the flake again to 150℃ returned the Dirac point to 45 V and lowered the height profile, therefore the process can be cycled repeatedly.
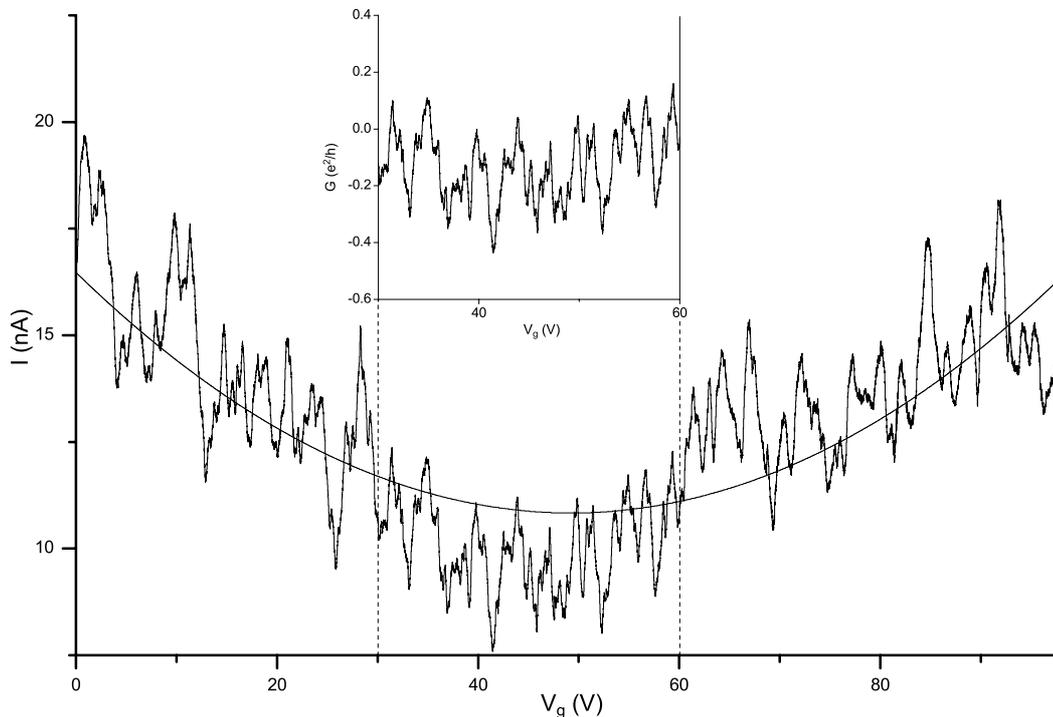
Figure 4.25: Universal conductance fluctuations in graphene. Gated 2-terminal current measurement (at 1.38 K and 0.2 mV bias) through a graphene flake heated to 440°C. (Inset) Fluctuations in the conductance after subtraction of a parabolic background, measured in units of e^2/h.

4.5   Discussion

In summary, the main observations of these experiments are that detectable holes are made in graphite and graphene when a negative bias greater than a threshold value V_{th} is applied to a tip in contact with the surface. Larger biases give wider and deeper holes and the flake must be connected to ground. There is no dependence on the tip-surface current, and no holes are seen below V_{th} even for a very large current, while with successful cuts any current is usually below 1 nA. With failed cuts the current follows a Fowler-Nordheim relation for field emission. Lower tip voltages leave depositions rather than holes, and the latter often have slightly raised edges. Cut widths vary exponentially with bias and V_{th} is larger for worn tips. No holes were achieved with positive bias, and tapping mode hole generation was too unstable to be reliable. Increased humidity lowered V_{th} and gave larger holes such that for any given bias there is also a humidity threshold. The range of V_{th} was -3 to -6 V. There is a 1 nm
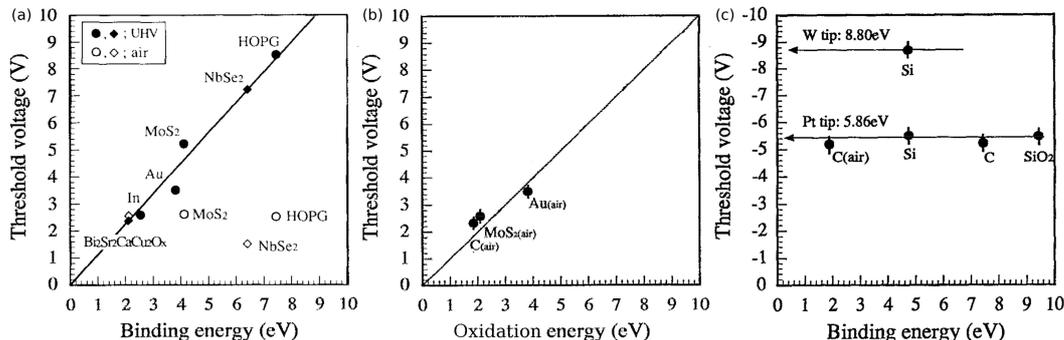
Figure 4.26: STM lithography thresholds for various materials. (a) V_{th} (applied to sample) vs. binding energy - the values are linear and nearly equal in UHV, but show no relationship in air. (b) Instead V_{th} in air is approximately equal to the oxidation energy. (c) With positive tip bias, V_{th} is independent of the substrate material and equal to the binding energy of the tip material (from [201, 232]).

thick water layer beneath a flake, which can be removed by annealing to around 120℃, and cutting is prone to introducing water puddles on top of the graphene surface.

4.5.1 Macroscopic mechanism

Various mechanisms for hole generation are proposed in the literature. Thermal oxidation [233] is widely discounted with STM LAO [198, 199, 205], since the maximum temperature change is < 5 K, which is implausible for oxidising graphite. Wang et al. [234] suggest electron diffusion, based on the hole profile, but their hole is only three atomic layers deep, so the data sample of widths is extremely small. Kondo et al. [201] compare LAO in UHV and ambient air on various materials. In UHV V_{th} is not just proportional to but nearly equal to the sample binding energy (7.43 eV for graphite), therefore the mechanism is sublimation of surface atoms by tunneling electrons (SITE). In air this relationship is lost, and V_{th} is around 2 V for all materials, but is now approximately linear with the oxidation energy [232] (Figure 4.26). The authors eliminate the possibility of field evaporation, according to which hole depth would be proportional to field strength, and ascribe a mechanism with a chemical reaction induced by tunneling electrons (CRITE). Albrecht et al. [199] have shown that it is specifically water that is crucial; hole generation failed in pure O_2 and pure N_2, and was as effective in low pressure water vapour (air equivalent) as in atmospheric air. The process is therefore one of two oxidation reactions involving water [203]:
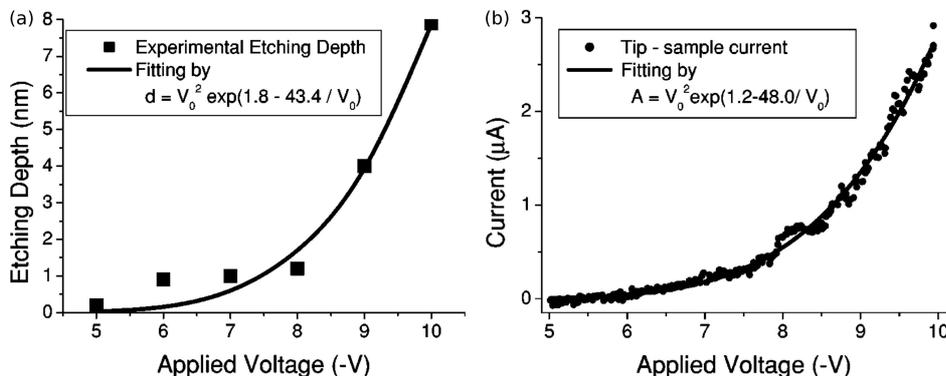
Figure 4.27: **A model for current-driven lithography.** Kim et al. [90] observe a Fowler-Nordheim field emission current for a biased tip on graphite, and use this to justify a current-driven process for the lithography.

$$C \text{ (graphite)} + H_2O \text{ (liquid)} + 1.82 \text{ eV} \rightarrow CO \text{ (gas)} + H_2 \text{ (gas)}$$
$$C \text{ (graphite)} + 2H_2O \text{ (liquid)} + 1.85 \text{ eV} \rightarrow CO_2 \text{ (gas)} + 2H_2 \text{ (gas)}$$
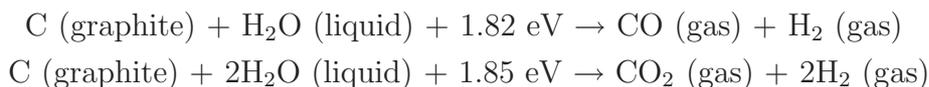
However, the current-independence, occasional cutting offset from the imaging tip and the dependence on tip sharpness observed here imply that the process is a function of electric field strength, not of bias-driven electron flow. Similarly Abe et al. [198] find a corresponding linear relationship between V_{th} and the tip-surface distance. They also describe capacitative current peaks at the start and end of the voltage pulse.

The authors of the previous AFM lithography of HOPG [90] support the CRITE mechanism, but their evidence is weak. They measure a current with a Fowler-Nordheim voltage dependence between the tip and graphite, to several volts beyond their measured V_{th}, but mention no simultaneous effect on the surface, nor do they record the current during cutting (Figure 4.27a). The exponent constant β is a function of the tip-sample distance and the material work functions. Their fitted $\beta = 48$ is within reasonable agreement to the value of 41 found here (see Figure 4.13). The latter was observed in the absence of cutting and, given that the substrates and distances are probably comparable, suggests that the authors are observing the same current effect. However, they claim that a current is nonetheless responsible for the cutting and go on to fit a Fowler-Nordheim type relation to the variation of cut depth with bias (Figure 4.27b). Three of their six data points are for cuts three layers deep and the fourth averages less than one layer deep, so their fit is reliable only for the remaining two. It is substantially different from the true best fit, and yet they claim a scientific coincidence with the constants in their fit to the current. Furthermore, a simple exponential fit
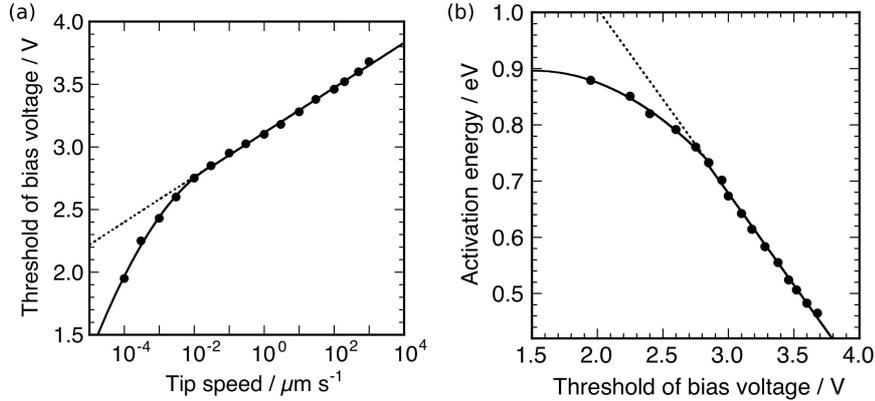
Figure 4.28: Relationship between tip velocity and threshold bias. (a) With STM lithography of graphite Hiura [200] find that V_{th} changes exponentially with the tip velocity during cutting and therefore deduce that (b) V_{th} is proportional to the activation energy of the reaction process.

to an estimation of their data gives a correlation coefficient 10% higher than the best Fowler-Nordheim fit.

A more plausible explanation is given by Hiura [200]. They propose the same chemical reaction but suggest that the electric field, rather than an electron flow, drives the reaction. They observe an exponential relationship between bias voltage and cut depth, and also with the volume of carbon removed, though they claim a linear relationship with the cut width. The latter contradicts an exponential volume relationship and is contrary to the observations here, but may arise from their narrow data sample. They associate the exponential relationship with a kinetic reaction rate constant which is exponential in the activation energy; the bias voltage enters the exponent because it modifies the activation energy. This interpretation is consistent with the minimum cut width being related to the tip apex width, since V_{th} is the bias at which a certain depth of carbon has been removed within the time taken for the tip to pass - it is a function of the reaction rate. For a given tip width the amount of carbon removed at V_{th} is always the same, so the product of rate constant κ and the interaction time is a constant C:

$$\frac{\kappa W_0}{v} = C$$

for a tip with diameter W_0 moving at speed v. The Arrhenius equation for a rate constant is

$$\kappa = A \exp\left(-\frac{Q}{kT}\right)$$

with a constant prefactor A and activation energy Q. Combining these two relations gives

$$Q = kT \left\{ \ln \left(\frac{AW_0}{C} \right) - \ln v \right\}$$

The authors also find that V_{th} is proportional to the logarithm of v, therefore the activation energy is inversely proportional to V_{th} (Figure 4.28). Fitted to their data the complete relation extrapolates to a Q at zero bias of 1.7 eV, very close to the proposed oxidation mechanism. With consolidated constants α and β the rate constant is

$$\kappa = \frac{\alpha}{W_0} e^{\beta V_{th}}$$

The authors neither investigate nor consider the role of water in the reaction rate. If water is the rate limiting reactant then the rate will be the product of the water concentration and κ. Relative humidity RH is a measure of the partial pressure of water in air, which is proportional to its mole fraction. The complete rate equation is suggested to be

$$r = \frac{\alpha' RH}{W_0} e^{\beta V_{th}}$$

For a fixed v the reaction time is always the same, therefore W_0/RH will be proportional to $\exp(\beta V_{th})$. An exponential fit to Figure 4.16 does increase the correlation coefficient by a further 15%. W_0 was measured for very few data points, and they do not show a clear relationship with V_{th}, but this is not unreasonable given the variability of other dependent variables and the measurement uncertainties. Except to the extent that it reduces the tip-sample distance, increased contact force would not be expected to have an effect on the rate. Furthermore, graphite is elastic to an STM tip pressing about 100 nm [199], which is also the largest 'Z move' that gave a significant change in cutting.

4.5.2 Surface water

Although classed as hydrophobic, graphene in ambient air is covered in a continuous water layer even at humidities as low as 10%, which increases in thickness linearly with increasing humidity [235]. In AFM studies Luna et al. [220] detected two distinct phases with different thickness. The lower, at 2 nm or shallower, is a short range ordered hexagonal ice-like lattice of bilayers, constrained by binding to the graphite surface - there is only 5% difference between the two intrinsic lattice spacings. This is the continuous layer, not usually observed during, but important for AFM imaging. They had cleaved the graphite in an airtight chamber, and hence were able to
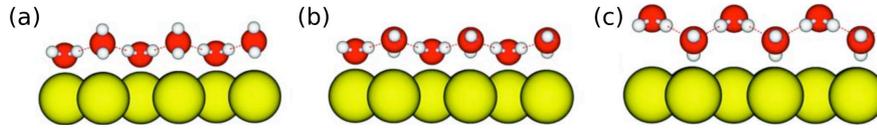
Figure 4.29: Water bilayer structure on Ag(111). Yellow, red and grey spheres indicate Ag, O and H atoms, respectively, shown in the preferred orientations under (a) positive electric field, (b) zero field and (c) negative field (from [237]).

detect this layer forming. At high humidities a phase up to 5 nm thick appeared in patches, which moved and coalesced on a scanning time scale, and which corresponds to the puddles observed here. They allow many hours for the surface to equalise after a change in humidity, whereas cutting measurements here were taken instantaneously, with humidity changing by 50–60% in an hour. This may help explain the lack of correlation in Figure 4.16 and also why puddles were not observed here forming spontaneously at high humidity. They also observe a significant tip oscillation phase shift between the two depths [236], confirming the impedance change inferred here.

Studying water on a hexagonal Ag(111) surface, Zhao et al. [237] predict that the molecules in an ice-like bilayer will orientate according to an applied electric field, due to their internal dipoles. With no field or a negative tip bias they will present oxygen towards the carbon surface, but with a positive bias they present hydrogen (Figure 4.29). The mean separation of the oxygen atoms from the silver surface is also 25 % (1 Å) greater under positive bias. This offers an explanation for why hole generation fails with positive tip bias, within the range measured, despite the field energy being sufficient to drive the reaction. Additionally, X-ray scattering has shown that the molecular density of water in the innermost layer on a positively charged silver surface is nearly double the negatively charged case [238]. In the case of graphite this arrangement would present two oxygen atoms to every three carbon atoms on the top surface. Given that the reaction also proceeds through several layers, it cannot be sustained by the water initially present, and subsequent physisorption of water onto the surface must be rapid to maintain the reaction rate. Hence some correlation with humidity is still seen in Figure 4.16 despite the lack of equilibrium at the surface. The oxygen:carbon ratio at the surface favours the first oxidation equation above, since it requires only one water molecule to each carbon atom.

4.5.3 Depositions and ridges

Apart from the exceptions discussed below, the only reported observation at positive voltages is the deposition of large lumps, several nanometres in size and distinct in form from the shallow ridges and mounds seen at negative bias [201, 202, 239]. V_{th} for these depositions is independent of substrate and atmosphere, but roughly equal to the sublimation energy of the tip material [232] (Figure 4.26c). The tip is also sometimes seen to degrade afterwards [202].

Mounds at low negative tip bias and ridges around holes are widely reported [199–201, 204, 205, 240, 241] and generally assumed to be a solid state carbon oxide. The height of raised hole edges, as here, is always ~ 3 Å [199], which is too shallow and regular to be material from the tip [205], and other protrusions range from 3 Å to a few nanometres. Graphene oxide layers are between 6 and 8 Å thick, 3–4 Å thicker than pure graphene [242, 243]; the oxygen is bonded out of plane as C–O–C in parallel with existing C–C bonds, and the depth is further increased by distortion of the hexagonal carbon lattice [244]. In air the protrusions resulting from lithography are very volatile, lasting from as little as 1 s [240] to a few minutes (as here), but submerged in water they are stable to continued imaging for over an hour [204]. Applying a small voltage (0.2 V) directly to a protrusion converted it into a single layer deep hole [204], so they are a metastable intermediate.

Potential cycling of an aqueous KNO_3 electrolyte covering a graphite surface causes an irreversible oxidation reaction involving the top two layers above -1.5 V, with an accelerating rate at -1.9 V [245]. The product is surface blisters tens of nanometres high and hundreds of nanometres in diameter. They are hollow and can be burst, revealing shallow craters in the graphite beneath, and when imaged dry are only a few nanometres high. The smallest are 3 Å high [246]. The blister roof gives a strong EDX signal for oxygen but atom-resolved AFM images show a hexagonal lattice identical to native HOPG, and the crater floor is pure graphite; the oxide is therefore on the interior of the roof. The blister volume can be accounted for by the gas pressure of the oxidation products. After continued cycling all trace of the lattice on the blister roof is lost and it fractures spontaneously due to oxide formation throughout [246]. The process is substantially more rapid cycling to -2 V than to -1.8 V. In a complementary experiment Alsmeyer and McCreery [247] observe that molecular intercalation in the graphite surface, as a result of potential cycling with weak aqueous acids, gives an increased Raman 'D' band due to the introduction of microcrystallite edges, and that intercalation always precedes lattice damage by oxidation. This is the same change

in Raman signal as seen in Section 4.4 when annealing the graphene flake, confirming that the latter was being oxidised.

4.5.4 Molecular mechanism

Though the protrusions resulting from LAO have not been identified directly the above evidence cumulatively supports the hypothesis of an unstable graphene oxide intermediate. Together with the chemical reaction proposed above, the evidence for an ordered water layer and the established role of electric field strength, it is possible to complete a detailed picture of the oxidation mechanism. The lowest V_{th} reported on pristine graphite is -2 V [200] but at existing defects and edges it drops to -1.5 V [202]. Similarly, exposing HOPG to a dilute acid solution, hydrogen-terminating the surface, reduces the subsequent V_{th} by 0.1 V [202]. STM LAO of silicon also has a V_{th} and requires the presence of water [192, 196], which dissociates with OH^- groups displacing protons on the H-passivated silicon surface. This creates a polarised Si–OH bond, making adjacent Si-Si bonds vulnerable to oxidation by further water molecules. As after the acid etch on HOPG, Si–F bonds are even more readily displaced by OH^-. At an existing graphite defect water adsorbs with a barrier of 1.6 eV [248], matching the V_{th} found by McCarley et al. [202]. Both the OH^- and the proton are then chemically bonded at the defect. As with silicon, the C–OH will be polarised, making adjacent C–C bonds vulnerable to oxidation, and the proton is substituted by another OH^-.

Just above V_{th} or in low humidity the reaction rate is sufficiently slow that the reaction only proceeds as far as formation of the oxide intermediate during the tip pass, creating the shallow ridges. These are unstable in air, and spontaneously evolve into gaseous CO over a few minutes, leaving behind a one layer deep trench. At slower tip speeds or higher bias the reaction completes during the tip pass, immediately revealing a trench. Once the lattice is substantially degraded the tip force alone may be sufficient to remove the remaining atoms, accounting for the accrual of dirt on the tip even from a very clean surface, and the build-up of larger amounts of material around cuts. Since the reaction always proceeds via the oxide intermediate the trench edges will be substantially oxygen terminated at the end of the reaction, giving the raised edges. Vacancies and dangling bonds may then rearrange over time, releasing the residual oxygen as CO. Since both the water dissociation energy and oxidation reaction energy are less than ~ 2 eV, the difference is assumed to reflect the higher energy needed initially to dissociate water in the absence of a defect. The linear relationship between V_{th} and activation energy is only valid beyond -2.5 V, suggesting

that the additional barrier height is around 0.9 eV.

The completed picture is consistent with all the experimental observations of this thesis. All that remains is to explain the few conflicts in the literature. The variation in V_{th} with STM can be accounted for entirely by comparing the various tip speeds and pulse lengths with the relationship found by Hiura [200]. The exception is Kim et al. [90], who used AFM, and their data appears to form a similar linear trend in combination with ours, therefore some difference between AFM and STM may be the cause. Three groups reported oxidation at positive tip bias. Venema [116] were working in UHV, so were effecting the SITE mechanism, which is ambipolar. They found $V_{th} = \pm 3.8$ V, which is lower than the ~ 8 V expected, but were cutting nanotubes, and explain the difference through electric field induced strain and localised electronic excitations. Penner et al. [204] was oxidising HOPG submerged in water, and found ambipolar hole generation. Ice has a much higher relative permittivity than water, and the bulk water might not terminate in the ice-like lattice on the graphite surface. The molecules nearest the graphite surface may therefore be less likely to polarise in an electric field, so the field direction is not important. Only the results of Mizutani et al. [203], who make holes with only positive tip bias in a humid atmosphere, are inexplicable with this model. The only apparent difference in their work is an exceptionally small feature size, with holes less than 2 nm wide.

4.6 Subsequent developments in the field

During the latter stages of this work (subsequent to the first lithography on SLG) progress has been made elsewhere etching graphene nanoribbons, and attempts at AFM LAO of graphene have also been reported. The smallest feature size achieved using chemical etching of SLG is 30 nm [33], and oxygen plasma etching can reliably reach 20 nm widths [175, 181, 249] but the error in the process is around 10 nm [249]. Despite the active channel width being narrower than the ribbon by up to 14 nm [249], the energy gap is only 0.03 eV, which is not sufficiently large to give observable gate modulated switching at room temperature [175]; the largest gap measured has been 0.2 eV [249]. Ponomarenko et al. [181] have achieved a point contact 10 nm wide by using repeated short etches to undercut the resist, and measuring the device transport to determine its size. They also measure a $\Delta E \sim 10$ meV for 40 nm wide quantum dot.

4. AFM LITHOGRAPHY OF GRAPHENE

Barthold et al. [250] have used a diamond coated tip to cut a 10 nm thick flake using contact force alone; at 50 μN it is three orders of magnitude greater than that tested here. They observe that the flake resistance rises after a scratch, but then recovers due to bonds reforming and dislocations migrating to the flake edge. The peak resistance rises with successive scratches, and eventually fails to recover. The resulting cut is 100 nm wide. Giesbers et al. [251] report AFM LAO on graphene. However, they are only able to achieve cuts starting from the edge of a flake, and with a bias of -25 V or more, despite 55% humidity and a slow 50 nm/s tip speed. Their minimum cut width on SLG is 30 nm. They assume that water droplets seen after cutting imply total dryness elsewhere on the flake, and hence failure to cut from the flake centre, but this is contrary to the literature evidence and the conclusions of this thesis. Weng et al. [252] find a V_{th} of -15 V in tapping mode, with a minimum cut width of around 70 nm. Their cut widths increase with cantilever amplitude, which may explain the failure of attempts at tapping mode operation here, since the opposite trend was expected. They create a 25 nm wide ribbon between cuts and a sub-micron ring device in FLG and measure its magnetoconductance, and find that a ground connection to the flake itself is not necessary, implying some leakage through their substrate. With STM LAO Tapasztó et al. [253] have achieved ribbons as narrow as 2.5 nm between cuts one layer deep and 2.5 nm wide on HOPG, which is the best resolution reported to date, giving a measured energy gap of 0.5 eV.

