Summer 8-12-2011

# Graphitic Surface Attachment by Single-Stranded DNA and Metal Nanoparticles

Luke A. Somers
*University of Pennsylvania*, lsomers@sas.upenn.edu

# Graphitic Surface Attachment by Single-Stranded DNA and Metal Nanoparticles

**Abstract**
Graphene and carbon nanotubes are extreme mechanical and electronic materials which have been the subjects of intense study and development since their discoveries. While many of their intrinsic properties have been discovered, their interactions with other materials are only beginning to be explored.

The noncovalent binding of single-stranded DNA oligonucleotides to carbon nanotubes and graphene has been seen to give rise to effective gas sensors. We examine similar systems to each of these in turn, imaging carbon nanotubes decorated with single-stranded DNA in Transmission Electron Microscope, and performing X-ray reflectivity of a single-stranded DNA film on graphite. The TEM study shows that the DNA bunches up along tubes but does not tend to clump on single tubes. Helical wrapping is not seen on single tubes. X-ray reflectivity shows that DNA on a graphite surface forms an inhomogeneous layer around 1.6 nm thick.

The differences between the various thicknesses of few-layer graphene are substantial though often underappreciated. These differences are highlighted in the system of several-nanometer metal particles on few-layer graphene flakes. We formed such particles by evaporation and annealing, then examined them in Scanning Electron Microscope. We found that gold nanoparticles were circular and experienced limited growth, with the radius varying as the number of layers to the 1/3 power. A theoretical explanation is given for this observation, based on an electrostatic interaction. This theory is also consistent with observations for titanium and silver nanoparticles. Ytterbium nanoparticles on graphene form instead into filaments. A related theory is presented showing that the same electrostatic interaction is capable of overcoming surface tension to deform particles from circularity.

**Degree Type**
Dissertation

**Degree Name**
Doctor of Philosophy (PhD)

**Graduate Group**
Physics & Astronomy

**First Advisor**
A. T. Charlie Johnson

**Keywords**
graphene, carbon nanotube, gold, ytterbium

**Subject Categories**
Condensed Matter Physics

# Graphitic Surface Attachment by Single-Stranded DNA and Metal Nanoparticles

Luke Somers

A Dissertation in Physics and Astronomy

Presented to the Faculties of the University of Pennsylvania
in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

2011

Supervisor of Dissertation

Signature: _____
Alan T. Charlie Johnson, Professor of Physics and Astronomy

Graduate Group Chair

Signature: _____

Dissertation Committee:
Gene Mele
Professor of Physics and Astronomy

Marija Drndic
Associate Professor of Physics and Astronomy

Paul Heiney
Professor of Physics and Astronomy

Gianlucca Piazza
Assistant Professor of Electrical and Systems Engineering

## Acknowledgements

This thesis is the culmination of my nearly seven years at Penn. In that time I've learned from, worked with, and taught many memorable people. The laboratory was a second home, and many friends have left, as I am now leaving. I am glad to have this opportunity to name some of those people who have had the most outstanding impacts.

First and foremost I must thank Charlie for all of the guidance and mentorship he has provided as my advisor. He granted me great freedom to devise solutions and even to set some questions. He did not however leave me alone, but pointed out the dead ends and traps, and offered ways forward.

As I started in the lab, Cristian Staii was in his post-defense time. In those few weeks he catapulted me into measurement sophistication.

I would also like to thank and commend my undergraduate collaborators, Sujit Datta and John Cinque.

Our postdocs Doug Strachan and Zhengtang Luo made me eat my peas until I liked them.

Doug Yates and Lolita Rotkina kept the electron microscopes running even with all of the abuse heaped upon them, a Herculean task, and used the rest of their energies pushing the boundaries of what could be clearly seen. Without them I would have been blind.

Jack Fischer provided the graphite sample used in portions of this work, and great assistance with an X-ray diffraction experiment not described here. I wish I had had the occasion to work with him more.

My collaborators on the X-ray reflectivity experiment were Kent Blasie, Joe Strzalka, Venkata Krishnan; their efforts to extract more from what we had were fairly drastic.

Gene Mele was a model collaborator, always ready to give time to clear up slight confusions and consider other possibilities. He also reshaped this dissertation, prompting me to produce a drastic improvement.

Paul Heiney made other vital corrections as a reader, and of course for the use of the X-ray system.

Lastly, without the patience and support of my wife, Ekaterina Shatalova, none of this would have been possible.

# Abstract

## Graphitic Surface Attachment by Single-Stranded DNA and Metal Nanoparticles

Luke Alexander Somers
Alan T. Charlie Johnson

Graphene and carbon nanotubes are extreme mechanical and electronic materials which have been the subjects of intense study and development since their discoveries. While many of their intrinsic properties have been discovered, their interactions with other materials are only beginning to be explored.

The noncovalent binding of single-stranded DNA oligonucleotides to carbon nanotubes and graphene has been seen to give rise to effective gas sensors. We examine similar systems to each of these in turn, imaging carbon nanotubes decorated with single-stranded DNA in Transmission Electron Microscope, and performing X-ray reflectivity of a single-stranded DNA film on graphite. The TEM study shows that the DNA bunches up along tubes but does not tend to clump on single tubes. Helical wrapping is not seen on single tubes. X-ray reflectivity shows that DNA on a graphite surface forms an inhomogeneous layer around 1.6 nm thick.

The differences between the various thicknesses of few-layer graphene are substantial though often underappreciated. These differences are highlighted in the system of several-nanometer metal particles on few-layer graphene flakes. We formed such particles by evaporation and annealing, then examined them in Scanning Electron Microscope. We found that gold nanoparticles were circular and experienced limited growth, with the radius varying as the number of layers to the 1/3 power. A theoretical explanation is given for this observation, based on an electrostatic interaction. This theory is also consistent with observations for titanium and silver nanoparticles. Ytterbium nanoparticles on graphene form instead into filaments. A related theory is presented showing that the same electrostatic interaction is capable of overcoming surface tension to deform particles from circularity.

# Table of Contents

# List of Figures

# 1: Introduction

## 1.1 Overview

The carbon forms Graphene and Carbon Nanotubes have occupied a large fraction of the attention of the condensed matter community since their discoveries. They possess exceptional tensile strength[26,99] and electronic mobility[27], among other exceptional properties. These all have great promise for technological applications, and