A radically different technique was discovered by Ci et al. [254]. At high temperatures and in a hydrogen atmosphere nickel nanoparticles cut FLG graphene via catalytic hydrogenation of the lattice, moving spontaneously along straight lines defined by the lattice angles. The cuts can be less than 10 nm wide, and the nanoparticles switch direction at defects and edges. Alternatively Li et al. [255] have made graphene nanoribbons less than 10 nm wide directly in suspension via polymer functionalisation.

Graphene flakes are now routinely annealed in UHV at up to 500 K to remove adsorbed gasses prior to transport measurements [256], therefore the lattice damage induced by annealing in air is probably due to oxidation. Sabio et al. [257] have calculated that water molecules between graphene and a SiO_2 substrate would give to a separation of around 1 nm, in agreement with the measured height change after annealing above 100℃. During the writing of this thesis, work in the Semiconductor Physics Group has tested current annealing on a graphene flake [258]. The results show a shift in the Dirac point, via an intermediary state with the initial and final curves overlaid in the bulk flake transport. Using SGM the flake is revealed to be divided into patches during this transition, each with either the initial or final Dirac point. Their

shape and dimensions are comparable to the patches observed during thermal annealing when the sub-flake water has been only partially expelled, and current annealing is therefore believed to have the same effect. Elsewhere EFM measurements have shown that the water layer on top of graphene is arranged with the oxygen pointing towards the flake surface even in the absence of an electric field, consistent with the modelling by Zhao et al. [237] and reinforcing the reason for failure of lithography at positive tip bias.

Recent work with LAO on HOPG is complementary to the conclusions above. Park et al. [259] see voltage rather than current dependence, and measure a Fowler-Nordheim type current only in the absence of surface changes. They see holes only above 32% humidity, increasing in width with additional water, V_{th} rises with damaged tips, and holes have slightly raised edges. Jiang and Guo [260] additionally observe mound formation and perform microzone Raman, which shows a broad hump after successful lithography, indicative of organic residuals (functional groups such as C–O and C–OH). They also support the molecular mechanism described above, based on the dissociation of lattice-bound water at carbon vacancies, and further add that the electric field will raise the hole concentration in the graphite beneath the tip, increasing bond polarisation and the likelihood of C–O bond formation.

4.7   Future directions

With the completion of the experiments described above, the reliability of cutting on graphite and graphene was nearly 100%, allowing for the changing V_{th}. That the work did not progress to produce devices for measurement was solely due to the earlier shortage of graphene material and subsequent lack of time, but continuation of the project should aim to demonstrate such devices. In particular, higher frequency current measurement may allow observation of 1D effects in the graphene transport even as its width is reduced during cutting.

Feature sizes of 20 nm were being produced repeatably, but finer features will require practice and reliably sharper tips. More robust tips and cutting in a cleaner, more controlled atmosphere will also help tip longevity and reduce the rate of change of V_{th}. Since 2 nm features have been demonstrated with STM, there is no reason to think that comparable sizes cannot be produced with appropriate AFM equipment, considering that only basic imaging tips and relatively out-dated equipment were employed here. Appropriate alignment patterns might be designed for locating and cutting SLG

by STM, but the flexibility of an AFM makes it a more attractive platform to explore.

In the light of the recent work current-annealing graphene [258] it would be interesting to combine the two experiments, performing SGM on thermally-annealed flakes and topographic AFM on current-annealed flakes, to confirm that the observations are equivalent. Similarly, isolation and comparison of the effects of water on top of and beneath flakes may help elucidate the role of adsorbed water in graphene transport. Exfoliating flakes under UHV and monitoring the return of water to beneath an annealed flake would assist this. It is disappointing not to have the post-annealing Raman data but it is expected to be reliably repeatable.

Graphite is clearly an appropriate test-bed for graphene cutting, and examining the effect of tip velocity (ignored here) and a more rigorous test for a mathematical relationship between humidity and V_{th} would further clarify the cutting mechanism. The unusually wide holes in FLG at high bias and apparent increased stability of the layers nearest the substrate warrant investigation, and the effectiveness of mechanical manipulations of flakes can almost certainly be improved. Interesting subsequent measurements include transport through folds and creases.

4.8 Conclusions

This work has substantially achieved its aims, using an AFM to define nanoscale features in graphene, as small as 5 nm, suitable for low dimensional electronic devices. It is expected that within a very short period of time the project will progress to fabricate and demonstrate such devices. It extends the historic work involving STM LAO of HOPG to describe the first resist-free lithography of SLG. The basic AFM functionality used here produces results comparable with state-of-the-art e-beam and etching facilities, and the resolution limits of SPM LAO are higher still. For one or even several devices the AFM process is quicker than conventional lithography, requiring only a single writing stage of a couple of minutes, in place of resist spinning, baking, e-beam alignment, exposure and etching, making it more useful for research purposes. Lithography by AFM also permits in situ measurement during fabrication and simultaneous modification of the pattern. E-beam lithography has advantages for bulk fabrication, once graphene is available on larger scales, but a multiple parallel scanning probe array such as the 'Millipede' system [261] may become equally effective.

An unexpected conclusion of this work is the crucial role of water in both lithography and electronic transport, and its constant presence on both silicon and graphitic

carbon. As well as confirming many reported results of STM LAO, this experiment has established the relationship between tip bias and hole width, and demonstrated a quantitative dependence of lithography on the ambient humidity. Together with a thorough analysis of the literature, these observations have enabled the proposition of a comprehensive and consistent mechanism explaining both hole and mound generation. The electric field due to the tip bias lowers the activation energy for chemical oxidation of the carbon by water adsorbed at the surface. The crystalline arrangement of the water molecules favours oxidation under negative bias, and the reaction proceeds via an unstable solid state oxide intermediate.

A parallel experiment has shown that the additional measured height of thin graphene flakes is due to a layer of water trapped between the flake and substrate. This water is expelled by heating to above 120℃, and may be responsible for doping of the graphene. Together these experiments advance the characterisation of graphene and its substrate interactions, and add to the available tools for controlling its structure and properties.

Chapter 5

Electrostatic Force Microscopy of Graphene

5.1 Introduction

Transport measurements of graphene using metal contact electrodes measure the bulk behaviour of the whole flake or, at the minimum, a region comparable in size to the electrode geometry. Although graphene is flat and has few defects over micron distances [208] some local variation of transport arising from defects or distortion is predicted at smaller scales. The conductance fluctuations described in Section 4.4 and the irregular conductance across the width of a cut flake found in Section 4.3.4 are both indications of inhomogeneous conductance, albeit in annealed flakes. The dispersion relation of graphene goes to zero at the Dirac point, therefore the density of states is expected to vanish. However, a non-zero minimum conductivity is always observed experimentally despite the zero average carrier density [212]. It is thought that disorder may cause local shifts in the Dirac point, creating a conducting random network of electron and hole puddles near the bulk conductance minimum [28]. The disorder itself may arise from local warping or rippling of the graphene, or from nearby charge doping.

Another expected local effect is an increase in the local density of states near the Fermi energy for certain graphene edge orientations, predicted [262] and observed experimentally [263, 264] on graphite. Edge states may also be apparent between bulk Landau levels, when the density of states in the centre of the flake is lower (see Section 2.2.1) [6].

A measure of the density of states of a system is its electronic compressibility κ,

defined by:

$$\kappa^{-1} = n_e^2 \frac{\partial^2 E_{tot}}{\partial n_e^2} = n_e^2 \frac{\partial \mu}{\partial n_e}$$

for electron density n_e, chemical potential μ and total energy per unit area E_{tot} [265]. In a 2D electron system κ can found by measuring the capacitance between the 2DES and a gate [266]. The inverse capacitance C^{-1} is the sum of two terms, a large geometric contribution due to the spatial separation between the gate and the 2DES, and a much smaller quantum capacitance term. The latter is proportional to $\partial \mu / \partial n_e$ but is usually overwhelmed by the geometric term [267, 268]. This can be overcome by measuring the perpendicular electric field penetrating through the 2DES with a second 2DES, whose capacitance with the gate is therefore sensitive to the quantum capacitance fluctuations in the first, and hence the density of states [17]. However this is not a local measurement.

In order to detect and characterise electron/hole puddles and localised edge states, this experiment aimed to make a local measurement of the screening of a graphene flake using an AFM tip in place of the second 2DES, and hence infer the local density of states. This chapter describes the electrostatic force microscopy technique used for the measurements and presents graphene conductance data and preliminary scanning probe results, with some evidence of the effects described above.

5.2 Electrostatic force microscopy

A conducting AFM tip oscillating at close range interacts with the sample through long-range Coulomb forces, and can be used to make local measurements of electrostatic force (electrostatic force microscopy, EFM) or the surface work function (Kelvin probe microscopy, KPM). These techniques use 'lift mode'; the AFM makes two passes per scan line, one to record the topography then one at constant separation for the electrical measurement.

EFM [269] uses the AFM tip as a local voltmeter. In dc-EFM a DC bias ΔV is applied between the tip and the sample, forming a capacitor between them, and the cantilever is oscillated near resonance. The electrostatic force is given by

$$F = \frac{1}{2}(\Delta V + \phi)^2 \frac{\partial C}{\partial z}$$

where C is the tip to surface capacitance and ϕ is the work function difference between the tip and sample [171]. A change in the local force acting on the cantilever

shifts the resonant frequency, with a corresponding change in oscillation amplitude [270]. Assuming the force gradient is negligible over the oscillation amplitude such that $\partial F/\partial z \ll k$ (k is the cantilever spring constant), it is proportional to a shift in the resonant frequency $\Delta\omega = \omega_0 - \omega_0'$ or phase $\Delta\varphi$ of the cantilever oscillation [94, 184]:

$$\Delta\omega = \frac{\omega_0}{2k}\frac{\partial F}{\partial z}$$

Therefore $\Delta\omega$ is related to the sample capacitance by [271]

$$\Delta\omega = -\frac{\omega_0}{4k}\frac{\partial^2 C}{\partial z^2}\Delta V^2$$

In ac-EFM an AC potential $V(\omega)$ is applied to the sample at the cantilever resonant frequency, in place of a mechanical drive oscillation. This creates an oscillating force on the cantilever

$$F(\omega) = \frac{1}{2}(\Delta V + \phi)V(\omega)\frac{\partial C}{\partial z}$$

and can be performed at lower voltages than dc-EFM, allowing measurements in linear response [171].

The total force experienced by the cantilever is dominated by contributions due to the large tip cone and the cantilever itself, and stray capacitative coupling to large features such as electrodes; only for very small tip-sample distances (below $\sim$5 nm) is it dominated by contributions from the sharp tip apex. However the force gradient at the tip apex is far larger than at the cone or cantilever to greater distances. Since the scanning resolution is determined by the size of the detector, the best resolution is therefore achieved when the signal source is the force gradient rather than the force itself [272]. Though the frequency shift is proportional to the force gradient, the oscillation amplitude is proportional to the force [269], so while both provide information about the local sample capacitance, frequency shift data has the higher spatial resolution.

For measuring, either the AC bias frequency is modulated to keep the tip at resonance and the resulting $\Delta\omega$ is recorded, or $\Delta\omega$ is kept constant by modulating the scan height, so that the tip position plots contours of constant force gradient. The frequency shift also has a faster response for feedback than the amplitude.

KPM [6, 172, 273] is similar to ac-EFM but the AC bias voltage is applied to the cantilever instead of the sample. The DC bias is adjusted by feedback to keep the tip and sample at the same potential. This method maps the electrostatic potential and hence local work function of the sample surface.

A different measurement technique is scanning gate microscopy (SGM) [86, 171, 274], which uses the AFM tip as a local gate to directly image scattering sites. The conductance of the sample is measured as a function of tip position and voltage, and changes when the tip locally depletes or gates the underlying electron system.

5.3   Apparatus

5.3.1   Dilution refrigerator

Measurement was done with an AFM mounted inside an Oxford Instruments Kelvinox dilution refrigerator, capable of temperatures as low as 17 mK (see Figure 5.1). This employs the spontaneous phase separation of a mixture of ^{3}He and ^{4}He below its tricritical point (860 mK) into ^{3}He-rich and ^{3}He-dilute phases. The enthalpy of ^{3}He in the two phases is different, so the cooling power derives from the endothermic transport of ^{3}He atoms from the rich phase into the dilute phase, analogous to evaporative cooling. Below 2.18 K ^{4}He is a superfluid and ^{3}He can flow through the dilute phase with little impedance. The cooling process is maintained by continuously evaporating ^{3}He from the dilute phase in a pumped still (it does so preferentially to the ^{4}He), recondensing it in a heat exchanger and returning it to the rich phase.

This cooling circuit is surrounded by a vacuum space ('sample space') also containing the device, which is in thermal contact with the mixing chamber. These elements together comprise the 'fridge', which is lowered into a ^{4}He bath containing a superconducting magnet capable of 10 T fields, all surrounded by a further vacuum space and hanging from ropes to insulate it from mechanical vibrations. Prior to cooling a small amount of helium gas is released into the sample space, which transfers cooling power from the helium jacket until it condenses at 4.2 K.

5.3.2   Low temperature AFM

The device was mounted on top of the coarse positioner, used for approximate alignment of the device with the tip, which was a stack of 3 (X, Y and Z) pairs of slip-stick blocks each containing a piezoelectric crystal (Figure 5.1c). In piezoelectric materials the unit cell has no centre of symmetry, so a physical displacement of the ions in the lattice polarises the crystal. Conversely, polarising the crystal causes it to expand or contract, which is used here to move a pair of blocks apart or together. The slip-stick mechanism exploits a difference between static and kinetic friction. For a slow expan-
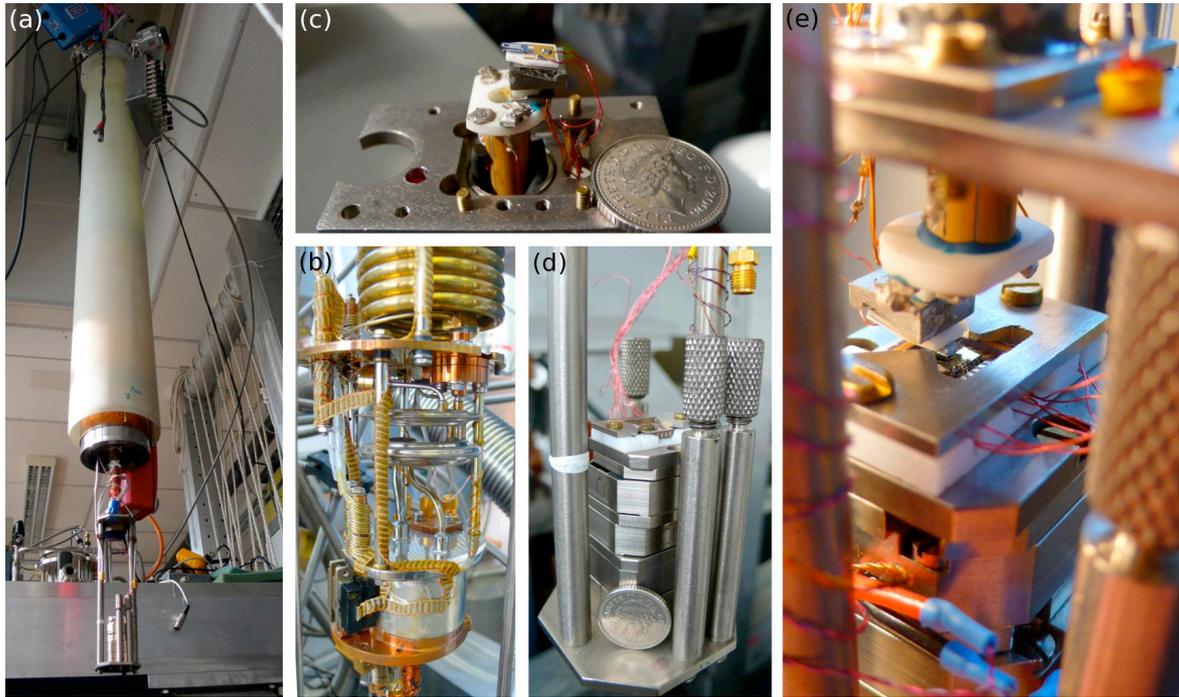
Figure 5.1: Low temperature scanning probe. (a) The whole length of the fridge, raised out of the helium bath. (b) The cooling circuit, located within the bottom of the plastic sheath in (a). (c) The AFM, showing the cantilever mount (the cantilever and tip are too small to be visible) and wires for resistance measurement, orange scan tube, and a 5 pence coin for comparison. (d) The coarse positioners, with a chip package mounted on top and soldered wires. (e) Field of view for manual alignment of the tip to the device beneath it, before sealing the sample space.

sion (or contraction) of the piezo the static friction is sufficient to move the adjacent block (the 'stick' phase). This is followed by sudden movement in the opposite direction, overcoming the friction such that the block remains static (the 'slip' phase). In this way the blocks are moved stepwise.

The tip hung above the sample from the scan tube, a piezoelectric cylinder with four quadrant electrodes on the outside for X–Y movement and a single electrode inside for Z movement (Figure 5.1c & e), which drove the scanning motion. Piezoelectric materials operate reliably from room temperature to cryogenic temperatures, but their range is substantially reduced at 4 K and below. The coarse positioners could be moved several microns per second at room temperature, but only a few nanometres at 20 mK, and they were becoming worn out and often froze during cooling, therefore most alignment

had to be done beforehand. The lateral range of the scan tube was about 100 μm at room temperature and 10 μm at 20 mK. The scan tube was also deteriorating, and had up to a 100% difference between the X and Y scan ranges, an irregular shear distortion of about 20% of the scan width and uneven stepping between scan lines. The distortion and line separation were hysteretic between upward and downward scanning. There was an unpredictable positional drift during cooling of up to 150 μm.

In place of a conventional laser detector, deflection of the cantilever was measured with a piezoresistive material coating its surface [186], and the cantilever was U-shaped to define a current path. When the cantilever bent, the resistance of the coating changed by a small amount, measured by a resistance bridge across the two arms of the cantilever and a strain gauge amplifier. This gave a power dissipation within the fridge of only 30 μW. The scan tube and cantilever mount were built by the Engineering Department and East Coast Scientific in Cambridge, and the cantilevers purchased from Park Scientific Instruments of California (now part of Veeco).

5.3.3 Measurement techniques

For topographic scanning the AFM was used in contact mode with no feedback. A single step of the coarse Z positioner was less than the cantilever deflection range, therefore the sample was raised step-wise until a deflection was detected. The target deflection for imaging was then achieved by vertical motion with the scan tube. The surface profile was recorded as the output of the strain gauge. Contact mode scanning over a graphene edge is liable to damage it, so was only used for approximate location of the flake.

Compressibility measurements and any scanning over the flake itself were done in ac-EFM mode. The flake was grounded and a DC bias applied to the back gate, with an added AC bias initially at the cantilever resonance ω_0 (Figure 5.2). This was found by manually sweeping the bias frequency and monitoring the cantilever response, and drove oscillation of the tip. Shifts in the resonance due to changing forces on the cantilever were measured using a phase-locked loop which modulated the drive frequency. The oscillation amplitude was measured by a lock-in amplifier and together with $\Delta\omega$ was recorded by a PC.

The tip measures the force curve normal to the sample surface, which on raised features is lifted correspondingly, therefore the EFM signal is a combination of topography and local force variation due to changes in the sample capacitance. The total force is a combination of long range Coulomb forces and short range van der Waals
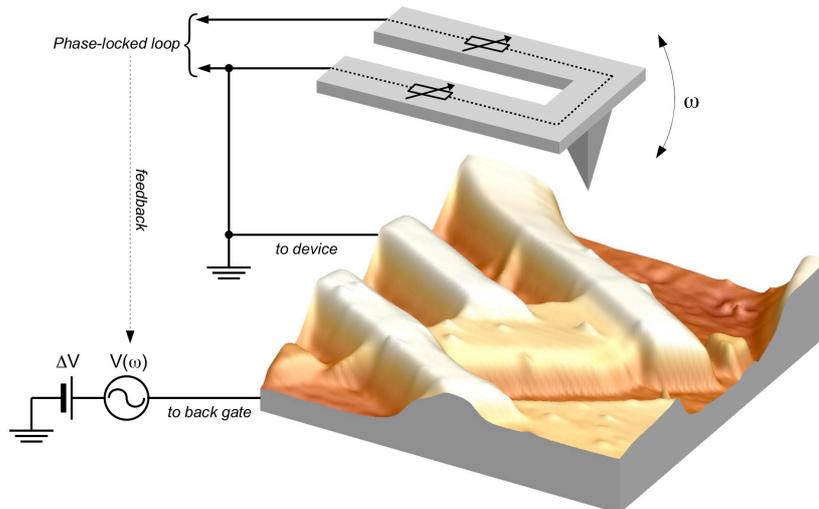
Figure 5.2: Schematic of EFM measurement. The device is grounded and a DC bias ΔV is applied to the back gate, plus an AC bias $V(\omega)$ to oscillate the cantilever. Deflection of the cantilever changes the resistance of its piezoresistive coating, enabling measurement of its oscillation. A phase-locked loop detects changes in the resonant frequency of the cantilever, and modulates the frequency of the driving oscillation to keep the cantilever on resonance.

forces. At more than a few nanometres from the surface the latter are comparable for most materials, so eliminating the effect of height variation with lift mode scanning also largely eliminates variations in the van der Waals contribution. The fridge and AFM had been designed and built for previous SGM work on GaAs devices [6, 274]. Since the samples involved were all flat epitaxial wafers there was no need for a topographic correction during the second, electrical measurement scan, and the controlling hardware did not have the facility for lift mode. Since contact mode measurement is not feasible on the device, topography cannot be eliminated from the data in these experiments, therefore only qualitative analysis is possible.

Figure 5.3 shows the frequency shift and amplitude response as the tip approaches the surface in EFM mode, together with the deflection during a contact mode approach and a schematic depiction of the force on the cantilever as a function of distance from the surface. There is a small jump to contact during the contact mode approach, and a large hysteresis on retraction, so adsorbed surface species such as water are not removed completely in UHV. The amplitude curve has a maximum about 100 nm from the surface when the force gradient switches sign. From this point the frequency
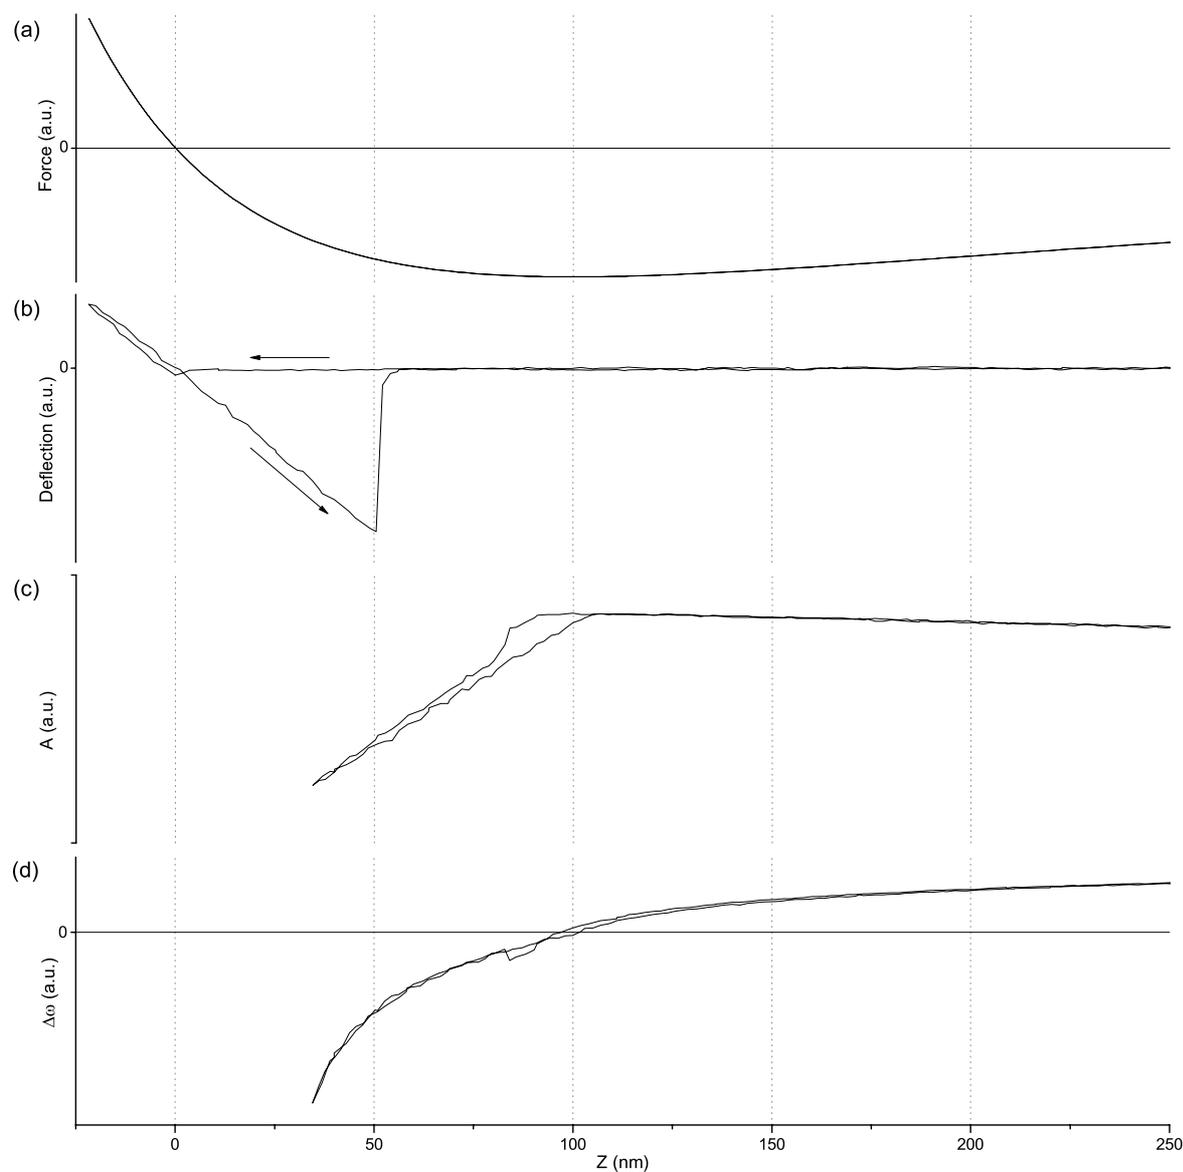
Figure 5.3: EFM tip approach curves. (a) Schematic force curve (negative is attractive) for the tip at height Z above the surface. (b) Deflection data from a contact mode approach/retraction - the tip sticks to the surface during retraction. (c) Corresponding amplitude and (d) frequency shift data. In these experiments the frequency shift is recorded as a voltage output, with a large offset such that it is always reported as positive (the curve above has been shifted). Scans were performed on both sides of the force minimum, and therefore at both positive and negative actual $\Delta\omega$.

shift also becomes more noisy and scanning is no longer stable. The tip height at the amplitude maximum is the combined range of repulsive surface forces and the oscillation amplitude (neglected in matching heights relative to the force curve), and is approaching the minimum tip-sample separation for effective EFM scanning - generally around 100 nm, the smallest value used was 75 nm.

When imaging contacts and graphene together with constant tip height, the additional separation needed to pass over the contacts gave a large separation from the graphene, and a correspondingly low image resolution. Scanning with feedback to keep $\Delta\omega$ constant by adjusting the tip height allowed a closer approach to the graphene and better resolution.

Shielded wires from the chip package travel the length of the fridge and are connected to measurement electronics via a break-out box. Keithley 2400 SMUs were used to apply both back gate and source-drain biases for conductance measurements.

5.3.4   Device design

Graphene devices were fabricated on silicon and packaged in a similar manner as for Chapter 4, but with certain modifications to accommodate the constraints of the fridge. The confined space and low temperature requirements of the fridge precluded the use of an optical microscope for AFM alignment; the tip could only be aligned by eye before the fridge was sealed, and the device then found by AFM scanning. To aid this a complex pattern of metal alignment marks was designed for the chips (see Figure A.8). The chip was divided into 4 square quadrants each 1.5 mm wide. The tip was aligned by eye to the quadrant that contained the device, but not over the device itself, then lowered for contact mode scanning. Each quadrant was divided into 25 sectors, each of which contained a grid of uniquely shaped alignment marks spaced every 30 μm. The scanning range at room temperature was sufficient to image several of these marks, identifying the current sector, before moving step-wise to the edge of the sector containing the device. A second alignment grid, of various arrow shapes pointing to the device, was laid together with the device contacts, spanning a 150 μm wide square centred on the flake (Figure A.8b). Once the edge of this was found the tip was withdrawn and the fridge cooled to base temperature.

After cooling, the AFM was again engaged in contact mode. The arrows in the secondary alignment pattern were spaced 7.5 μm apart, so that at least one was visible in every low temperature scan, and in concentric rings of changing arrow shapes, so that the distance from the device was apparent. The tip was moved stepwise towards

the device until imaging within a few microns of the flake. Scanning the flake itself in contact mode risks damaging it, therefore the tip was then disengaged, re-engaged in EFM mode, and moved over the flake. Thus, with calibration of the coarse positioner movement, the device could be reached in not many more than 3 imaging steps.

Contact pads were on only two sides of the chip, to allow tip access unobstructed by bond wires. Devices were patterned with as many contacts as the flake size and the number of bond pads permitted, laid in Ti/Au to ensure good lift-off reliability (see Figure A.7 for an example). They were arranged to give four terminal measurements both along straight edges and in Hall bar configurations, with additional redundancies where possible. Overlap with the flake was minimised to leave a large area clear for EFM scanning. Redundant contacts allowed for a couple of inevitable failures during lift-off and wire bonding. However, the method for mounting the devices in the fridge was highly compromised, and many more contacts were unusable as a result. Less than one in three were ultimately available for measurement.

5.3.5   Device mounting

The design of the sample space left insufficient room for a chip package socket beneath the scanning tip, therefore trailing wires were soldered directly onto the package (Figure 5.1e). This was suitable for plain GaAs and Si samples but the graphene was too fragile, and the first device to be mounted (G10, previously thermally annealed and measured in Section 4.4) blew up during soldering. As an alternative, short lengths of wire were soldered onto a package before the chip was mounted and wire bonded. These wires were then soldered to the existing wires in the fridge. While this protected the device, it meant that the package could not be mounted in any other dipping probe or cryostat for testing. Replacement of the positioner piezo stack after it fractured (see below) later permitted some rearrangement of the sample space and room was made for an additional chip package. Therefore a socket was improvised where one package was soldered in place, and the package containing the device was sat on top and joined to the lower with silver glue connections. While this was effective, the solder and glue connections were fragile and delicate to make, and several failed during mounting and cooling.

At this time CAPE were unable to produce new graphene, therefore a wafer with 7 wholly or part single layer flakes was purchased from Graphene Industries at Manchester University [213]. Also, the lithography e-beam system was about to begin several months of repair, so only a single writing slot was available. Two finished devices

were unusable because of current leakage through the oxide to the back gate. This was observed intermittently with the other devices and appeared to be a recurring problem on this wafer sample. Two were lost because of e-beam writing misalignment and damage while removing residual PMMA (see below) and a completed device was destroyed when the tape holding the cantilever mount in place came unfastened during cooling and the tip dropped onto the device.

The two devices that were measured, S2211 (Figure A.7) and G24 (Figure 5.9), had 10 and 11 contacts respectively. Failures during lift-off and wire bonding were not exceptional (5 contacts lost in total) but the additional solder and silver interconnects added great fragility to the circuit; 3 solder joints broke outright once the package was fastened, and a further four were not conducting. As a result neither device was capable of a full 4 contact measurement once in the fridge. The largest flake was lost entirely after preliminary measurement, as a result of an erroneous helium fill prior to this project which had caused a build-up of nitrogen ice around the magnet. This resulted in an impact jolt when lowering the fridge, and subsequently two magnet quenches during measurement. The shock from these destabilised and later cracked the piezo stack, causing the probe tip to swing across the chip surface, slicing through the remaining contacts and the graphene. Manual repairs using silver glue functioned at room temperature but not after cooling the device.

5.4 Bulk transport measurements

The intention had been to measure the magnetoconductance of each flake for comparison with observations from EFM measurement, but the loss of contacts hampered this and G24 could not be measured electrically. Room temperature two-terminal measurement of the gated conductance of S2211 gives more typical profiles for graphene than observed after thermal annealing in air (compare with Figure 2.6). The conductance minima are sharper and deeper, with the bulk carrier switching from electrons to holes over a range of about 3 V, compared with $\sim$ 30 V in Figure 4.25. Figure 5.4 shows curves for different current paths around the device. A path along the edge of Flake A (11–13) has a lower resistance and a broader Dirac point at a fractionally higher back gate value than a path across the flake width (11–3), even though two pairs of contacts are equidistant. The Dirac point of Flake B, subject to the same environment and processing, is about 1.5 V lower, so flake geometry or local effects do play some role in determining carrier concentrations. A current path through Flake C, lying across
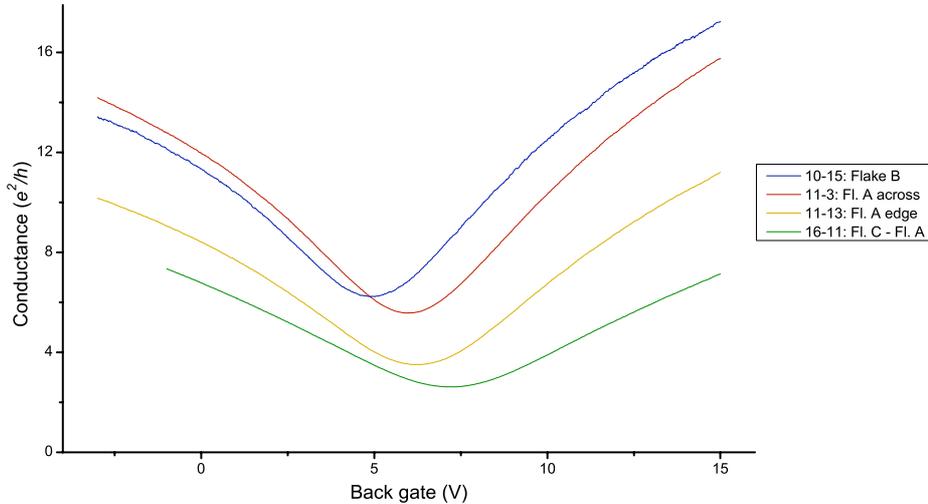
Figure 5.4: Room temperature graphene conductance. Two-terminal conductance measurements of various current paths through device S2211, labelled according to the scheme in Figure A.7. Different flakes have conductance minima at different back gate values.

Flake A, shows a broadened minima where the curves for both flakes are overlaid.

The tip crash described above subsequently damaged many of the contacts, and after cooling to 20 mK the only circuit possible was a two-terminal measurement of two flakes in series (contacts 13–15 through Flakes A and B). The resulting conductance and magnetoresistance data in Figure 5.5 show quantum Hall plateaux, but not at the normal intervals for SLG ($R_{xy}^{-1} = \pm 4(n + 1/2)e^2/h$ for non-negative integer n [37] - see Section 2.4.2) because of the addition of series conductances. The data can be quantitatively explained if the Dirac point of one of the flakes has shifted negative by several volts. The plateaux at $V_g = 8$ V (Figure 5.5c) are approximately at 1.8, 4.5, 6.4 and 7.9 e^2/h, equal to the values for filling factors $\nu = 2, 6, 10$ and 14 in one flake intersecting with a plateau for $\nu = 18$ in the other. At negative V_g above 5 T there is interference between the $\nu = 2$ plateau of the first flake and the lower filling factors (2, 6, 10, 14) of the second, with combined conductance values of 1, 1.5, 1.67 and 1.75 e^2/h. The zero field conductance is at $4e^2/h$ as expected [212].

The quantum Hall plateaux are further distorted because a two-terminal measurement is affected by its geometric characteristics, with contributions from both the longitudinal and Hall conductivities depending on the aspect ratio (Figure 5.5d). A sample of arbitrary shape can always be described with a rectangle of width W and
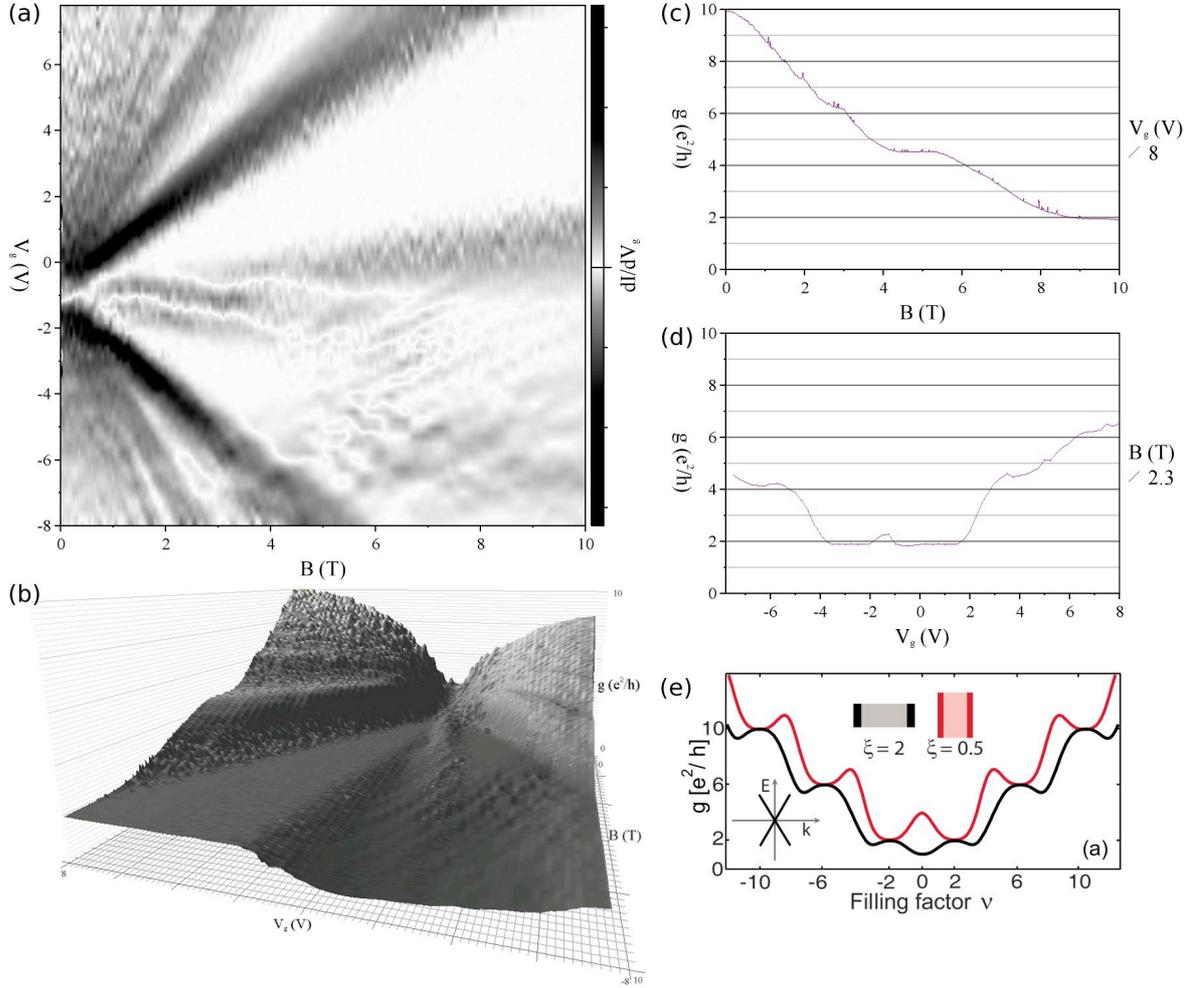
Figure 5.5: Low temperature graphene conductance. (a) Differential plot of two-terminal gated magnetoconductance measurement of two flakes in series at 20 mK. The current path is from contacts 13 to 15, via Flakes A and B and the floating contact 14, shown in Figure A.7a. (b) Conductance surface plot. (c) Magnetoconductance at $V_g = 8$ V showing quantum Hall plateaux. (d) Gated conductance at $B = 2.3$ T with a peak within the conductance minimum and at the plateaux edges, resulting from the measurement geometry. (e) Modelling of two-terminal conductance for two different contact geometries by Williams et al. [275].

length L by conformal mapping, without knowing the actual current path [276]. The finite longitudinal conductivity leads to a modification at the charge neutrality point and N-shaped distortions of the plateaux, which are of opposite sign for $\xi = L/W > 1$ and $\xi < 1$ [275] (see Figure 5.5e). The peak at the centre of the charge neutrality point and small maxima at the leading edges of the plateaux (and hence minima at the densities of compressible quantum Hall states) indicate for one of the flakes a nearly square geometry with W slightly greater than L. This corresponds more obviously with Flake B, so Flake A now has the more negative Dirac point.

5.5 EFM measurement

5.5.1 Bulk flake behaviour

Figure 5.6 is an EFM scan (positive $\Delta\omega$) showing regions of graphene (Figure A.7: Flake B), contacts and substrate. The grounded metal and graphene screen the tip from the electric field more than the oxide, so the tip experiences a smaller force. This reduces $\Delta\omega$ and the feedback loop lifts the tip to keep it constant. This AFM was not calibrated so tip heights Z are estimated to within 5 nm by comparison with features of known height in contact mode scans. To within this accuracy the EFM lift height moving between the contacts and graphene is the same as their real height difference, so they are screening by approximately equal amounts, but the tip lifts an additional 105 nm compared with over the substrate. An alignment mark appears only a third the height of the contacts (Figure 5.6a top-right) even though it is the same thickness, because it is not grounded but does still add a small moderating dipole field.

The measured $\Delta\omega$ (Figure 5.6c) is an error signal resulting from a slow feedback response, and it gives an approximate differential image. The width of edges is a combination of the feedback response time and the EFM resolution at this separation. The amplitude signal (Figure 5.6b) has much lower resolution because sensing is over the whole cantilever, and the alignment mark is too small to be detected. This signal is also a residual, since the amplitude shift is reduced by the feedback on $\Delta\omega$, but the amplitude response has a longer range so it is not fully nulled. It still therefore gives a qualitatively correct result, with a lower amplitude over the metal and graphene. Figure (Figure 5.6d) is scanned at a lower $\Delta\omega$ set point, giving a smaller separation where the resolution is sufficient to detect lumps of dirt 200 nm wide. These appear at their true height (10–30 nm) even on the graphene, so the EFM signal is still dominated by Coulomb forces.
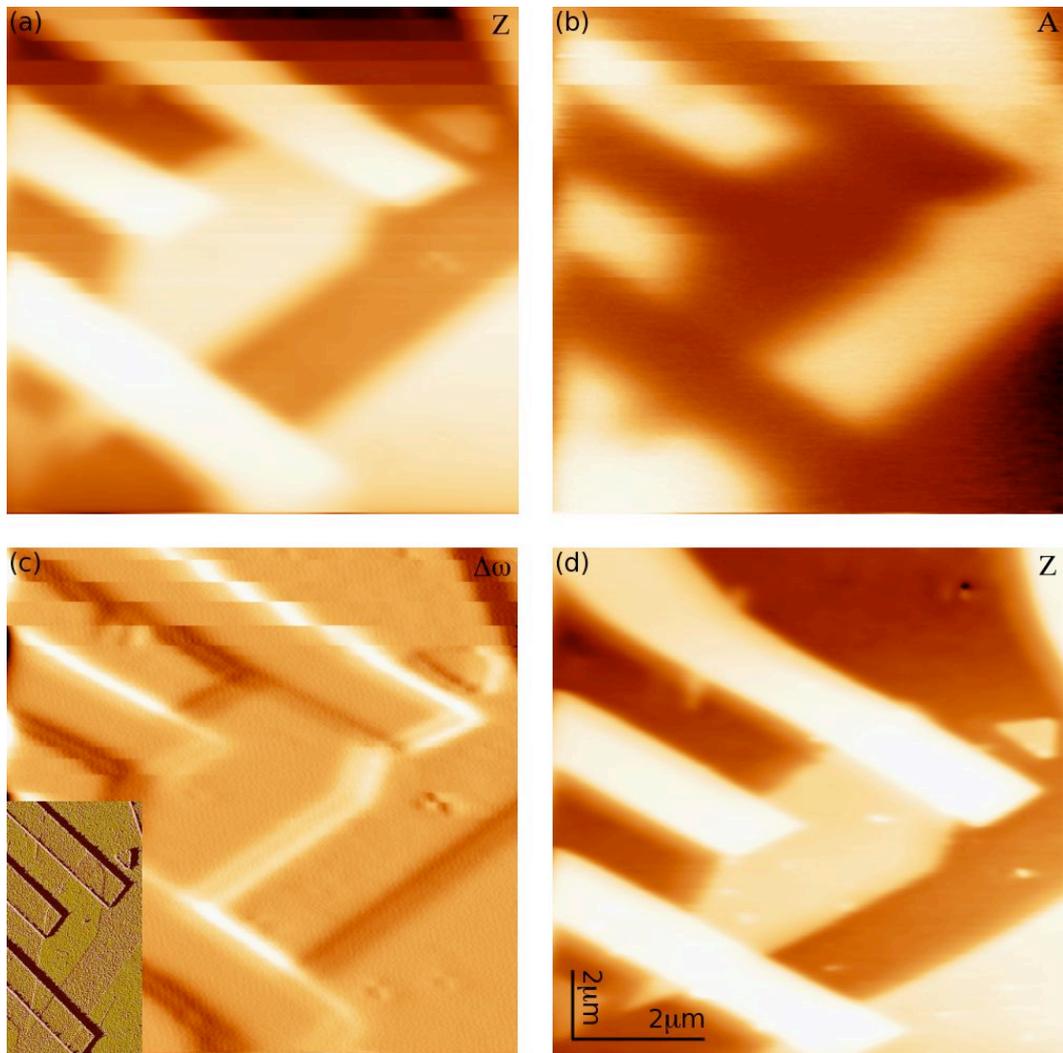
Figure 5.6: EFM scans of graphene. (a) Tip height as a response to feedback on $\Delta\omega$, showing S2211 Flake B and three contacts; (b) amplitude signal and (c) residual $\Delta\omega$. (d) A scan of the same region at lower $\Delta\omega$ and hence smaller tip-sample separation. Colour scales are arbitrary, and brighter equals higher tip height/greater frequency shift/larger amplitude.
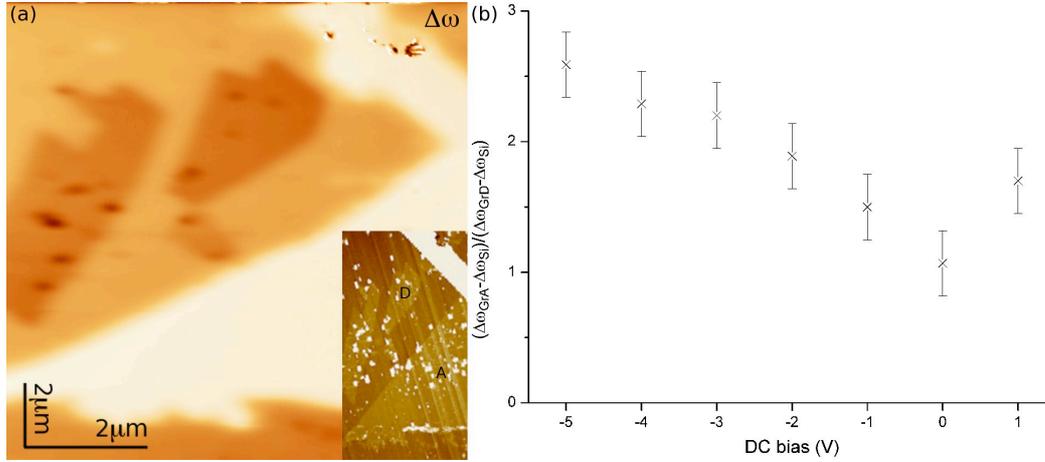
Figure 5.7: EFM with changing DC offset (i). (a) $\Delta\omega$ (without feedback) at $\Delta V = -3$ V of three flakes, shown in the inset topographic AFM scan. The dirt and jagged edge at the bottom are the result of a tip crash tearing the flake. (b) The ratio of the difference in frequency shift of the grounded and ungrounded flakes relative to the substrate at various values of DC bias, showing a minimum near 0 V.

Figure 5.7a (without height feedback, negative $\Delta\omega$) is one of a series taken with different DC bias, which alters the signal over graphene by changing both the EFM DC offset and the gating electric field. The DC offset has a strong influence on the response curve of $\Delta\omega$ and to maintain consistent contrast and resolution the scans were taken at different heights, therefore absolute values of $\Delta\omega$ cannot be compared.

Also in the image is an ungrounded flake (D), which gives a lower $\Delta\omega$ than the substrate. The signal at any point on the surface is the sum of influences between the back gate and the tip, so is typically a measure of the cumulative screening effect of silicon plus graphene, or silicon plus contact metal. However, above this flake the screening is less than over the silicon alone. This means that the ungrounded flake is adding to the local electric field and is therefore charged. This may be due to doping by, for example, water trapped beneath it (see Section 4.4). Figure 5.7b shows the ratio of the difference in frequency shift of the grounded and ungrounded flakes relative to the substrate at various values of DC bias. This eliminates the effect of changing DC bias on the $\Delta\omega$ response curve and, assuming that the screening of the ungrounded flake is constant, gives a measure of the changing screening of the grounded flake. This appears to have a minimum just below $V_g = 0$ V, reasonably consistent with the transport measurements for this device.
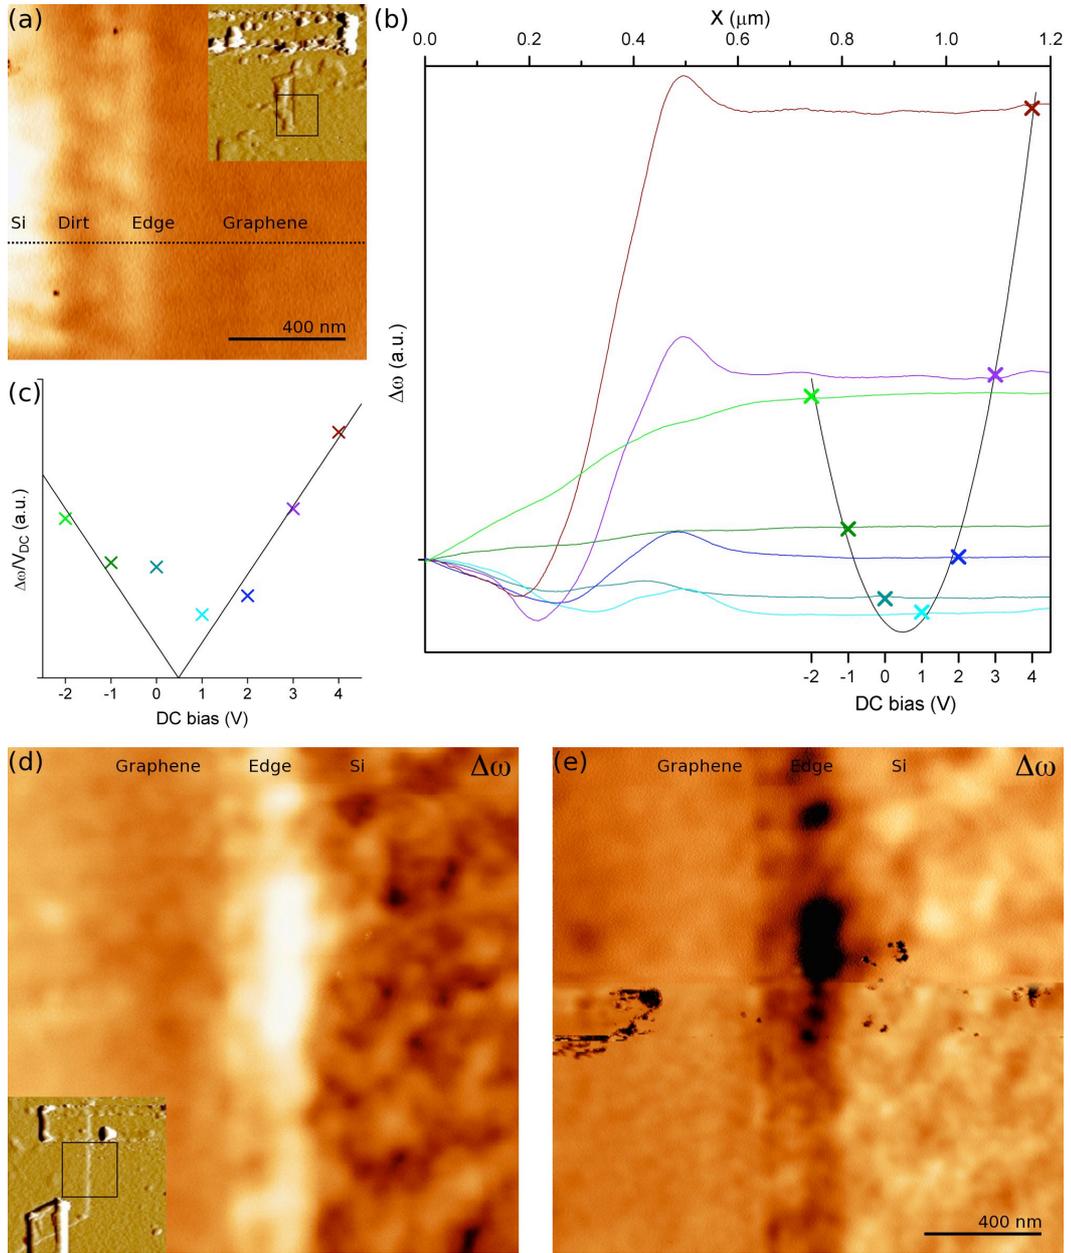
Figure 5.8: EFM with changing DC offset (ii). (a) $\Delta\omega$ of a graphene edge, shown in the inset topographic scan. (b) Horizontal sections of $\Delta\omega$ at different DC bias, normalised to the the substrate, with a parabolic fit to the mean values over the graphene against changing bias. (c) The mean $\Delta\omega$ values divided by and plotted against the square root of the parabolic fit. (d) & (e) $\Delta\omega$ scans over a different edge showing an inversion of the graphene and substrate and a significant edge feature not visible in the topography. The surface texture is due to residual PMMA.

The experiment was repeated on device G24 scanning at nearly constant height. Figure 5.8b shows sections through the $\Delta\omega$ scans, normalised to the frequency shift over the substrate on the far left. The relative shift over the graphene (plateaux on the right) closely follows a parabolic curve with changing bias. In ac-EFM the frequency shift is proportional to the DC bias ΔV. The graphene carrier concentration is proportional to V_g [27], and if the quantum capacitance of graphene is proportional to the carrier concentration [277], so also are its derivatives. Therefore, since in this experiment ΔV and V_g are equivalent at an offset, the minimum of the parabola is an estimate of the Dirac point position of the flake, at $V_g = 0.48$ V. Dividing the data by the square root of the fitted curve (Figure 5.8c) shows that the fit is weakest for the points nearest the Dirac point. This is expected because the actual flake conductance has a finite minimum and so a parabolic fit is not valid in this region.

Two additional features are visible at positive bias: there is a peak at the very edge of the graphene and a trough a short distance from the edge. The latter coincides with a 10 nm high ridge of dirt pushed off the graphene by contact mode scanning (see Section 5.5.2 below) and, as with the ungrounded flake above, it appears to be adding to the force signal. The peak at the graphene edge implies additional screening and may be some form of edge state, perhaps due to adsorbed species terminating the edge of the graphene lattice.

The presence of the piled dirt prevented closer, higher resolution scanning. A pair of scans over a different edge (Figure 5.8d–e) shows the inversion of the graphene and silicon screening, and a repetition of the edge feature.

5.5.2   Cleaning graphene

After lithography to lay contact metal a layer of residual PMMA about 2 nm thick always remains on the graphene. Its surface is rough, with an RMS height variation of 2 nm across length scales of 5–10 nm, and occasional lumps tens of nanometres high. The latter prevent close EFM scanning of the surface, giving an optimum resolution of a few hundred nanometres, and even where there are no large lumps the PMMA texture influences the EFM signal.

Initially cleaning was attempted in the same manner as for AFM lithography (Section 4.3.2), using the AFM tip to push the dirt out of the scan region. This was not so effective here, possibly because the tips on this AFM are designed for contact mode use so are more flexible and liable to run over the top of the dirt. Subsequent conventional AFM imaging revealed that the dirt was moved, but rarely clear of the flake, being

instead piled into higher lumps that made EFM scanning almost impossible (Figure 5.9a). Only one sufficiently large area was cleared for measurement, which is described in the next section.

After removal from the fridge this same device was cleaned with various resist strippers, usually bathing the chip for 20 minutes at 70℃. Posistrip, used previously for GaAs devices, actually increased the dirt thickness to about 4 nm and if not removed carefully left additional crystalline deposits (Figure 5.9b). Shipley 1165 appeared to give better results and removed the larger dirt particles, but a subsequent tip crash rubbed the flake surface, revealing that the layer of continuous dirt was still 3 nm deep (Figure 5.9c–d). Therefore no effective solution was found, although no further devices were measured.

5.5.3   High resolution EFM and graphene texture

To observe variation in the EFM signal over length scales of a few nanometres, the minimum separation and a perfectly clean sample were needed. This was only achieved in one instance, and features were only revealed above the signal noise with scan speeds as slow as 150 nm/s. Figure 5.10 shows the averages of three such EFM scans. Scanning is at constant height but with feedback modulating the drive amplitude to keep the cantilever amplitude constant. This gives an amplitude signal that is exceptionally sensitive to the topography, as revealed by comparison with a corresponding contact mode scan. The same features are present in both, though the noise level is much less in the amplitude signal, revealing fluctuations as small as 15 nm wide. The frequency shift is largely unaffected by the additional feedback and shows a completely different texture, with features around 50 nm wide.

Since the topography is still present in the data, it is instructive to convert the frequency shift to a feedback response in tip height, based on the response during a static tip approach. The scanning height here is 110 nm, at which a metal gate gives a frequency shift equivalent to a 15 nm change in tip height relative to the substrate, compared with 6 nm for bulk graphene. Since the gate is in fact 45 nm higher than the graphene this $\Delta\omega$ measurement is almost independent of topography. The variation in EFM height change for the features in Figure 5.10a is over 3 nm, so local variation in the graphene screening is at least 50% of the bulk value. These measurements were done with a positive DC bias of 0.5 V, similar to the estimated position of the charge neutrality point from Figure 5.8. Therefore not only is the observed variation in $\Delta\omega$ reasonable evidence for local variation in the graphene screening on a different length
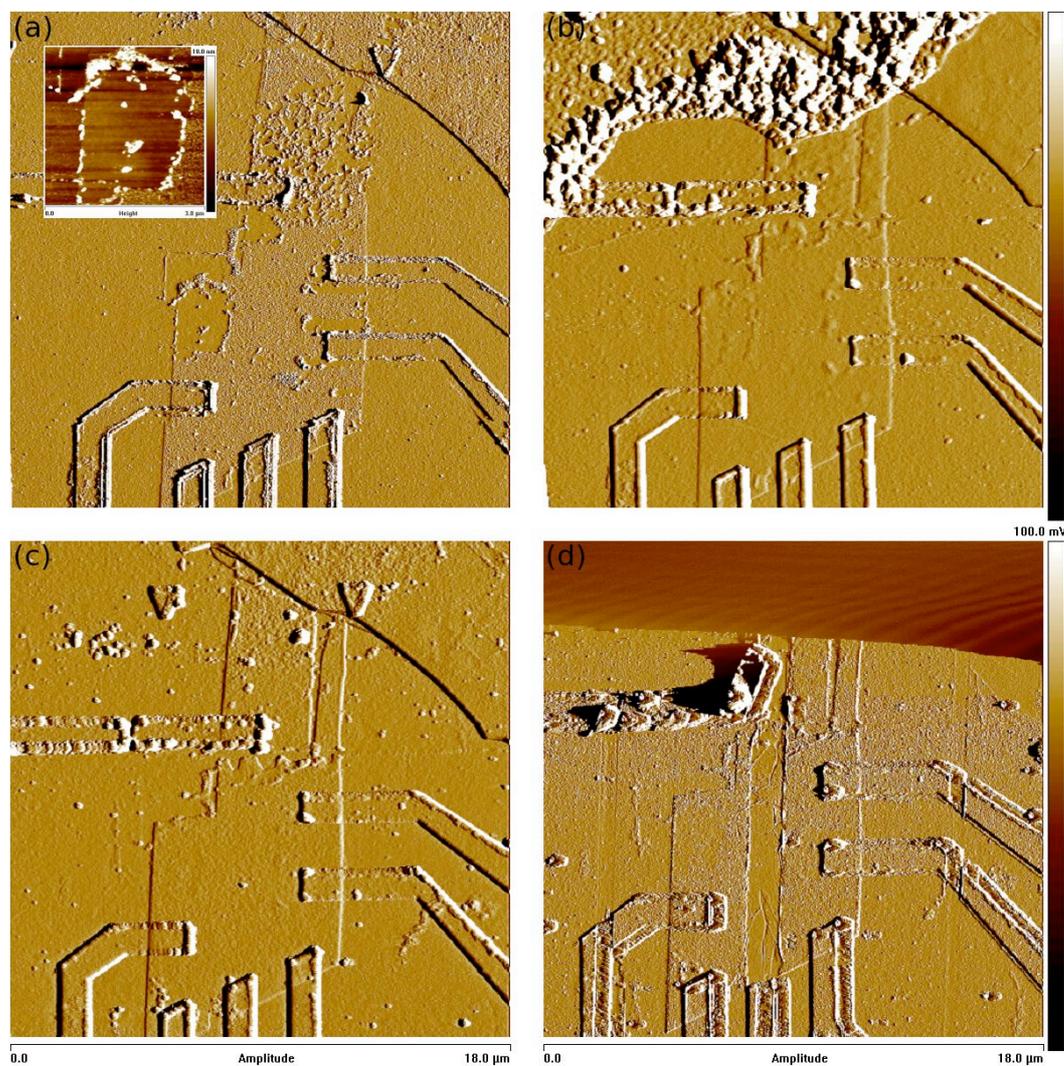
Figure 5.9: Cleaning residual PMMA from graphene. (a) The flake after contact mode sweeping in the low temperature AFM and (Inset) a close-up of cleaned graphene and piled debris. (b) Washing in Posistrip left additional residues. (c) Shipley 1165 appeared to remove much of the dirt. (d) A sharper scanning tip shows that the texture is still rougher than on graphene that has been rubbed clean.
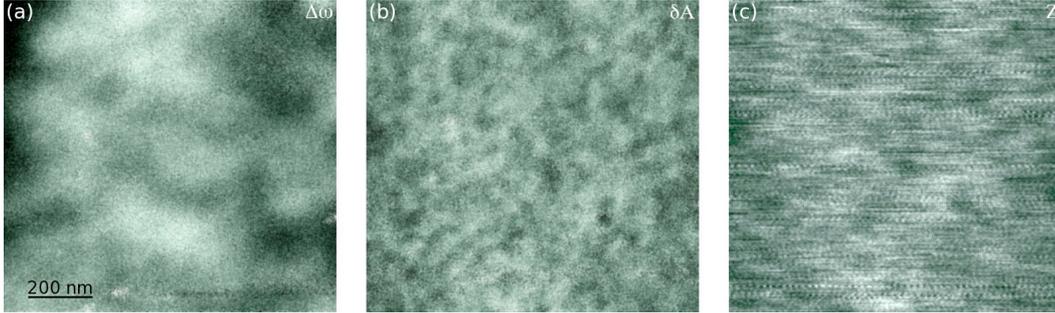
Figure 5.10: High resolution EFM. Averages of three consecutive slow EFM scans to reduce the signal noise, showing (a) features on a lateral scale of ∼50 nm in the $\Delta\omega$ signal and (b) on a lateral scale of ∼15 nm in a nulled amplitude signal. (c) The latter correlate closely with the flake topography, as measured by contact mode AFM. The colour scale is inverted and high features are darker.

scale to its topographic variation, it may in fact be arising here from electron and hole puddles.

Undulations of graphene, a few tens of nanometres wide, are frequently observed by AFM; Figure 4.8c–d shows corrugations of the same lateral size as Figure 5.10b. This is generally attributed to the graphene conforming to the substrate roughness, as demonstrated after annealing a flake in Section 4.4, and Ishigami et al. [278] find that the correlation lengths of the texture of graphene and its SiO_2 substrate are consistent with the graphene morphology being determined by the silicon roughness. However, additional intrinsic ripples are predicted; warping on a lateral scale of ∼ 10 nm would increase the elastic energy but suppress thermal vibrations, minimising the free energy overall [28, 279]. Stolyarova et al. [280] have observed sub-nanometre high fluctuations on such a length scale, which were not consistent with the underlying substrate texture. Meyer et al. [281] have used electron diffraction to detect intrinsic corrugations in suspended graphene, which had a 5° mean variation of geometric surface normal. They were able to estimate a maximum length scale of 25 nm.

The features in Figure 5.10b cannot be attributed to either substrate conformation or intrinsic ripples, and no predictions have been made for length scale of electronic fluctuations within graphene, independent of topographical variation, so no correlation of the $\Delta\omega$ data can be made here.

5.5.4 Static EFM spot measurements

In assessing local structure within the graphene an applied magnetic field is of more interest than the gating electric field, since features such as edge states and current loops are expected while passing through Landau levels. After the scans in Figure 5.10 the magnetic field was ramped to 1 T. While the amplitude data was unchanged, the frequency shift signal was swamped by noise and scanning was not sufficiently stable for imaging over even longer time periods. In other scans at larger separations, changing the magnetic field appeared to have a more substantial effect on the cantilever response than the DC bias. To mitigate this, EFM measurements were made with the tip static, at a fixed height over a single spot on the sample.

Figure 5.11 shows measurements taken with the tip over both graphene and the substrate, while ramping the magnetic field at intervals in the DC bias. The data over silicon is considered to be the background to the graphene data, combining the response of the substrate and of the cantilever. The $\Delta\omega$ are qualitatively similar, with the graphene showing only a steeper gradient. There are some low field features, but these did not vary with changing bias.

The amplitude curves are markedly different, and over the graphene exhibit periodic oscillations which are absent over the substrate. These move to higher field with increasing back gate, in the same manner as Shubnikov-de Haas oscillations arising from quantum Hall behaviour in magnetoresistance data. In the latter case maxima occur in the resistance at fields B_N when the Fermi energy E_f crosses a quantised energy level E_N, where N is the peak index [31] (see Figure A.2 for an example). The quantised energy levels in graphene are at

$$E_N(B) = \sqrt{2NeBv_f^2\hbar}$$

therefore

$$B_N = \frac{E_f^2}{2Nev_f^2\hbar} = \frac{\hbar k_f^2}{2Ne} = \frac{hn}{4Ne}$$

and hence the carrier density n can be found from a plot of N against $1/B_N$ [34].

The positions of the maxima in $1/B$ at $V_g = 0$ V in the EFM data do indeed give a linear plot (inset to Figure 5.11); a resistance maximum arises between Landau levels because transport is suppressed away from the flake edges, so it is reasonable for an EFM measurement to detect the corresponding reduction in screening. EFM is therefore a companion measurement to magnetoresistance. The gradient of the line gives a carrier density of 1.47×10^{11} cm^{-2}.
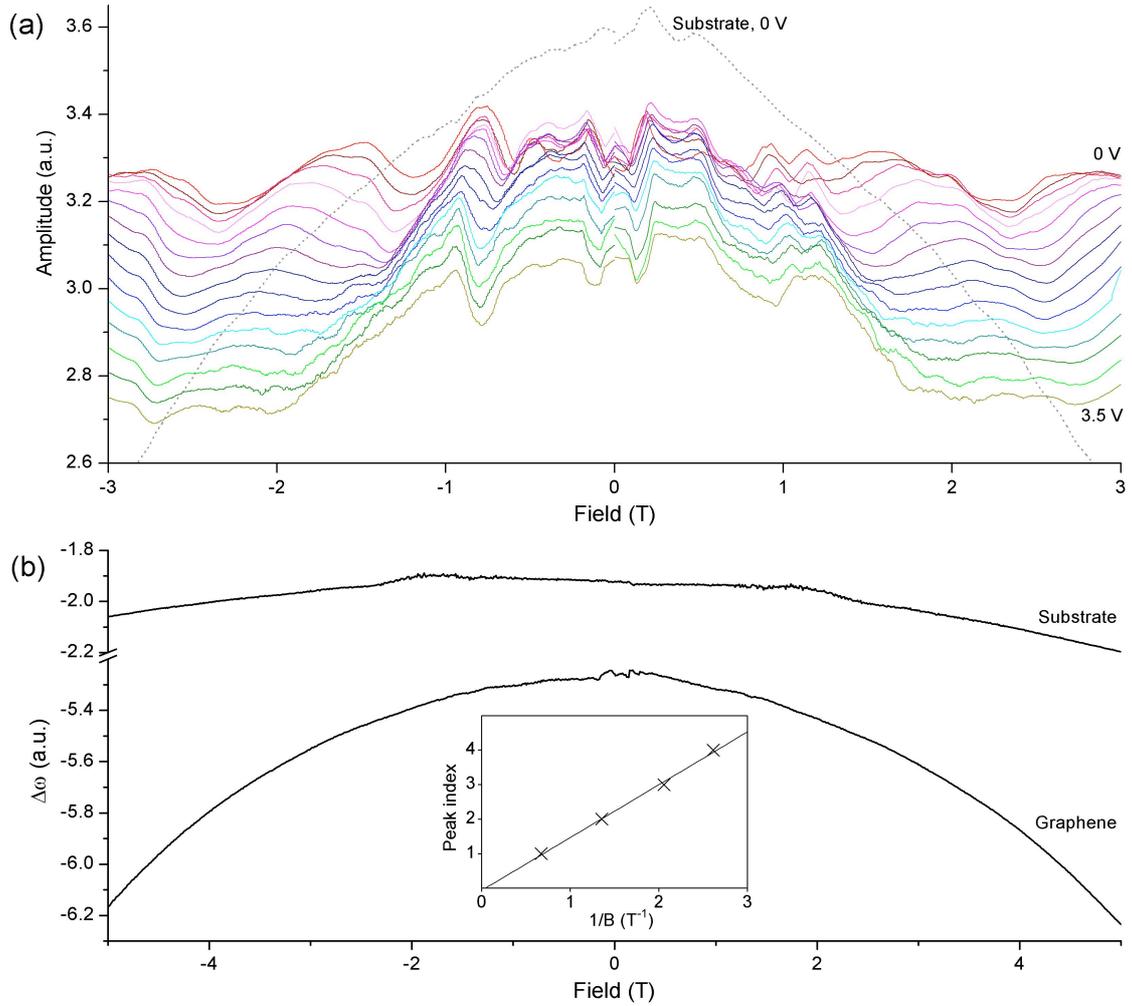
Figure 5.11: Shubnikov-de Haas oscillations in EFM spot measurements.
(a) The graphene amplitude signal over magnetic field sweeps at DC bias intervals
of 0.25 V. The dotted line is the background signal for the substrate. (b) $\Delta\omega$ for
both graphene and the substrate at $\Delta V = 0$ V. (Inset) Plot of the peak positions in
the amplitude signal against $1/B$.

147

According to theory, the gate voltage induces a carrier density in graphene of $n = \epsilon_0 \epsilon V_g / te$) where ϵ_0 and ϵ are the permittivities of free space and SiO_2, respectively, and t is the oxide thickness (here 300 nm), giving an expected variation of $n = 7 \times 10^{10}$ cm^{-2} per volt [27]. Because the minimum conductance is finite this model is only valid away from the Dirac point. Taking the gradient for the peak spacings at $\Delta V = 2.25$ V gives $n = 1.69 \times 10^{11}$ cm^{-2}, which implies that the Dirac point is at $\Delta V = -0.02$ V. This is 0.5 V lower than the estimate for the bulk flake made in Section 5.5.1, a difference which is within the experimental error. Additionally the local variation of carrier density around $\Delta V = 2.25$ V is only 5.1×10^{10} cm^{-2} per volt, so the DC bias is still too low for an accurate estimate.

It was hoped to measure larger field and bias ranges over several spots to confirm this interpretation and to look for variations in the local carrier density. However, subsequent to this measurement, all electrical contact to the flake was lost and no further working devices were available for several months.

5.6 Conclusions

This project was brought to a halt prematurely because new devices could not be fabricated within a reasonable timeframe. Since then the AFM has been rebuilt with a static tip and an LED and fibre-optic cable for detecting deflection, which makes scanning easier and more reliable. The coarse positioners were replaced with a new scanning unit, which moves the sample instead of the tip, and allows space for a new package socket. This change is particularly important as chip packages no longer need to be soldered in place, and can therefore be measured with other cryostats. These improvements also remove some of the problems which led to constrained device design and failures. The rebuilt system has been used for SGM measurements which found local shifts in the Dirac point of a current-annealed flake, on micron length scales [258].

The low yield of working devices was disappointing, and repetition of all the measurements described here is desirable to confirm the conclusions drawn. Nevertheless, EFM seems to be an effective technique for measuring both bulk and local properties of graphene. Good qualitative conclusions were possible, despite the absence of lift mode operation to permit complete quantitative assessment - the shallow graphene thickness and height variation proved to have a much weaker effect on the EFM signal than the screening effect. A more critical problem was the residual PMMA which, in the absence of an effective cleaning process, determines the scanning resolution limit.

Where cleaning was achieved, interesting sub-micron scale electronic structure was detected. Concurrent with the work reported here, Martin et al. [282] have measured the local compressibility of graphene using an aluminium-based SET attached to the tip of a scanning probe. They found local variations in the Dirac point of order 2 V due to electron and hole puddles on a length scale of 150 nm, which was approximately equal to the spatial resolution of the SET. By comparison with bare silicon they deduced that the puddles could not be accounted for solely by trapped charge in the substrate oxide. In contrast Zhang et al. [283] performed a higher resolution experiment using STM and imaged charge puddles on a scale of 20 nm, which did correlate with charge donors beneath the graphene. The local variation in compressibility detected here by EFM corresponds more closely to the latter, and did not correlate with the surface texture.

Even though the range of magnetic field sweeps for spot measurement was small, the realistic estimate of the carrier density obtained is strong evidence that the fluctuations in the EFM amplitude signal correspond to Shubnikov-de Haas oscillations in the flake. Spot EFM measurements therefore appear to be a good indicator of the local carrier concentration and, used together with bulk transport data over a large number of spots, could be used to build a map of the shape and size of electron and hole puddles.

The weaknesses of the measurements achieved here mean that the observations of local variation in the graphene screening and carrier density are not conclusive. However, improvements can easily be made to the measurement process, and this experiment has shown that EFM is a useful and potentially very informative tool for use in studies of graphene.

Chapter 6

Conclusions

The two principle chapters of this thesis describe two independent but related experiments. Although using different substrates and technologies, both aimed primarily to fabricate functioning low-dimensional devices using recently discovered carbon structures. The final goal in both was specific electronic transport behaviour. Many fabrication tools and experimental techniques were shared by the projects, but their role and emphasis differed. With carbon nanotubes device design became the primary challenge. New modifications were needed for fabrication processes, AFMs were used solely for imaging and low temperature measurement was the indicator of success. With graphene an AFM became both fabrication tool and measurement apparatus, laboratory space and experimental interface. Little modification of clean room stages was needed and (low temperature) electrical measurement was a secondary tool. While experimental skills and technical knowledge could be carried between the experiments, the steps of measurement and fabrication were specific to each.

Nearly 20 years after they were brought to widespread attention, carbon nanotube research is now a maturing field. The number of papers involving nanotubes published on arXiv.org has been roughly constant at a little over 200 per year throughout the work of this thesis. The limits of handling and measuring nanotubes are increasingly well understood, the targets for their use are becoming more realistic, and the applications more refined. This thesis aimed to add to those applications by introducing nanotubes to a single new element, that is a new substrate, and to the extent that functioning devices were produced, it has been successful.

In contrast, publishing on graphene continues to increase exponentially, with nearly as many papers currently being posted online in 2 months as were released in the whole of 2006, and four times as many as on nanotubes. After years of theoretical

predictions, the actual properties and interactions of graphene are now only beginning to be understood. This unique material requires a new set of tools to interact with it, and an AFM may be the first which can manipulate graphene on its functional scale. If nanotechnology is the ability to interact directly with materials on the nanoscale, not just the existence of nanometre-sized features, then AFM lithography is one of the few true nanoscale technologies.

In both projects, the physical interactions between materials emerged to play a dominant role. With nanotubes, the interface with a GaAs substrate influenced both fabrication and transport in unexpected ways, and the pervasive presence of water in connection with graphene turned out to be critical for lithography as well as influential for transport. It is with regard to these relationships and their importance that the projects departed from their initial goals, and the degree to which the experimental approach was redirected determined the success of the outcomes.

6.1 Carbon nanotube devices on GaAs

This work successfully fabricated and measured nanotube devices with GaAs substrates, but the yield and performance of the devices themselves was disappointing. The experiment and its aims were planned based on the assumption that dispersion techniques used for nanotubes on silicon would apply equally on GaAs. The approach taken was then poorly prepared for the reality, and the facilities available for the dispersing process were not wholly adequate for a detailed quantitative study. The most significant mistake of the project was to press on regardless, but the subsequent solutions to the problems created by the poor dispersions are key achievements of the work. That devices with basic functionality were nonetheless completed is a proof of the feasibility of the project's aims.

This project prompted a brief investigation of the nanoscale surface morphology of GaAs, and also consideration of the influence of a GaAs substrate on nanotube transport. Both questions are the subject of relatively little study in the literature, and theoretical modelling specifically relating nanotubes to GaAs would be invaluable in supporting any future progress in this field. Limited conclusions were made regarding the comparatively weak binding energy and the inducement of disorder by the GaAs, by combining observations from the effectiveness of dispersions, AFM manipulations and transport measurements, which might otherwise have been considered unrelated. The final design and fabrication process was relatively efficient, and included a large

number of measures to ensure reliability. It is a shame that while these may have all worked, the devices could not be measured as a result of inappropriate storage.

Transport measurements could have been better planned, but ultimately the device robustness set the main limit on the quantity of data collected. Importantly however, enough evidence was gathered to prove that the basic premise behind the experiment, that transport in nanotubes and a nearby GaAs 2DEG would be mutually dependent, was valid and simple to achieve. Subject to the lessons learned from this work, and the modifications proposed, it seems very probable that detection of charge transport between the two systems can be achieved without any further developments in fabrication.

6.2 AFM lithography of graphene

This experiment was proposed without any knowledge of the preceding STM work with graphite, but even considered as a whole the literature was not a conclusive or complete body, after 15 years of work. The step from lithography of graphite to graphene turned out to be much less difficult than expected - in fact the process eventually transferred almost without modification. However, the first attempts here on graphene were ambiguous precisely because the process was poorly understood, and further study on graphite proved invaluable.

A totally complete parametric study was not achieved because the initial abundance of variables was too great. However, progress was made, particularly regarding the quantitative effect of humidity and the expected cut profile. Current measurements were important in determining that the process is electric field dependent, thus narrowing the field of candidate mechanisms. The actual lithography achieved exceeded the target specifications, and proved in a single experiment that the technique is both viable and effective for defining device structures in graphene. Though not reported here, actual devices fabrication requires no further development, and its demonstration is an important step for the immediate future.

The comparative simplicity of the experimental goals also allowed more flexibility to investigate related observations, and the detection of water beneath graphene on silicon substrates is of benefit to wider studies of its transport and morphology. AFM lithography also proved its effectiveness in comparison with the established e-beam and etching technologies, and has the potential to become a common, accessible and powerful tool for graphene research and, potentially, commercial exploitation.

6.3   Two new platforms for low dimensional devices

Research into low dimensional electron transport is fundamental to the future of computer technology, and both these projects address the limits of the field directly. There is significant overlap in the capabilities of the materials. As already discussed, the energy scales and structure of nanotubes more readily avails them for 1D transport than graphene, but the latter gains advantages when more than a few channels are required. The methods for defining 0D regions in both materials are equally effective and require comparable effort. However, the key advantage of graphene is that one continuous flake can be patterned to contained multiple 2D, 1D and 0D regions and interconnects all in a single lithography step. In contrast many nanotubes are required to form a multi-component circuit, and no single process for defining barriers and interconnects of suitably small sizes is yet scalably compatible with nanotube dispersions.

Of the two projects, AFM lithography of graphene has achieved the more significant developments relative to the existing technologies, and offers the more general range of applications. While GaAs as a substrate for nanotube devices offers certain advantages over silicon, they were not expressed here, and no new measurement technique or device function was demonstrated at this stage that cannot already be achieved with proven device structures. Without any substantial modifications to the microscope or additional processing of the graphene, AFM lithography immediately matched the very best abilities of the prior techniques. Thus future progress might be expected to overtake the limits of chemical and plasma etching altogether, plus the basic facilities needed for AFM lithography are common and non-specialist. Graphene also has the benefit of intense focus within the field, and the better chance of imminent developments which will further enhance fabrication capabilities.

In order to achieve the experimental aims these projects have linked a wide range of scientific fields and techniques. Carbon nanotubes and GaAs, and graphene and LAO lithography were both combined for the first time. The latter in particular was demonstrated as a useful and effective technique, and a detailed study of the lithography process supports it as a reliable and valuable addition to the field.

Appendix A

Device Patterns and Designs

A.1 Nanotube and GaAs device patterns

Figure A.1: Hall bar pattern. Coloured lines are feature outlines and filled colours show the topmost features: mesa, ohmic contacts and alignment features.
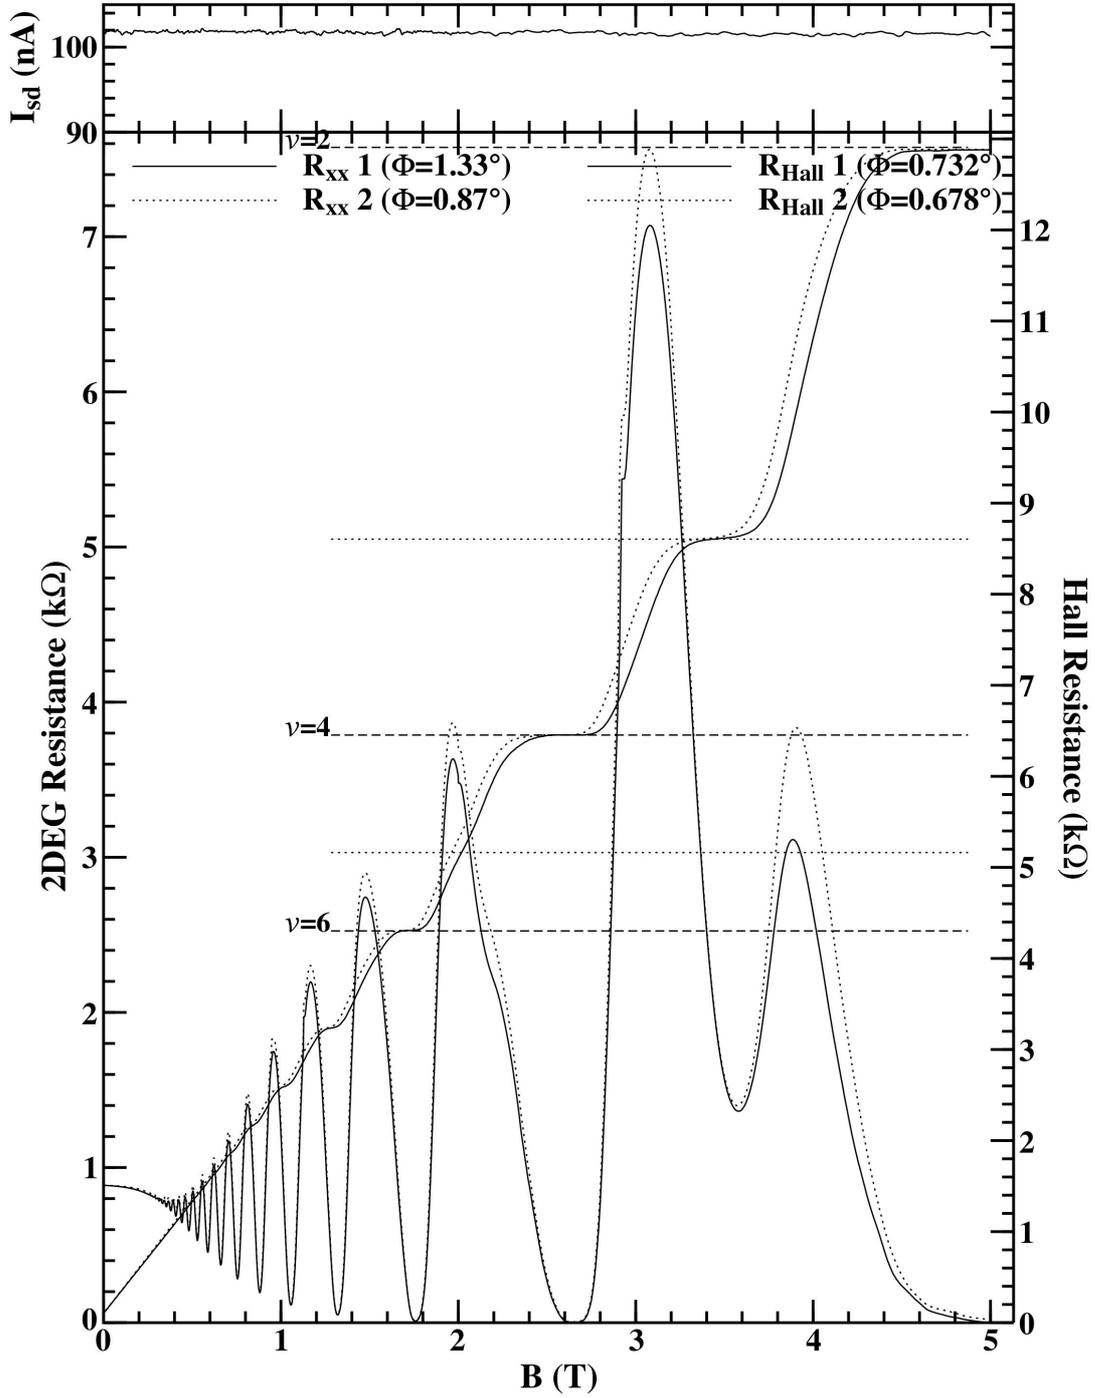
Figure A.2: Measurements of the quantum Hall effect. Assessment results, at 1.5 K in the dark, for a Hall bar annealed to around 500°C in the CAPE nanotube CVD chamber, showing quantum Hall plateaux in the Hall resistance and Shubnikov-de Haas oscillations in the 2DEG ('diagonal') resistance.

156
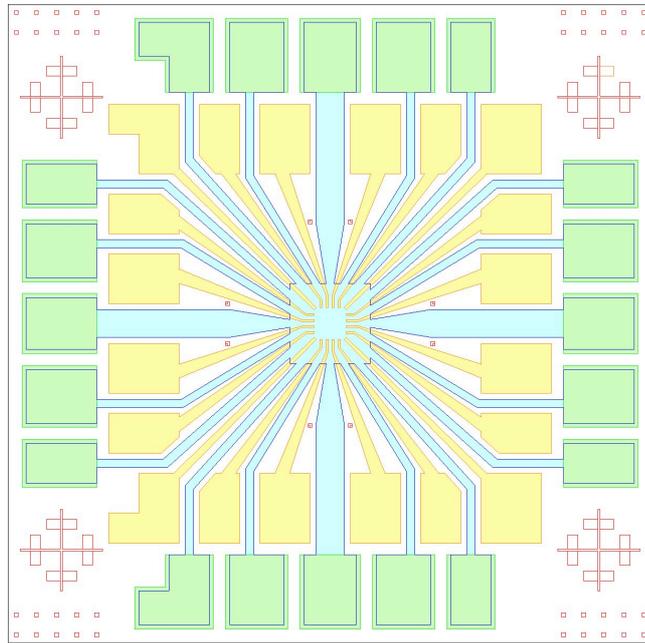
Figure A.3: The mesa pattern used for nanotube devices. Colours as for Fig. A.1 plus nanotube contacts. The pattern is 2 mm wide.
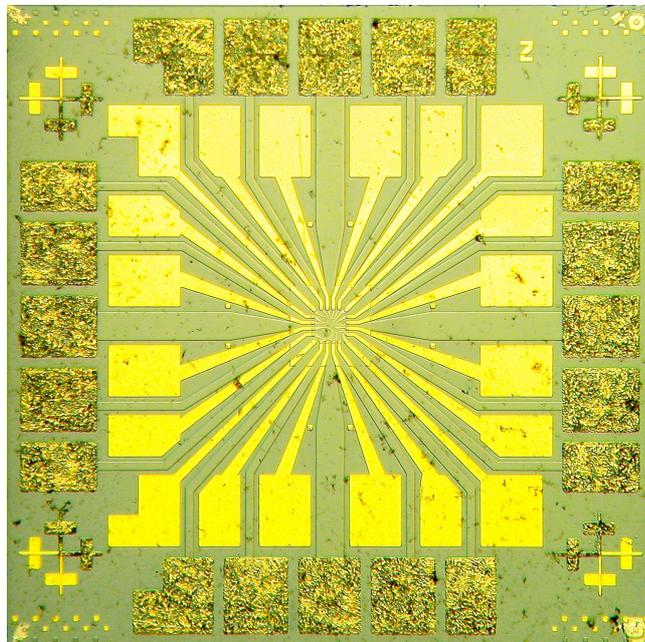
Figure A.4: A completed mesa device, including nanotube contacts. The dark specks are large nanotube bundles.

157

Figure A.5: First generation nanotube device design (top) with nanotubes, e-beam written contacts and the alignment grid, and (bottom) the completed device.
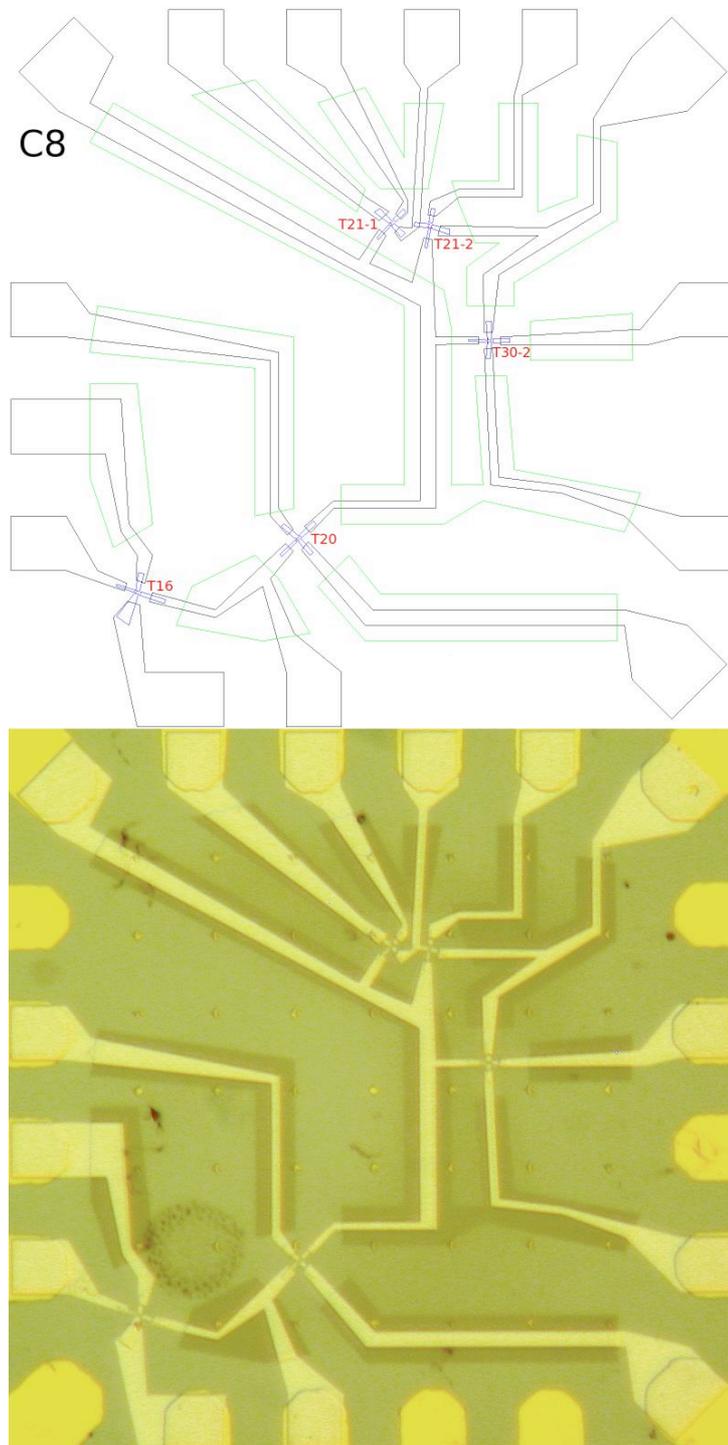
Figure A.6: Second generation nanotube device design (top) for a 3 stage process: (i) nanotube contacts (ii) cross-linked resist and (iii) connections to the optical features, and (bottom) the completed device.

A.2   Completed nanotube devices

Devices are labelled according to the mesa and location, so E7t4-3 is nanotube no. 3 in AFM scan no. 4 of mesa no. 7 from chip E.

E1: initially all devices shorted to 2DEG

E1t5-2	0.7μm tube, 2.6μm channel	*No conduction*
E1t11-2	1.2μm tube, 1.4μm channel	*Tube broken*
E1t12-2	1.2μm tube, 1.4μm channel	*Tube broken*
E1t18	0.8μm tube, 3μm channel	*Bonding failed*
E1t24-1/2	Crossed 2.3μm and 1.3μm tubes	*Blew up*
E1t24-3	1.9μm tube, 2.5μm channel	*No conduction*
E1t27-1/2	Singly-contacted 1.3μm tube touching	*Side tube broken,*
	1.7μm tube, side gate at 1μm	*bonding failed*
E1t27-3	1.4μm tube, 5μm channel	*Tube removed*

E2:

E2t1-2	1.1μm tube, 3μm channel	***Measured***
E2t4-1	0.8μm tube, 0.7μm channel	*No conduction, blew up*
E2t4-2/3	Crossed 0.8μm and 1.2μm tubes	*No conduction, blew up*
E2t8	1.5μm tube, 0.9μm channel	***Measured****, blew up*
E2t9-1	1μm tube, 1.9μm channel	*Blew up*
E2t9-2	0.7μm tube, 1.4μm channel	*Blew up*
E2t10-3	0.9μm kinked tube, 0.9μm channel	*Bonding failed*
E2t13a/b	Triply-contacted 0.7μm+1.3μm	***Measured****, blew up*
	sections, 2.4μm channel	

E6: initially all devices shorted to 2DEG

E6t3-1	0.7μm tube, long channel	*Tube removed*
E6t3-3a/b	Triply-contacted 0.7μm+1μm sections, 1μm channel	*Tube broken*
E6t6-1	1.9μm tube, side gate at 0.6μm	*Bonding failed*
E6t6-2	0.8μm tube	*Bonding failed*
E6t25-2a/b	Triply-contacted 1μm+2μm sections	*Middle bond failed, blew up*
E6t25-3	Singly-contacted 3.3μm tube touching 1.2μm kinked tube, side gate at 0.4μm	*Middle bond failed, blew up*
E6t28	2μm tube, 0.3μm wide side gates at 0.7μm and 0.6 μm, 2μm channel	*Tube removed*

E7:

E7t0-1	0.9μm tube, 1.8μm channel	*Bonding failed*
E7t0-2	2 parallel 1.8μm tubes, 0.3μm apart, side gate at 0.3μm to top tube	*Bonding failed*
E7t4-1	1.2μm tube	*Blew up*
E7t4-3	0.7μm tube	***Measured***, *blew up*
E7t4-4	0.7μm kinked tube	*Blew up*
E7t12	0.8μm tube, 1.3μm channel	*Tube broken*

C5: 4 devices with 300 nm contacted nanotube length and 750 nm wide split gate. *No conduction, 4 blew up.*

C6: 4 devices with 300 nm contacted nanotube length and 750 nm wide split gate. *No conduction, 1 blew up.*

C7: 4 devices with 300 nm contacted nanotube length and 750 nm wide split gate. *No conduction, 3 blew up.*

C8: 5 devices with 300 nm contacted nanotube length and 750 nm wide split gate. *No conduction, 3 blew up.*

A.3 Graphene device patterns
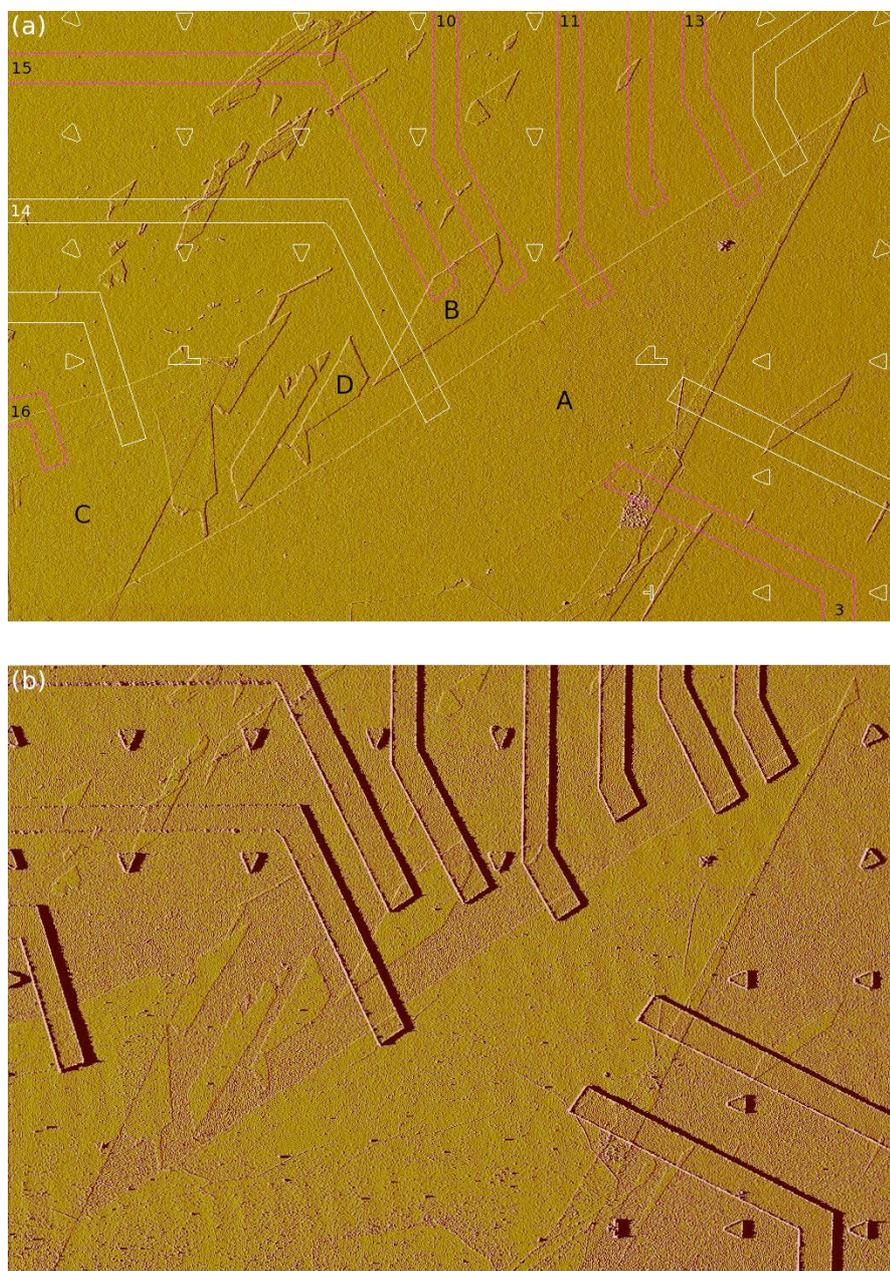
Figure A.7: Contacts to a large graphene flake. (a) Pattern for device S2211. Contact and flake labels are referred to in the text. The pink contacts were successfully bonded and soldered in the fridge. The field of view is 57 μm wide. (b) The same region after metal evaporation. Note the increased roughness due to residual PMMA after an acetone wash.
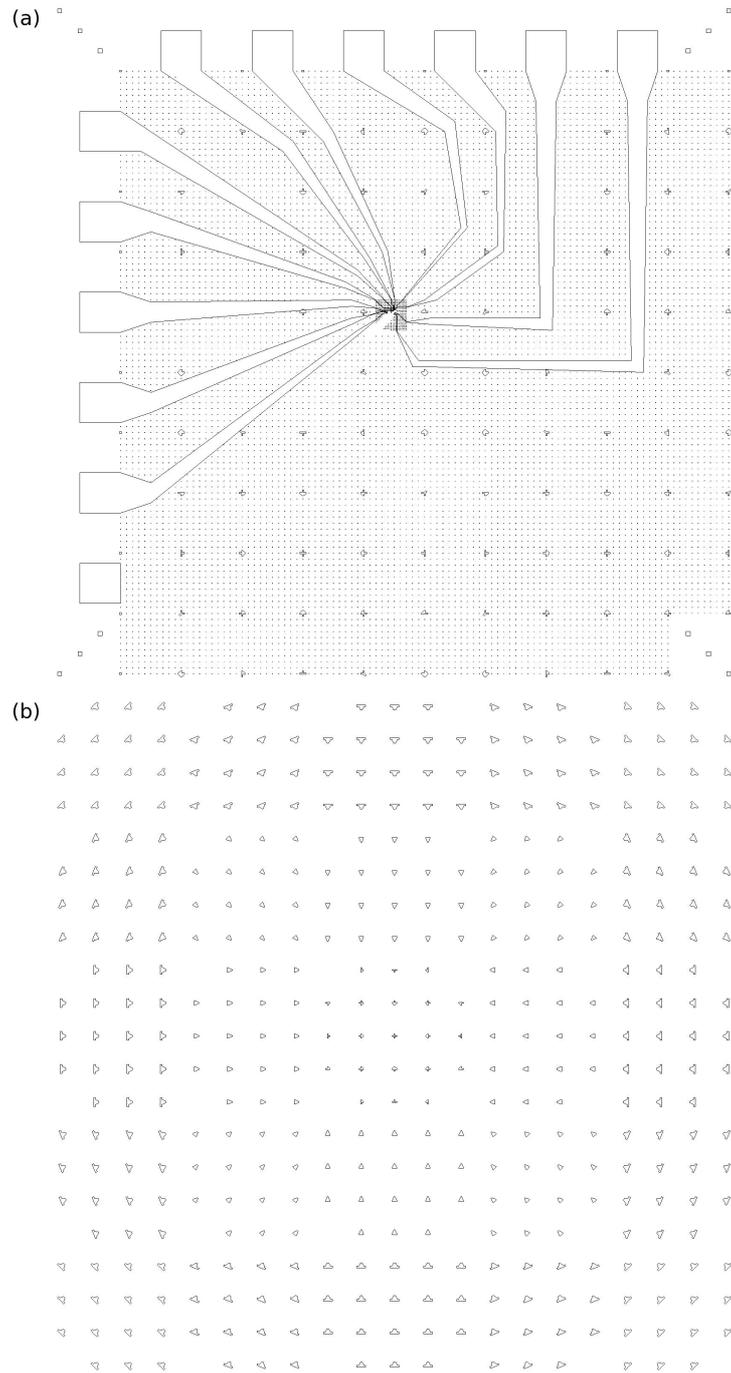
Figure A.8: Alignment mark pattern for the low temperature AFM. (a) A complete device pattern 3.7 mm wide, with bond pads and flake contacts. The shapes of the optical alignment marks every 100 μm are visible, with smaller copies in the adjacent squares. (b) The low temperature alignment grid, seen surrounding the device above, with marks every 7.5 μm.

A.4 Completed graphene devices

Each chip is numbered; additional letters are for multiple flakes on one chip. MFC refers to multiple finger contacts, MEC to multiple edge contacts, SC to a singly contacted flake, (C) to a CAPE flake and (G) to a Graphene Industries flake. Sizes are for an approximated rectangular shape.

AFM lithography:

G01a (C)	$15 \times 6\mu$m MFC FLG	**Cutting**, *tip crash*
G01b (C)	$10 \times 9\mu$m MFC SLG	**First SLG cutting**, *tip crash*
G04a (C)	$4 \times 2\mu$m MEC FLG	*Cuts ambiguous, tip crash*
G04b (C)	$4 \times 1\mu$m MEC FLG	*Cuts ambiguous*
G04c (C)	$4 \times 2\mu$m SC BLG	**Cutting**
G05	Graphite	**Cutting**
G06	Graphite	**Cutting**
G07 (C)	$4 \times 2\mu$m MEC SLG	**SLG cutting**
G08	$11 \times 1.5\mu$m MFC SLG	**SLG cutting**
G09 (C)	$4 \times 1\mu$m SC SLG	*Cuts ambiguous*
G10a (C)	$1.5 \times 0.5\mu$m SLG	*Annealed, blew up during soldering*
G10b (C)	$2.5 \times 1.5\mu$m SLG	*Annealed, blew up during soldering*
G11a (C)	$6 \times 2\mu$m SC FLG	**First cutting**, *blew up during probing*
G11b (C)	Two $8 \times 1\mu$m MFC SLGs	*Tip crash, blew up during probing*
G40 (C)	$33 \times 17\mu$m MEC SLG	*Destroyed during lift-off*
GPE04 (C)	$4 \times 2\mu$m MFC FLG	*Tip crash, annealed*

Low-temperature EFM:

S2211 (G)	$55 \times 17\mu$m SLG	**Measured**, *tip crash*
G23 (G)	$35 \times 7\mu$m SLG	*Destroyed during resist-strip*
G24 (G)	$10 \times 7\mu$m SLG	**Measured**
G25 (G)	$28 \times 12\mu$m SLG	*Tip crash*
G26a (G)	$21 \times 7\mu$m SLG	*E-beam misaligned*
G26b (G)	$12 \times 10\mu$m BLG	*Back gate leaked*
G27 (G)	$37 \times 21\mu$m SLG	*Back gate leaked*

References

[1] H. W. Kroto, J. R. Heath, S. C. O'Brien, R. F. Curl, and R. E. Smalley. C_{60}: *Buckminsterfullerene.* Nature, **318** 162–163, 1985.

[2] C. G. Smith. *Low-dimensional quantum devices.* Rep. Prog. Phys., **59** 235–282, 1996.

[3] H. E. Beere. *Selective area growth of III-V semiconductor compounds using Ga^+ FIB deposition during MBE growth.* PhD thesis, University of Cambridge, 1999.

[4] K. v. Klitzing, G. Dorda, and M. Pepper. *New Method for High-Accuracy Determination of the Fine-Structure Constant Based on Quantized Hall Resistance.* Phys. Rev. Lett., **45**(6) 494–497, 1980.

[5] C. W. J. Beenakker and H. van Houten. *Quantum Transport in Semiconductor Nanostructures.* Solid State Physics, **44** 1, 1991.

[6] R. Crook. *Combining a Low-Temperature Scanning Charged Probe with Low-Dimensional Semiconductor Devices.* PhD thesis, University of Cambridge, 2000.

[7] J. Nygård. *Experiments on Mesoscopic Electron Transport in Carbon Nanotubes.* PhD thesis, University of Copenhagen, 2000.

[8] C. G. Smith, M. Pepper, H. Ahmed, J. E. F. Frost, D. G. Hasko, D. C. Peacock, D. A. Ritchie, and G. A. C. Jones. *The transition from one-to zero-dimensional ballistic transport.* J. Phys. C: Solid State Physics, **21**(24) L893–L898, 1988.

[9] L. I. Glazman and R. I. Shekhter. *Coulomb oscillations of the conductance in a laterally confined heterostructure.* J. Phys.: Cond. Matt., **1** 5811–5815, 1989.

[10] P. Leek. *Charge Transport in Carbon Nanotubes using Surface Acoustic Waves.* PhD thesis, University of Cambridge, 2006.

REFERENCES

[11] H. C. Casey and M. B. Panish. *Heterostructure Lasers. Part A: Fundamental principles.* Academic Press, 1978.

[12] J. J. Harris, C. T. Foxon, D. E. Lacklison, and K. W. J. Barnham. *Scattering mechanisms in (Al,Ga)As/GaAs 2DEG structures.* Superlattices and Microstructures, **2**(15) 563–568, 1986.

[13] M. A. Herman and H. Sitter. *Molecular Beam Epitaxy, Fundamentals and Current Status.* Springer Series in Material Science 7, 1988.

[14] T. J. Thornton, M. Pepper, H. Ahmed, D. Andrews, and G. J. Davies. *One-Dimensional Conduction in the 2D Electron Gas of a GaAs-AlGaAs Heterojunction.* Phys. Rev. Lett., **56**(11) 1198–1201, 1986.

[15] D. A. Wharam, T. J. Thornton, R. Newbury, M. Pepper, H. Ahmed, J. E. F. Frost, D. G. Hasko, D. C. Peacock, D. A. Ritchie, and G. A. C. Jones. *One-dimensional transport and the quantisation of the ballistic resistance.* J. Phys. C: Solid State Physics, **21**(8) L209–L214, 1988.

[16] M. Field, C. G. Smith, M. Pepper, D. A. Ritchie, J. E. F. Frost, G. A. C. Jones, and D. G. Hasko. *Measurements of Coulomb Blockade with a Noninvasive Voltage Probe.* Phys. Rev. Lett., **70**(9) 1311–1314, 1993.

[17] M. Field, C. G. Smith, M. Pepper, K. M. Brown, E. H. Linfield, M. P. Grimshaw, D. A. Ritchie, and G. A. C. Jones. *Coulomb Blockade as a Noninvasive Probe of Local Density of States.* Phys. Rev. Lett., **77**(2) 350–353, 1996.

[18] T. W. Odom, J. L. Huang, and C. M. Lieber. *STM studies of single-walled carbon nanotubes.* J. Phys.: Cond. Matt., **14**(6) 143–167, 2002.

[19] L. V. Radushkevich and V. M. Lukyanovich. О структуре углерода, образующегося при термическом разложении окиси углерода на железном контакте. J. of Phys. Chem. (USSR), **26** 88–95, 1952.

[20] S. Iijima. *Helical microtubules of graphitic carbon.* Nature, **354** 56–58, 1991.

[21] M. Monthioux and V. L. Kuznetsov. *Who should be given the credit for the discovery of carbon nanotubes?* Carbon, **44** 1621–1624, 2006.

[22] S. Iijima and T. Ichihashi. *Single-shell carbon nanotubes of 1-nm diameter.* Nature, **363** 603–605, 1993.

[23] D. S. Bethune, C. H. Klang, M. S. de Vries, G. Gorman, R. Savoy, J. Vazquez, and R. Beyers. *Cobalt-catalysed growth of carbon nanotubes with single-atomic-layer walls.* Nature, **363** 605–607, 1993.

[24] L. X. Zheng, M. J. O'Connell, S. K. Doorn, X. Z. Liao, Y. H. Zhao, E. A. Akhadov, M. A. Hoffbauer, B. J. Roop, Q. X. Jia, et al. *Ultralong single-wall carbon nanotubes.* Nature Materials, **3** 673–676, 2004.

[25] A. Thess, R. Lee, P. Nikolaev, H. Dai, P. Petit, J. Robert, C. Xu, Y. Lee, S. Kim, et al. *Crystalline Ropes of Metallic Carbon Nanotubes.* Science, **273** 483, 1996.

[26] E. Wong, P. Sheehan, and C. Lieber. *Nanobeam Mechanics: Elasticity, Strength, and Toughness of Nanorods and Nanotubes.* Science, **277** 1971, 1997.

[27] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov. *Electric Field Effect in Atomically Thin Carbon Films.* Science, **306** 666–669, 2004.

[28] A. K. Geim and K. S. Novoselov. *The rise of graphene.* Nature Materials, **6**(3) 183–191, 2007.

[29] A. Krishnan, E. Dujardin, M. M. J. Treacy, J. Hugdahl, S. Lynum, and T. W. Ebbesen. *Graphitic cones and the nucleation of curved carbon surfaces.* Nature, **388** 451–454, 1997.

[30] S. Cho and M. S. Fuhrer. *Charge transport and inhomogeneity near the minimum conductivity point in graphene.* Phys. Rev. B, **77**(8) 81402, 2008.

[31] C. Berger, Z. Song, X. Li, X. Wu, N. Brown, D. Maud, C. Naud, and W. de Heer. *Magnetotransport in high mobility epitaxial graphene.* physica status solidi (a), **204**(6) 1746–1750, 2007.

[32] J. B. Oostinga, H. B. Heersche, X. Liu, A. F. Morpurgo, and L. M. K. Van-dersypen. *Gate-induced insulating state in bilayer graphene devices.* Nature Materials, **7**(2) 151–157, 2007.

[33] B. Özyilmaz, P. Jarillo-Herrero, D. Efetov, and P. Kim. *Electronic transport in locally gated graphene nanoconstrictions.* Appl. Phys. Lett., **91** 192107, 2007.

REFERENCES

[34] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, M. I. Katsnelson, I. V. Grigorieva, S. V. Dubonos, and A. A. Firsov. *Two-dimensional gas of massless Dirac fermions in graphene.* Nature, **438** 197–200, 2005.

[35] K. S. Novoselov, Z. Jiang, Y. Zhang, S. V. Morozov, H. L. Stormer, U. Zeitler, J. C. Maan, G. S. Boebinger, P. Kim, and A. K. Geim. *Room-Temperature Quantum Hall Effect in Graphene.* Science, **315** 1379, 2007.

[36] T. W. Odom, H. Jin-Lin, P. Kim, and C. M. Lieber. *Structure and Electronic Properties of Carbon Nanotubes.* J. of Phys. Chem. B, **104**(13) 2794–2809, 2000.

[37] Y. Zhang, Y. W. Tan, H. L. Stormer, and P. Kim. *Experimental observation of quantum Hall effect and Berry's phase in graphene.* Nature, **438** 201–204, 2005.

[38] R. Saito, M. Fujita, G. Dresselhaus, and M. S. Dresselhaus. *Electronic structure of chiral graphene tubules.* Appl. Phys. Lett., **60**(18) 2204–2206, 1992.

[39] N. Hamada, S. Sawada, and A. Oshiyama. *New One-Dimensional Conductors: Graphitic Microtubules.* Phys. Rev. Lett., **68**(10) 1579–1581, 1992.

[40] J. W. Mintmire, B. I. Dunlap, and C. T. White. *Are Fullerene Tubules Metallic?* Phys. Rev. Lett., **68**(5) 631–634, 1992.

[41] C. L. Kane and E. J. Mele. *Size, Shape, and Low Energy Electronic Structure of Carbon Nanotubes.* Phys. Rev. Lett., **78**(10) 1932–1935, 1997.

[42] J. W. G. Wildoer, L. C. Venema, A. G. Rinzler, R. E. Smalley, and C. Dekker. *Electronic structure of atomically resolved carbon nanotubes.* Nature, **391** 59–62, 1998.

[43] C. White and T. Todorov. *Carbon nanotubes as long ballistic conductors.* Nature, **393** 240–242, 1998.

[44] Z. Yao, C. Dekker, and P. Avouris. *Electrical Transport Through Single-Wall Carbon Nanotubes.* Topics in Applied Physics, **80** 147–171, 2001.

[45] S. J. Tans, M. H. Devoret, H. Dai, A. Thess, R. E. Smalley, L. J. Geerligs, and C. Dekker. *Individual single-wall carbon nanotubes as quantum wires.* Nature, **386** 474–477, 1997.

[46] P. Avouris, Z. Chen, and V. Perebeinos. *Carbon-based electronics.* Nature Nanotechnology, **2**(10) 605–615, 2007.

[47] S. J. Tans, M. H. Devoret, R. J. A. Groeneveld, and C. Dekker. *Electron-electron correlations in carbon nanotubes.* Nature, **394** 761–764, 1998.

[48] A. Bachtold, P. Hadley, T. Nakanishi, and C. Dekker. *Logic Circuits with Carbon Nanotube Transistors.* Science, **294**(5545) 1317–1320, 2001.

[49] C. Schönenberger. *Charge and spin transport in carbon nanotubes.* Semicond. Sci. Technol., **21**(11) 1–9, 2006.

[50] M. R. Gräber, M. Weiss, S. Oberholzer, and C. Schönenberger. *Defining and controlling double quantum dots in single-walled carbon nanotubes.* Semicond. Sci. Technol., **21**(11) 64–68, 2006.

[51] A. Y. Kasumov, R. Deblock, M. Kociak, B. Reulet, H. Bouchiat, I. I. Khodos, Y. B. Gorbatov, V. T. Volkov, C. Journet, and M. Burghard. *Supercurrents Through Single-Walled Carbon Nanotubes.* Science, **284** 1508–1511, 1999.

[52] L. E. Hueso, G. Burnell, J. L. Prieto, L. Granja, C. Bell, D. J. Kang, M. Chhowalla, S. N. Cha, J. E. Jang, G. A. J. Amaratunga, and N. D. Mathur. *Electrical transport between epitaxial manganites and carbon nanotubes.* Appl. Phys. Lett., **88** 083120, 2006.

[53] D. Deutsch. *Quantum theory, the Church-Turing principle and the universal quantum computer.* Proc. R. Soc. Lond. A, **400** 97–117, 1985.

[54] D. P. DiVincenzo. *The Physical Implementation of Quantum Computation.* Fortschr. Phys., **48**(9-11) 771–783, 2000.

[55] P. W. Shor. *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer.* SIAM Review, **26** 303–332, 1997.

[56] L. K. Grover. *Quantum Mechanics Helps in Searching for a Needle in a Haystack.* Phys. Rev. Lett., **79** 325–328, 1993.

[57] A. W. Rushforth, C. G. Smith, M. D. Godfrey, H. E. Beere, D. A. Ritchie, and M. Pepper. *Noninvasive detection of the evolution of the charge states of a double dot system.* Phys. Rev. B, **69** 113309, 2004.

REFERENCES

[58] C. G. Smith, S. Gardelis, A. W. Rushforth, R. Crook, J. Cooper, D. A. Ritchie, E. H. Linfield, Y. Jin, and M. Pepper. *Realization of quantum-dot cellular automata using semiconductor quantum dots.* Superlattices and Microstructures, **34**(3-6) 195–203, 2003.

[59] T. Fujisawa, T. Hayashi, H. D. Cheong, Y. H. Jeong, and Y. Hirayama. *Rotation and phase-shift operations for a charge qubit in a double quantum dot.* Physica E, **21**(2-4) 1046–1052, 2004.

[60] A. G. Huibers, J. A. Folk, S. R. Patel, C. M. Marcus, C. I. Duruöz, and J. S. Harris Jr. *Low-Temperature Saturation of the Dephasing Time and Effects of Microwave Radiation on Open Quantum Dots.* Phys. Rev. Lett., **83**(24) 5090–5093, 1999.

[61] S. Roche, J. Jiang, F. Triozon, and R. Saito. *Quantum Dephasing in Carbon Nanotubes due to Electron-Phonon Coupling.* Phys. Rev. Lett., **95** 076803, 2003.

[62] J. Appenzeller, R. Martel, P. Avouris, H. Stahl, U. T. Hunger, and B. Lengeler. *Phase-coherent transport in ropes of single-wall carbon nanotubes.* Phys. Rev. B, **64**(12) 121404, 2001.

[63] H. I. Jørgensen, K. Grove-Rasmussen, J. R. Hauptmann, and P. E. Lindelof. *Single wall carbon nanotube double quantum dot.* Appl. Phys. Lett., **89** 232113, 2006.

[64] J. Fischer, B. Trauzettel, and D. Loss. *Hyperfine interaction and electron-spin decoherence in graphene and carbon nanotube quantum dots.* Phys. Rev. B, **80**(15) 155401, 2009.

[65] T. W. Ebbesen and P. M. Ajayan. *Large-scale synthesis of carbon nanotubes.* Nature, **358** 220–222, 1992.

[66] C. Journet, W. Maser, P. Bernier, A. Loiseau, M. de La Chapelle, S. Lefrant, P. Deniard, R. Leek, and J. Fischerk. *Large-scale production of single-walled carbon nanotubes by the electric-arc technique.* Nature, **388** 756, 1997.

[67] T. Guo, P. Nikolaev, A. Thess, D. T. Colbert, and R. E. Smalley. *Catalytic growth of single-walled manotubes by laser vaporization.* Chem. Phys. Lett., **243**(1-2) 49–54, 1995.

[68] T. Guo, P. Nikolaev, A. G. Rinzler, D. Tománek, D. T. Colbert, and R. E. Smalley. *Self-assembly of tubular fullerenes.* J. Phys. Chem., **99**(27) 10694–10697, 1995.

[69] M. José-Yacamán, M. Miki-Yoshida, L. Rendon, and J. G. Santiesteban. *Catalytic growth of carbon microtubules with fullerene structure.* Appl. Phys. Lett., **62** 657, 1993.

[70] M. J. Bronikowski, P. A. Willis, D. T. Colbert, K. A. Smith, and R. E. Smalley. *Gas-phase production of carbon single-walled nanotubes from carbon monoxide via the HiPco process: A parametric study.* J. Vac. Sci. Technol. A, **19(4)** 1800–1805, 2001.

[71] J. Kong, H. T. Soh, A. M. Cassell, C. F. Quate, and H. Dai. *Synthesis of individual single-walled carbon nanotubes on patterned silicon wafers.* Nature, **395** 878–881, 1998.

[72] M. Su, Y. Li, B. Maynor, A. Buldum, J. P. Lu, and J. Liu. *Lattice-Oriented Growth of Single-Walled Carbon Nanotubes.* J. Phys. Chem. B, **104**(28) 6505–6508, 2000.

[73] N. Arjmandi, P. Sasanpour, and B. Rashidian. *CVD Synthesis of Small-Diameter Single-Walled Carbon Nanotubes on Silicon.* Scientia Iranica, Transactions D, **16**(1) 61–64, 2009.

[74] S. Maruyama, R. Kojima, Y. Miyauchi, S. Chiashi, and M. Kohno. *Low-temperature synthesis of high-purity single-walled carbon nanotubes from alcohol.* Chem. Phys. Lett., **360**(3-4) 229–234, 2002.

[75] T. Hertel, R. E. Walkup, and P. Avouris. *Deformation of carbon nanotubes by surface van der Waals forces.* Phys. Rev. B, **58**(20) 13870–13873, 1998.

[76] S. Sapmaz, P. Jarillo-Herrero, L. P. Kouwenhoven, and H. S. J. van der Zant. *Quantum dots in carbon nanotubes.* Semicond. Sci. Technol., **21**(11) 52–63, 2006.

[77] J. Nygård, D. H. Cobden, M. Bockrath, P. L. McEuen, and P. E. Lindelof. *Electrical transport measurements on single-walled carbon nanotubes.* Appl. Phys. A, **69**(3) 297–304, 1999.

REFERENCES

[78] J. Park, S. Rosenblatt, Y. Yaish, V. Sazonova, H. Üstünel, S. Braig, T. A. Arias, P. W. Brouwer, and P. L. McEuen. *Electron-Phonon Scattering in Metallic Single-Walled Carbon Nanotubes.* Nano Lett., **4**(3) 517–520, 2004.

[79] J. Cao, Q. Wang, and H. Dai. *Electron transport in very clean, as-grown suspended carbon nanotubes.* Nature Materials, **4**(10) 745–749, 2005.

[80] J. Nygård and D. H. Cobden. *Quantum dots in suspended single-wall carbon nanotubes.* Appl. Phys. Lett., **79** 4216–4218, 2001.

[81] M. Bockrath, D. H. Cobden, J. Lu, A. G. Rinzler, R. E. Smalley, L. Balents, and P. L. McEuen. *Luttinger-liquid behavior in carbon nanotubes.* Nature, **397** 598–601, 1999.

[82] A. Bezryadin, A. R. M. Verschueren, S. J. Tans, and C. Dekker. *Multiprobe Transport Experiments on Individual Single-Wall Carbon Nanotubes.* Phys. Rev. Lett., **80**(18) 4036–4039, 1998.

[83] T. Hertel, R. Martel, and P. Avouris. *Manipulation of Individual Carbon Nanotubes and their Interaction with Surfaces.* J. of Phys. Chem. B, **102**(6) 910–915, 1998.

[84] J. Lefebvre, J. F. Lynch, M. Llaguno, M. Radosavljevic, and A. T. Johnson. *Single-wall carbon nanotube circuits assembled with an atomic force microscope.* Appl. Phys. Lett., **75** 3014–3016, 1999.

[85] M. S. Fuhrer, J. Nygård, L. Shih, M. Forero, Y. G. Yoon, M. S. C. Mazzoni, H. J. Choi, J. Ihm, S. G. Louie, A. Zettl, and P. L. McEuen. *Crossed Nanotube Junctions.* Science, **288** 494–497, 2000.

[86] M. T. Woodside and P. L. McEuen. *Scanned Probe Imaging of Single-Electron Charge States in Nanotube Quantum Dots.* Science, **296** 1098–1101, 2002.

[87] A. Guillén-Cervantes, Z. Rivera-Alvarez, M. López-López, E. López-Luna, and I. Hernández-Calderón. *GaAs surface oxide desorption by annealing in ultra high vacuum.* Thin Solid Films, **373**(1-2) 159–163, 2000.

[88] Unidym. URL `http://www.unidym.com`.

[89] D. H. Cobden, M. Bockrath, P. L. McEuen, A. G. Rinzler, and R. E. Smalley. *Spin Splitting and Even-Odd Effects in Carbon Nanotubes.* Phys. Rev. Lett., **81**(3) 681–684, 1998.

[90] D.-H. Kim, J.-Y. Koo, and J.-J. Kim. *Cutting of multiwalled carbon nanotubes by a negative voltage tip of an atomic force microscope: A possible mechanism.* Phys. Rev. B, **68**(11) 113406, 2003.

[91] J. Nygård and D. H. Cobden. *Transport phenomena in nanotube quantum dots from strong to weak confinement.* arXiv, **cond-mat** 0105.289, 2001.

[92] Z. Yao, C. L. Kane, and C. Dekker. *High-Field Electrical Transport in Single-Wall Carbon Nanotubes.* Phys. Rev. Lett., **84**(13) 2941–2944, 2000.

[93] K. L. Lu, R. M. Lago, Y. K. Chen, M. L. H. Green, P. J. F. Harris, and S. C. Tsang. *Mechanical Damage of Carbon Nanotubes by Ultrasound.* Carbon, **34**(6) 814–816, 1996.

[94] T. S. Jespersen and J. Nygård. *Probing induced defects in individual carbon nanotubes using electrostatic force microscopy.* Appl. Phys. A, **88**(2) 309–313, 2007.

[95] A. Bachtold, M. Henny, C. Terrier, C. Strunk, C. Schönenberger, J.-P. Salvetat, J.-M. Bonard, and L. Forro. *Contacting carbon nanotubes selectively with low-ohmic contacts for four-probe electric measurements.* Appl. Phys. Lett., **73**(2) 274–276, 1998.

[96] D. P. Wang, B. R. Perkins, A. J. Yin, A. Zaslavsky, J. M. Xu, R. Beresford, and G. L. Snider. *Carbon nanotube gated lateral resonant tunneling field-effect transistors.* Appl. Phys. Lett., **87** 152102, 2005.

[97] M. Bockrath, D. H. Cobden, P. L. McEuen, N. G. Chopra, A. Zettl, A. Thess, and R. E. Smalley. *Single-Electron Transport in Ropes of Carbon Nanotubes.* Science, **275** 1922–1925, 1997.

[98] K. D. Ausman, R. Piner, O. Lourie, R. S. Ruoff, and M. Korobov. *Organic Solvent Dispersions of Single-Walled Carbon Nanotubes: Toward Solutions of Pristine Nanotubes.* J. Phys. Chem. B, **104**(38) 8911–8915, 2000.

REFERENCES

[99] J. Chen, M. A. Hamon, H. Hu, Y. Chen, A. M. Rao, P. C. Eklund, and R. C. Haddon. *Solution Properties of Single-Walled Carbon Nanotubes.* Science, **282** 95–98, 1998.

[100] D. W. Schaefer, J. Zhao, J. M. Brown, D. P. Anderson, and D. W. Tomlin. *Morphology of dispersed carbon single-walled nanotubes.* Chem. Phys. Lett., **375**(3-4) 369–375, 2003.

[101] B. Vigolo, A. Penicaud, C. Coulon, C. Sauder, R. Pailler, C. Journet, P. Bernier, and P. Poulin. *Macroscopic Fibers and Ribbons of Oriented Carbon Nanotubes.* Science, **290** 1331–1334, 2000.

[102] G. S. Duesberg, M. Burghard, J. Muster, and G. Philipp. *Separation of carbon nanotubes by size exclusion chromatography.* Chem. Commun., **1998**(3) 435–436, 1998.

[103] K. D. Ausman, M. J. O'Connell, P. Boul, L. M. Ericson, M. J. Casavant, D. A. Walters, C. Huffman, R. Saini, Y. Wang, E. Haroz, E. W. Billups, and R. E. Smalley. *Roping and Wrapping Carbon Nanotubes.* In *XVth International Winterschool on Electronic Properties of Novel Materials*, 2000.

[104] M. F. Islam, E. Rojas, D. M. Bergey, A. T. Johnson, and A. G. Yodh. *High Weight Fraction Surfactant Solubilization of Single-Wall Carbon Nanotubes in Water.* Nano Lett., **3**(2) 269–273, 2003.

[105] J. M. Bonard, T. Stora, J. P. Salvetat, F. Maier, T. Stoeckli, C. Duschl, L. Forro, W. A. de Heer, and A. Chatelain. *Purification and size-selection of carbon nanotubes.* Adv. Mat., **9**(10) 827–831, 1997.

[106] J. Ebbecke, C. J. Strobl, and A. Wixforth. *Acoustoelectric current transport through single-walled carbon nanotubes.* Phys. Rev. B, **70** 233401, 2004.

[107] J. Li, Q. Zhang, Y. Yan, S. Li, and L. Chen. *Fabrication of Carbon Nanotube Field-Effect Transistors by Fluidic Alignment Technique.* IEEE Transactions on Nanotechnology, **6**(4) 481–484, 2007.

[108] S. Niyogi, M. A. Hamon, D. E. Perea, C. B. Kang, B. Zhao, S. K. Pal, A. E. Wyant, M. E. Itkis, and R. C. Haddon. *Ultrasonic Dispersions of Single-Walled Carbon Nanotubes.* J. Phys. Chem. B, **107**(34) 8799–8804, 2003.

[109] D. E. Johnston, M. F. Islam, A. G. Yodh, and A. T. Johnson. *Electronic devices based on purified carbon nanotubes grown by high-pressure decomposition of carbon monoxide.* Nature Materials, **4**(8) 589–592, 2005.

[110] S. Iijima, C. Brabec, A. Maiti, and J. Bernholc. *Structural flexibility of carbon nanotubes.* J. Chem. Phys., **104**(5) 2089–2092, 1996.

[111] S. Hasegawa, K. Arakawa, H. Oooka, and H. Nakashima. *Macrostep and mound formation during AlGaAs growth on vicinal GaAs(110) studied by scanning tunneling microscopy.* Appl. Surf. Sci., **162** 430–434, 2000.

[112] M. Yamamoto, M. Higashiwaki, S. Shimomura, N. Sano, and S. Hiyamizu. *Surface Corrugation of GaAs Layers Grown on (775) B-Oriented GaAs Substrates by Molecular Beam Epitaxy.* Jpn. J. Appl. Phys., **36** 6285–6289, 1997.

[113] L. F. Chibotaru, S. A. Bovin, and A. Ceulemans. *Bend-induced insulating gap in carbon nanotubes.* Phys. Rev. B, **66**(16) 161401, 2002.

[114] H. W. C. Postma, A. Sellmeijer, and C. Dekker. *Manipulation and Imaging of Individual Single-Walled Carbon Nanotubes with an Atomic Force Microscope.* Adv. Mat., **12**(17) 1299–1302, 2000.

[115] M. J. Biercuk, N. Mason, J. M. Chow, and C. M. Marcus. *Locally Addressable Tunnel Barriers within a Carbon Nanotube.* Nano Lett., **4**(12) 2499–2502, 2004.

[116] L. C. Venema. *Electronic structure of carbon nanotubes.* PhD thesis, Delft University of Technology, 2000.

[117] J. Y. Park, Y. Yaish, M. Brink, S. Rosenblatt, and P. L. McEuen. *Electrical cutting and nicking of carbon nanotubes using an atomic force microscope.* Appl. Phys. Lett., **80** 4446–4448, 2002.

[118] S. Weber. *JNanoTube Modeler*, 2000. URL http://jcrystal.com/steffen weber/JAVA/jnano/jnano.html.

[119] W. Orellana, R. H. Miwa, and A. Fazzio. *Stability and electronic properties of carbon nanotubes adsorbed on Si(001).* Surf. Sci., **566** 728–732, 2004.

[120] Y. H. Kim, M. J. Heben, and S. B. Zhang. *Nanotube Wires on Commensurate InAs Surfaces: Binding Energies, Band Alignments, and Bipolar Doping by the Surfaces.* Phys. Rev. Lett., **92**(17) 176102, 2004.

REFERENCES

[121] L. A. Wade, I. R. Shapiro, Z. Ma, S. R. Quake, and C. P. Collier. *Correlating AFM Probe Morphology to Image Resolution for Single-Wall Carbon Nanotube Tips.* Nano Lett., **4**(4) 725–732, 2004.

[122] A. R. Clawson. *Guide to references on III–V semiconductor chemical etching.* Materials Science & Engineering, **31**(1-6) 1–438, 2001.

[123] M. Terrones, F. Banhart, N. Grobert, J. C. Charlier, H. Terrones, and P. M. Ajayan. *Molecular Junctions by Joining Single-Walled Carbon Nanotubes.* Phys. Rev. Lett., **89**(7) 075505, 2002.

[124] B. W. Smith and D. E. Luzzi. *Electron irradiation effects in single wall carbon nanotubes.* J. Appl. Phys., **90** 3509–3515, 2001.

[125] S. Suzuki and Y. Kobayashi. *Conductivity Decrease in Carbon Nanotubes Caused by Low-Acceleration-Voltage Electron Irradiation.* Jpn. J. Appl. Phys., **44**(49) L1498–L1501, 2005.

[126] A. Vijayaraghavan, K. Kanzaki, S. Suzuki, Y. Kobayashi, H. Inokawa, Y. Ono, S. Kar, and P. M. Ajayan. *Metal-Semiconductor Transition in Single-Walled Carbon Nanotubes Induced by Low-Energy Electron Irradiation.* Nano Lett., **5**(8) 1575–1579, 2005.

[127] S. Suzuki, K. Kanzaki, Y. Homma, and S. Fukuba. *Low-Acceleration-Voltage Electron Irradiation Damage in Single-Walled Carbon Nanotubes.* Jpn. J. Appl. Phys., **43** 1118–1120, 2004.

[128] J. Appenzeller, R. Martel, P. Avouris, H. Stahl, and B. Lengeler. *Optimized contact configuration for the study of transport phenomena in ropes of single-wall carbon nanotubes.* Appl. Phys. Lett., **78** 3313–3315, 2001.

[129] Y. F. Chen and M. S. Fuhrer. *Electric Field-Dependent Charge-Carrier Velocity in Semiconducting Carbon Nanotubes.* Phys. Rev. Lett., **95** 236803, 2005.

[130] P. Jarillo-Herrero, S. Sapmaz, C. Dekker, L. P. Kouwenhoven, and H. S. J. van der Zant. *Electron-hole symmetry in a semiconducting carbon nanotube quantum dot.* Nature, **429** 389–392, 2004.

[131] D. Mann, A. Javey, J. Kong, Q. Wang, and H. Dai. *Ballistic Transport in Metallic Nanotubes with Reliable Pd Ohmic Contacts.* Nano Lett., **3**(11) 1541–1544, 2003.

[132] A. Javey, J. Guo, Q. Wang, M. Lundstrom, and H. Dai. *Ballistic carbon nanotube field-effect transistors.* Nature, **424** 654–657, 2003.

[133] Y. Nosho, Y. Ohno, S. Kishimoto, and T. Mizutani. *n-type carbon nanotube field-effect transistors fabricated by using Ca contact electrodes.* Appl. Phys. Lett., **86**(7) 073105, 2005.

[134] B. Babić and C. Schönenberger. *Observation of Fano-resonances in single-wall carbon nanotubes.* Phys. Rev. B, **70** 195408, 2004.

[135] H. Maki, T. Sato, and K. Ishibashi. *Electrical properties of carbon nanotube sheets with Ni and Mg electrode metals.* Trans. Mater. Res. Soc. Jpn, **31**(1) 215–217, 2006.

[136] M. Bockrath, J. Hone, A. Zettl, P. L. McEuen, A. G. Rinzler, and R. E. Smalley. *Chemical doping of individual semiconducting carbon-nanotube ropes.* Phys. Rev. B, **61**(16) R10606, 2000.

[137] D. A. Wharam, M. Pepper, H. Ahmed, J. E. F. Frost, D. G. Hasko, D. C. Peacock, D. A. Ritchie, and G. A. C. Jones. *Addition of the one-dimensional quantised ballistic resistance.* J. Phys. C: Solid State Physics, **21**(24) L887–L891, 1988.

[138] B. E. Luquette and B. M. Nichols. *Fabrication and Characterization of Schottky Diodes using Single Wall Carbon Nanotubes.* Technical report, Army Research Laboratory (Adelphi, MD), 2008.

[139] P. L. McEuen, M. Bockrath, D. H. Cobden, Y. G. Yoon, and S. G. Louie. *Disorder, Pseudospins, and Backscattering in Carbon Nanotubes.* Phys. Rev. Lett., **83**(24) 5098–5101, 1999.

[140] V. Skákalová, A. B. Kaiser, Y. S. Woo, and S. Roth. *Electronic transport in carbon nanotubes: From individual nanotubes to thin and thick networks.* Phys. Rev. B, **74**(8) 85403, 2006.

[141] T. Tsukamoto, S. Moriyama, D. Tsuya, M. Suzuki, T. Yamaguchi, Y. Aoyagi, and K. Ishibashi. *Carbon nanotube quantum dots fabricated on a GaAs/AlGaAs two-dimensional electron gas substrate.* J. Appl. Phys., **98**(7) 076106, 2005.

REFERENCES

[142] A. A. M. Staring, H. van Houten, C. W. J. Beenakker, and C. T. Foxon. *Coulomb-blockade oscillations in disordered quantum wires.* Phys. Rev. B, **45**(16) 9222–9236, 1992.

[143] K. Tsukagoshi, K. Nakazato, H. Ahmed, and K. Gamo. *Electron pump in multiple-tunnel junctions.* Phys. Rev. B, **56**(7) 3972–3975, 1997.

[144] A. A. Maarouf, C. L. Kane, and E. J. Mele. *Electronic structure of carbon nanotube ropes.* Phys. Rev. B, **61**(16) 11156–11165, 2000.

[145] S. Okada and A. Oshiyama. *Electronic Structure of Semiconducting Nanotubes Adsorbed on Metal Surfaces.* Phys. Rev. Lett., **95** 206804, 2005.

[146] E. S. Snow, J. P. Novak, P. M. Campbell, and D. Park. *Random networks of carbon nanotubes as an electronic material.* Appl. Phys. Lett., **82** 2145–2147, 2003.

[147] T. Fukuda, F. Arai, and L. Dong. *Assembly of Nanodevices With Carbon Nanotubes Through Nanorobotic Manipulations.* Proc. IEEE, **91**(11) 1803–1818, 2003.

[148] D. C. Cox, R. D. Forrest, P. R. Smith, V. Stolojan, and S. R. P. Silva. *Study of the current stressing in nanomanipulated three-dimensional carbon nanotube structures.* Appl. Phys. Lett., **87** 033102, 2005.

[149] M. Yu, M. J. Dyer, G. D. Skidmore, H. W. Rohrs, X. Lu, K. D. Ausman, J. R. Von Ehr, and R. S. Ruoff. *Three-dimensional manipulation of carbon nanotubes under a scanning electron microscope.* Nanotechnology, **10** 244–252, 1999.

[150] P. Beecher, P. Servati, A. Rozhin, A. Colli, V. Scardaci, S. Pisana, T. Hasan, A. J. Flewitt, et al. *Ink-jet printing of carbon nanotube thin film transistors.* J. Appl. Phys., **102** 043710, 2007.

[151] S. Hofmann, C. Ducati, J. Robertson, and B. Kleinsorge. *Low-temperature growth of carbon nanotubes by plasma-enhanced chemical vapor deposition.* Appl. Phys. Lett., **83** 135–137, 2003.

[152] S. Hofmann, B. Kleinsorge, C. Ducati, A. C. Ferrari, and J. Robertson. *Low-temperature plasma enhanced chemical vapour deposition of carbon nanotubes.* Diamond & Related Materials, **13**(4-8) 1171–1176, 2004.

[153] M. Cantoro, S. Hofmann, S. Pisana, V. Scardaci, A. Parvez, C. Ducati, A. C. Ferrari, A. M. Blackburn, K. Wang, and J. Robertson. *Catalytic Chemical Vapor Deposition of Single-Wall Carbon Nanotubes at Low Temperatures*. Nano Lett., **6**(6) 1107–1112, 2006.

[154] I. A. Prudaev and S. S. Khludkov. *Diffusion and solubility of electrically active iron atoms in gallium arsenide*. Russian Physics Journal, **51**(11) 1157–1160, 2008.

[155] J. R. Hauptmann, A. Jensen, J. Nygård, and P. E. Lindelof. *Single Wall Carbon Nanotubes in Epitaxial Grown Semiconductor Heterostructures*. Electronic Properties of Novel Nanostructures, **786** 495–498, 2005.

[156] A. Jensen, J. R. Hauptmann, J. Nygard, J. Sadowski, and P. E. Lindelof. *Hybrid Devices from Single Wall Carbon Nanotubes Epitaxially Grown into a Semiconductor Heterostructure*. Nano Lett., **4**(2) 349–352, 2004.

[157] J. Nygård, A. Jensen, J. R. Hauptmann, J. Sadowski, and P. E. Lindelof. *Integration of Carbon Nanotubes with Semiconductor Technology by Epitaxial Encapsulation*. In *AIP Conference Proceedings*, volume 723, page 552, 2004.

[158] S. Stobbe, P. E. Lindelof, and J. Nygård. *Integration of carbon nanotubes with semiconductor technology: fabrication of hybrid devices by III-V molecular beam epitaxy*. Semicond. Sci. Technol., **21**(11) 10–16, 2006.

[159] J. C. Meyer, J. Cech, B. Hornbostel, and S. Roth. *Progress in single-walled carbon nanotube based nanoelectromechanical systems*. physica status solidi (b), **243**(13) 3500–3504, 2006.

[160] C. W. Liang and S. Roth. *Electrical and Optical Transport of GaAs/Carbon Nanotube Heterojunctions*. Nano Lett., **8**(7) 1809–1812, 2008.

[161] K. Ishibashi. *Personal communication*, 2006.

[162] S. Gustavsson, I. Shorubalko, R. Leturcq, T. Ihn, K. Ensslin, and S. Schön. *Detecting terahertz current fluctuations in a quantum point contact using a nanowire quantum dot*. Phys. Rev. B, **78**(3) 035324, 2008.

[163] R. Engel-Herbert, Y. Takagaki, and T. Hesjedal. *Growth of carbon nanotubes on GaAs*. Materials Letters, **61**(23-24) 4631–4634, 2007.

REFERENCES

[164] S. Dittmer, O. A. Nerushev, and E. E. B. Campbell. *Low ambient temperature CVD growth of carbon nanotubes.* Appl. Phys. A, **84**(3) 243–246, 2006.

[165] L. B. Ruppalt, P. M. Albrecht, and J. W. Lyding. *Deposition and STM Investigation of Single-Walled Carbon Nanotubes on GaAs (110).* In *4th IEEE Conference on Nanotechnology, 2004*, pp. 13–15, 2004.

[166] L. B. Ruppalt, P. M. Albrecht, and J. W. Lyding. *Atomic resolution scanning tunneling microscope study of single-walled carbon nanotubes on GaAs (110).* J. Vac. Sci. Technol. B, **22** 2005–2007, 2004.

[167] L. B. Ruppalt, P. M. Albrecht, and J. W. Lyding. *UHV-STM study of single-walled carbon nanotubes applied to the GaAs(110) and InAs(110) surfaces.* J. Phys. IV France, **132** 31–34, 2006.

[168] L. B. Ruppalt and J. W. Lyding. *Charge transfer between semiconducting carbon nanotubes and their doped GaAs(110) and InAs(110) substrates detected by scanning tunnelling spectroscopy.* Nanotechnology, **18**(21) 215202–215202, 2007.

[169] D. Gonzalez, A. G. Nasibulin, S. D. Shandakov, H. Jiang, P. Queipo, A. S. Anisimov, T. Tsuneta, and E. I. Kauppinen. *Spontaneous Charging of Single-Walled Carbon Nanotubes: A Novel Strategy for the Selective Substrate Deposition of Individual Tubes at Ambient Temperature.* Chem. Mater., **18**(21) 5052–5057, 2006.

[170] X. Liu, C. Lee, S. Han, C. Li, and C. Zhou. *Molecular Nanoelectronics*, Carbon Nanotubes: Synthesis, Devices, and Integrated Systems, pp. 1–20. ASP, 2003.

[171] A. Bachtold, M. S. Fuhrer, S. Plyasunov, M. Forero, E. H. Anderson, A. Zettl, and P. L. McEuen. *Scanned Probe Microscopy of Electronic Transport in Carbon Nanotubes.* Phys. Rev. Lett., **84**(26) 6082–6085, 2000.

[172] Y. Miyato, K. Kobayashi, K. Matsushige, and H. Yamada. *Surface potential investigation on single wall carbon nanotubes by Kelvin probe force microscopy and atomic force microscope potentiometry.* Nanotechnology, **18** 084008, 2007.

[173] M. Bockrath, W. Liang, D. Bozovic, J. H. Hafner, C. M. Lieber, M. Tinkham, and H. Park. *Resonant Electron Scattering by Defects in Single-Walled Carbon Nanotubes.* Science, **291** 283–285, 2001.

[174] R. M. Westervelt. *Graphene Nanoelectronics.* Science, **320** 324–325, 2008.

[175] Z. Chen, Y. M. Lin, M. J. Rooks, and P. Avouris. *Graphene nano-ribbon electronics.* Physica E, **40**(2) 228–232, 2007.

[176] V. H. O. Barone and G. E. Scuseria. *Electronic Structure and Stability of Semiconducting Graphene Nanoribbons.* Nano Lett., **6**(12) 2748–2754, 2006.

[177] Y. W. Son, M. L. Cohen, and S. G. Louie. *Energy Gaps in Graphene Nanoribbons.* Phys. Rev. Lett., **97**(21) 216803, 2006.

[178] G. Liang, N. Neophytou, D. E. Nikonov, and M. S. Lundstrom. *Performance Projections for Ballistic Graphene Nanoribbon Field-Effect Transistors.* IEEE Transactions on Electron Devices, **54**(4) 677–682, 2007.

[179] Y. Ouyang, Y. Yoon, J. K. Fodor, and J. Guo. *Comparison of performance limits for carbon nanoribbon and carbon nanotube transistors.* Appl. Phys. Lett., **89** 203107, 2006.

[180] L. Brey and H. A. Fertig. *Electronic states of graphene nanoribbons studied with the Dirac equation.* Phys. Rev. B, **73**(23) 235411, 2006.

[181] L. A. Ponomarenko, F. Schedin, M. I. Katsnelson, R. Yang, E. H. Hill, K. S. Novoselov, and A. K. Geim. *Chaotic Dirac Billiard in Graphene Quantum Dots.* Science, **320** 356–358, 2008.

[182] G. Binnig, H. Rohrer, C. Gerber, and E. Weibel. *Surface Studies by Scanning Tunneling Microscopy.* Phys. Rev. Lett., **49**(1) 57–61, 1982.

[183] G. Binnig, C. F. Quate, and C. Gerber. *Atomic Force Microscope.* Phys. Rev. Lett., **56**(9) 930–933, 1986.

[184] Y. Martin, C. C. Williams, and H. K. Wickramasinghe. *Atomic force microscope–force mapping and profiling on a sub 100-Å scale.* J. Appl. Phys., **61**(10) 4723–4729, 1987.

[185] A. A. Baski. *Fabrication of Nanoscale Structures using STM and AFM.* Advanced Semiconductor and Organic Nano-Techniques, **3** 189–224, 2002.

[186] M. Tortonese, R. C. Barrett, and C. F. Quate. *Atomic resolution with an atomic force microscope using piezoresistive detection.* Appl. Phys. Lett., **62**(8) 834–836, 1993.

REFERENCES

[187] C. Berger, Z. Song, X. Li, X. Wu, N. Brown, C. Naud, D. Mayou, T. Li, J. Hass, et al. *Electronic Confinement and Coherence in Patterned Epitaxial Graphene.* Science, **312** 1191–1196, 2006.

[188] R. Magno and B. R. Bennett. *Nanostructure patterns written in III–V semiconductors by an atomic force microscope.* Appl. Phys. Lett., **70** 1855–1857, 1997.

[189] H. W. Schumacher, U. F. Keyser, U. Zeitler, R. J. Haug, and K. Eberl. *Nanomachining of mesoscopic electronic devices using an atomic force microscope.* Appl. Phys. Lett., **75** 1107–1109, 1999.

[190] J. Regul, U. F. Keyser, M. Paesler, F. Hohls, U. Zeitler, R. J. Haug, A. Malave, E. Oesterschulze, D. Reuter, and A. D. Wieck. *Fabrication of quantum point contacts by engraving GaAs/AlGaAs heterostructures with a diamond tip.* Appl. Phys. Lett., **81** 2023–2025, 2002.

[191] M. Wendel, S. Kühn, H. Lorenz, J. P. Kotthaus, and M. Holland. *Nanolithography with an atomic force microscope for integrated fabrication of quantum electronic devices.* Appl. Phys. Lett., **65** 1775–1777, 1994.

[192] M. J. Murphy. *Low-Dimensional Electron Systems Fabricated with an Atomic Force Microscope.* PhD thesis, University of Cambridge, 2007.

[193] H. C. Day and D. R. Allee. *Selective area oxidation of silicon with a scanning force microscope.* Appl. Phys. Lett., **62** 2691–2693, 1993.

[194] J. A. Dagata, J. Schneir, H. H. Harary, C. J. Evans, M. T. Postek, and J. Bennett. *Modification of hydrogen-passivated silicon by a scanning tunneling microscope operating in air.* Appl. Phys. Lett., **56** 2001, 1990.

[195] B. Irmer, M. Kehrle, H. Lorenz, and J. P. Kotthaus. *Fabrication of Ti/TiO tunneling barriers by tapping mode atomic force microscopy induced local oxidation.* Appl. Phys. Lett., **71** 1733–1735, 1997.

[196] A. E. Gordon, R. T. Fayfield, D. D. Litfin, and T. K. Higman. *Mechanisms of surface anodization produced by scanning probe microscopes.* J. Vac. Sci. Technol. B, **13** 2805–2808, 1995.

[197] A. Majumdar, P. I. Oden, J. P. Carrejo, L. A. Nagahara, J. J. Graham, and J. Alexander. *Nanometer-scale lithography using the atomic force microscope.* Appl. Phys. Lett., **61** 2293–2295, 1992.

[198] T. Abe, K. Hane, and S. Okuma. *Nanometer-scale pit formation by scanning tunneling microscopy on graphite surface and tip current measurements.* J. Appl. Phys., **75(2)** 1228–1230, 1994.

[199] T. R. Albrecht, M. M. Dovek, M. D. Kirk, C. A. Lang, C. F. Quate, and D. P. E. Smith. *Nanometer-scale hole formation on graphite using a scanning tunneling microscope.* Appl. Phys. Lett., **55** 1727–1729, 1989.

[200] H. Hiura. *Tailoring graphite layers by scanning tunneling microscopy.* Appl. Surf. Sci., **222**(1-4) 374–381, 2004.

[201] S. Kondo, M. Lutwyche, and Y. Wada. *Nanofabrication of layered materials with the scanning tunneling microscope.* Appl. Surf. Sci., **75**(1-4) 39–44, 1994.

[202] R. L. McCarley, S. A. Hendricks, and A. J. Bard. *Controlled Nanofabrication of Highly Oriented Pyrolytic Graphite with the Scanning Tunneling Microscope.* J. Phys. Chem., **96**(25) 10089–10092, 1992.

[203] W. Mizutani, J. Inukai, and M. Ono. *Making a Monolayer Hole in a Graphite Surface by Means of a Scanning Tunneling Microscope.* Jpn. J. Appl. Phys., **29**(5) L815–L817, 1990.

[204] R. M. Penner, J. M. Heben, and N. S. Lewis. *Mechanistic investigations of nanometer-scale lithography at liquid-covered graphite surfaces.* Appl. Phys. Lett., **58**(13) 1389–1391, 1991.

[205] S. L. Tang, A. J. McGhie, and M. E. Lewittes. *Nanostructure fabrication with the scanning tunneling microscope by tunneling medium manipulation.* Appl. Phys. Lett., **60** 1821–1923, 1992.

[206] R. R. Nair, P. Blake, A. N. Grigorenko, K. S. Novoselov, T. J. Booth, T. Stauber, N. M. R. Peres, and A. K. Geim. *Fine Structure Constant Defines Visual Transparency of Graphene.* Science, **320** 1308, 2008.

[207] P. Blake, E. W. Hill, A. H. C. Neto, K. S. Novoselov, D. Jiang, R. Yang, T. J. Booth, and A. K. Geim. *Making graphene visible.* Appl. Phys. Lett., **91** 063124, 2007.

REFERENCES

[208] A. C. Ferrari, J. C. Meyer, V. Scardaci, C. Casiraghi, M. Lazzeri, F. Mauri, S. Piscanec, D. Jiang, K. S. Novoselov, S. Roth, and A. K. Geim. *Raman Spectrum of Graphene and Graphene layers.* Phys. Rev. Lett., **97** 187401, 2006.

[209] H. Hibino, H. Kageshima, F. Maeda, M. Nagase, Y. Kobayashi, and H. Yamaguchi. *Microscopic thickness determination of thin graphite films formed on SiC from quantized oscillation in reflectivity of low-energy electrons.* Phys. Rev. B, **77** 75413, 2008.

[210] H. Hibino, H. Kageshima, F. Maeda, M. Nagase, Y. Kobayashi, and H. Yamaguchi. *Thickness Determination of Graphene Layers Formed on SiC Using Low-Energy Electron Microscopy.* e-J. Surf. Sci. and Nanotech., **6** 107–110, 2008.

[211] A. Reina, X. Jia, J. Ho, D. Nezich, H. Son, V. Bulovic, M. S. Dresselhaus, and J. Kong. *Large Area, Few-Layer Graphene Films on Arbitrary Substrates by Chemical Vapor Deposition.* Nano Lett., **9**(1) 30–35, 2009.

[212] K. S. Novoselov, D. Jiang, F. Schedin, T. J. Booth, V. V. Khotkevich, S. V. Morozov, and A. K. Geim. *Two-dimensional atomic crystals.* Proc. Natl. Acad. Sci., **102**(30) 10451–10453, 2005.

[213] Graphene Industries, 2007. URL `http://grapheneindustries.com`.

[214] K. V. Emtsev, A. Bostwick, K. Horn, J. Jobst, G. L. Kellogg, L. Ley, J. L. McChesney, T. Ohta, S. A. Reshanov, et al. *Towards wafer-size graphene layers by atmospheric pressure graphitization of silicon carbide.* Nature Materials, **8**(3) 203–207, 2009.

[215] C. Oshima and A. Nagashima. *Ultra-thin epitaxial films of graphite and hexagonal boron nitride on solid surfaces.* J. Phys.: Cond. Matt., **9** 1–20, 1997.

[216] U. Starke and C. Riedl. *Epitaxial graphene on SiC(0001) and SiC(000$\bar{1}$): from surface reconstructions to carbon electronics.* J. Phys.: Cond. Matt., **21**(13) 134016, 2009.

[217] X. Liang, Z. Fu, and S. Y. Chou. *Graphene Transistors Fabricated via Transfer-Printing In Device Active-Areas on Large Wafer.* Nano Lett., **7**(12) 3840–3844, 2007.

[218] D. Tyndall. *Exfoliation and AFM-Lithography of Graphene.* Master's thesis, University of Cambridge, 2009.

[219] NGS Naturgraphit GmbH. URL `http://www.graphit.de`.

[220] M. Luna, J. Colchero, and A. M. Baro. *Study of Water Droplets and Films on Graphite by Noncontact Scanning Force Microscopy.* J. Phys. Chem. B, **103**(44) 9576–9581, 1999.

[221] R. H. Fowler and L. Nordheim. *Electron Emission in Intense Electric Fields.* Proc. R. Soc. Lond. A, **119**(781) 173–181, 1928.

[222] X. Lu, M. Yu, H. Huang, and R. S. Ruoff. *Tailoring graphite with the goal of achieving single sheets.* Nanotechnology, **10** 269–272, 1999.

[223] H. V. Roy, C. Kallinger, B. Marsen, and K. Sattler. *Manipulation of graphitic sheets using a tunneling microscope.* J. Appl. Phys., **83** 4695–4699, 1998.

[224] H. Hiura, T. W. Ebbesen, J. Fujita, K. Tanigaki, and T. Takada. *Role of sp^3 defect structures in graphite and carbon nanotubes.* Nature, **367** 148–151, 1994.

[225] A. M. B. do Rego and L. F. V. Ferreira. *Photonic and electronic spectroscopies for the characterization of organic surfaces and organic molecules adsorbed on surfaces.* Handbook of Surfaces and Interfaces of Materials, **2** 275–313, 2001.

[226] W. Kim, A. Javey, O. Vermesh, Q. Wang, Y. Li, and H. Dai. *Hysteresis Caused by Water Molecules in Carbon Nanotube Field-Effect Transistors.* Nano Lett., **3**(2) 193–198, 2003.

[227] T. Bakos, S. N. Rashkeev, and S. T. Pantelides. *H_2O and O_2 molecules in amorphous SiO_2: Defect formation and annihilation mechanisms.* Phys. Rev. B, **69**(19) 195206, 2004.

[228] R. H. Doremus. *Diffusion of water in silica glass.* J. Mater. Res., **10**(9) 2379–2389, 1995.

[229] E. McCann, K. Kechedzhi, V. Falko, H. Suzuura, T. Ando, and B. L. Altshuler. *Weak-Localization Magnetoresistance and Valley Symmetry in Graphene.* Phys. Rev. Lett., **97**(14) 146805, 2006.

[230] P. A. Lee and A. D. Stone. *Universal Conductance Fluctuations in Metals.* Phys. Rev. Lett., **55**(15) 1622–1625, 1985.

REFERENCES

[231] D. W. Horsell, A. K. Savchenko, F. V. Tikhonenko, K. Kechedzhi, I. V. Lerner, and V. I. Falko. *Mesoscopic conductance fluctuations in graphene.* Solid State Comm., **149** 1041–1045, 2009.

[232] S. Kondo, S. Heike, M. Lutwyche, and Y. Wada. *Surface modification mechanism of materials with scanning tunneling microscope.* J. Appl. Phys., **78**(1) 155–160, 1995.

[233] S. Myhra. *Tip-induced oxidative nano-machining of conducting diamond-like carbon (DLC).* Appl. Phys. A, **80**(5) 1097–1104, 2005.

[234] C. Wang, C. Bai, X. Li, G. Shang, I. Lee, X. Wang, X. Qiu, and F. Tian. *Evidence of diffusion characteristics of field emission electrons in nanostructuring process on graphite surface.* Appl. Phys. Lett., **69** 348–350, 1996.

[235] J. Freund, J. Halbritter, and J. K. H. Horber. *How Dry are Dried Samples? Water Adsorption Measured by STM.* Microscopy Research and Technique, **44**(5) 327–338, 1999.

[236] A. Gil, J. Colchero, M. Luna, J. Gomez-Herrero, and A. M. Baro. *Adsorption of Water on Solid Surfaces Studied by Scanning Force Microscopy.* Langmuir, **16**(11) 5086–5092, 2000.

[237] J. Zhao, C. T. Chan, and J. G. Che. *Effects of an electric field on a water bilayer on Ag(111).* Phys. Rev. B, **75**(8) 85435, 2007.

[238] M. F. Toney, J. N. Howard, J. Richer, G. L. Borges, J. G. Gordon, O. R. Melroy, D. G. Wiesler, D. Yee, and L. B. Sorensen. *Distribution of water molecules at Ag(111)/electrolyte interface as studied with surface X-ray scattering.* Surf. Sci., **335** 326–332, 1995.

[239] R. H. Bernhardt, G. C. McGonigal, R. Schneider, and D. J. Thomson. *Mechanisms for the deposition of nanometer-sized structures from organic fluids using the scanning tunneling microscope.* J. Vac. Sci. Technol. A, **8** 667–671, 1990.

[240] J. P. Rabe, S. Buchholz, and A. M. Ritcey. *Reactive graphite etch and the structure of an adsorbed organic monolayera scanning tunneling microscopy study.* J. Vac. Sci. Technol. A, **8** 679–683, 1990.

[241] C. Wang, X. Li, G. Shang, X. Qiu, and C. Bai. *Threshold behavior of nanometer scale fabrication process using scanning tunneling microscopy.* J. Appl. Phys., **81** 1227–1230, 1997.

[242] C. Gómez-Navarro, R. T. Weitz, A. M. Bittner, M. Scolari, A. Mews, M. Burghard, and K. Kern. *Electronic Transport Properties of Individual Chemically Reduced Graphene Oxide Sheets.* Nano Lett., **7**(11) 3499–3503, 2007.

[243] S. Stankovich, R. D. Piner, S. B. T. Nguyen, and R. S. Ruoff. *Synthesis and exfoliation of isocyanate-treated graphene oxide nanoplatelets.* Carbon, **44**(15) 3342–3347, 2006.

[244] K. Mkhoyan, A. W. Contryman, J. Silcox, D. A. Stewart, G. Eda, C. Mattevi, S. Miller, and M. Chhowalla. *Atomic and Electronic Structure of Graphene-Oxide.* Nano Lett., **9**(3) 1058–1063, 2009.

[245] C. A. Goss, J. C. Brumfield, E. A. Irene, and R. W. Murray. *Imaging the Incipient Electrochemical Oxidation of Highly Oriented Pyrolytic Graphite.* Anal. Chem., **65**(10) 1378–1389, 1993.

[246] B. Zhang and E. Wang. *Effects of anodic oxidation on the surface structure of highly oriented pyrolytic graphite revealed by in situ electrochemical scanning tunnelling microscopy in H_2SO_4 solution.* Electrochimica Acta, **40**(16) 2627–2633, 1995.

[247] D. C. Alsmeyer and R. L. McCreery. *In Situ Raman Monitoring of Electrochemical Graphite Intercalation and Lattice Damage in Mild Aqueous Acids.* Anal. Chem., **64**(14) 1528–1533, 1992.

[248] A. Allouche and Y. Ferro. *Dissociative adsorption of atmospheric molecules at vacancies on the graphite (0001) surface of samples exposed to plasma.* Journal of Nuclear Materials, **363** 117–121, 2007.

[249] M. Y. Han, B. Özyilmaz, Y. Zhang, and P. Kim. *Energy Band-Gap Engineering of Graphene Nanoribbons.* Phys. Rev. Lett., **98**(20) 206805, 2007.

[250] P. Barthold, T. Lüdtke, and R. J. Haug. *Nanomachining of multilayer graphene using an atomic force microscope.* arXiv, **cond-mat** 0807.4033, 2008.

REFERENCES

[251] A. J. M. Giesbers, U. Zeitler, S. Neubeck, F. Freitag, K. S. Novoselov, and J. C. Maan. *Nanolithography and manipulation of graphene using an atomic force microscope.* Solid State Comm., **147**(9-10) 366–369, 2008.

[252] L. Weng, L. Zhang, Y. P. Chen, and L. P. Rokhinson. *Atomic force microscope local oxidation nanolithography of graphene.* Appl. Phys. Lett., **93** 093107, 2008.

[253] L. Tapasztó, G. Dobrik, P. Lambin, and L. P. Biró. *Tailoring the atomic structure of graphene nanoribbons by scanning tunnelling microscope lithography.* Nature Nanotechnology, **3**(7) 397–401, 2008.

[254] L. Ci, Z. Xu, L. Wang, W. Gao, F. Ding, K. F. Kelly, B. I. Yakobson, and P. M. Ajayan. *Controlled Nanocutting of Graphene.* Nano Res., **1**(2) 116–122, 2008.

[255] X. Li, X. Wang, L. Zhang, S. Lee, and H. Dai. *Chemically Derived, Ultrasmooth Graphene Nanoribbon Semiconductors.* Science, **319** 1229–1232, 2008.

[256] J. H. Chen, C. Jang, S. Xiao, M. Ishigami, and M. S. Fuhrer. *Intrinsic and extrinsic performance limits of graphene devices on SiO_2.* Nature Nanotechnology, **3**(4) 206–209, 2008.

[257] J. Sabio, C. Seoanez, S. Fratini, F. Guinea, A. H. C. Neto, and F. Sols. *Electrostatic interactions between graphene layers and their environment.* Phys. Rev. B, **77**(19) 195409, 2008.

[258] M. R. Connolly, K. L. Chiou, C. G. Smith, D. Anderson, G. A. C. Jones, A. Lombardo, A. Fasoli, and A. C. Ferrari. *Scanning gate microscopy of current-annealed single layer graphene.* arXiv, **cond-mat** 0911.3832, 2009.

[259] J. G. Park, C. Zhang, R. Liang, and B. Wang. *Nano-machining of highly oriented pyrolytic graphite using conductive atomic force microscope tips and carbon nanotubes.* Nanotechnology, **18**(40) 405306, 2007.

[260] Y. Jiang and W. Guo. *Convex and concave nanodots and lines induced on HOPG surfaces by AFM voltages in ambient air.* Nanotechnology, **19** 345302, 2008.

[261] A. Knoll, P. Bächtold, J. Bonan, G. Cherubini, M. Despont, U. Drechsler, U. Dürig, B. Gotsmann, W. Häberle, et al. *Integrating nanotechnology into a working storage device.* Microelectronic Engineering, **83**(4-9) 1692–1697, 2006.

[262] M. Fujita, K. Wakabayashi, K. Nakada, and K. Kusakabe. *Peculiar Localized State at Zigzag Graphite Edge.* J. Phys. Soc. Jap., **65**(7) 1920–1923, 1996.

[263] Y. Kobayashi, K. Fukui, T. Enoki, K. Kusakabe, and Y. Kaburagi. *Observation of zigzag and armchair edges of graphite using scanning tunneling microscopy and spectroscopy.* Phys. Rev. B, **71**(19) 193406, 2005.

[264] Y. Niimi, T. Matsui, H. Kambara, K. Tagami, M. Tsukada, and H. Fukuyama. *Scanning tunneling microscopy and spectroscopy of the electronic local density of states of graphite surfaces near monoatomic step edges.* Phys. Rev. B, **73**(8) 085421, 2006.

[265] J. P. Eisenstein, L. N. Pfeiffer, and K. W. West. *Compressibility of the two-dimensional electron gas: Measurements of the zero-field exchange energy and fractional quantum Hall gap.* Phys. Rev. B, **50**(3) 1760–1778, 1994.

[266] T. P. Smith, B. B. Goldberg, P. J. Stiles, and M. Heiblum. *Direct measurement of the density of states of a two-dimensional electron gas.* Phys. Rev. B, **32**(4) 2696–2699, 1985.

[267] J. P. Eisenstein, L. N. Pfeiffer, and K. W. West. *Negative Compressibility of Interacting Two-Dimensional Electron and Quasiparticle Gases.* Phys. Rev. Lett., **68**(5) 674–677, 1992.

[268] D. L. John, L. C. Castro, and D. L. Pulfrey. *Quantum capacitance in nanoscale device modeling.* J. Appl. Phys., **96**(9) 5180–5184, 2004.

[269] Y. Martin, D. W. Abraham, and H. K. Wickramasinghe. *High-resolution capacitance measurement and potentiometry by force microscopy.* Appl. Phys. Lett., **52** 1103–1105, 1988.

[270] T. R. Albrecht, P. Grütter, D. Horne, and D. Rugar. *Frequency modulation detection using high-Q cantilevers for enhanced force microscope sensitivity.* J. Appl. Phys., **69**(2) 668–673, 1991.

[271] V. Palermo, M. Palma, and P. Samori. *Electronic Characterization of Organic Thin Films by Kelvin Probe Force Microscopy.* Adv. Mat., **18** 145–164, 2006.

[272] A. Gil, J. Colchero, J. Gomez-Herrero, and A. M. Baro. *Electrostatic force gradient signal: resolution enhancement in electrostatic force microscopy and improved Kelvin probe microscopy.* Nanotechnology, **14**(2) 332–340, 2003.

[273] M. Nonnenmacher, P. O'Boyle, and H. K. Wickramasinghe. *Kelvin probe force microscopy*. Appl. Phys. Lett., **58**(25) 2921–2923, 1991.

[274] R. Crook, C. G. Smith, W. R. Tribe, S. J. O'Shea, M. Y. Simmons, and D. A. Ritchie. *Quantum-dot electron occupancy controlled by a charged scanning probe*. Phys. Rev. B, **66**(12) 121301, 2002.

[275] J. R. Williams, D. A. Abanin, L. DiCarlo, L. S. Levitov, and C. M. Marcus. *Quantum Hall conductance of two-terminal graphene devices*. Phys. Rev. B, **80** 045408, 2009.

[276] D. A. Abanin and L. S. Levitov. *Conformal invariance and shape-dependent conductance of graphene samples*. Phys. Rev. B, **78**(3) 035416, 2008.

[277] T. Fang, A. Konar, H. Xing, and D. Jena. *Carrier statistics and quantum capacitance of graphene sheets and ribbons*. Appl. Phys. Lett., **91** 092109, 2007.

[278] M. Ishigami, J. H. Chen, W. G. Cullen, M. S. Fuhrer, and E. D. Williams. *Atomic Structure of Graphene on SiO_2*. Nano Lett., **7**(6) 1643–1648, 2007.

[279] A. Fasolino, J. H. Los, and M. I. Katsnelson. *Intrinsic ripples in graphene*. Nature Materials, **6** 858–861, 2007.

[280] E. Stolyarova, K. T. Rim, S. Ryu, J. Maultzsch, P. Kim, L. E. Brus, T. F. Heinz, M. S. Hybertsen, and G. W. Flynn. *High-resolution scanning tunneling microscopy imaging of mesoscopic graphene sheets on an insulating surface*. Proc. Natl. Acad. Sci., **104**(22) 9209–9212, 2007.

[281] J. C. Meyer, A. K. Geim, M. I. Katsnelson, K. S. Novoselov, T. J. Booth, and S. Roth. *The structure of suspended graphene sheets*. Nature, **446**(7131) 60–63, 2007.

[282] J. Martin, N. Akerman, G. Ulbricht, T. Lohmann, J. H. Smet, K. von Klitzing, and A. Yacoby. *Observation of electronhole puddles in graphene using a scanning single-electron transistor*. Nature Physics, **4** 144–148, 2008.

[283] Y. Zhang, V. W. Brar, C. Girit, A. Zettl, and M. F. Crommie. *Origin of spatial charge inhomogeneity in graphene*. Nature Physics, **5**(10) 722–726, 2009.

Postscript

A self-portrait etched in HOPG by AFM lithography.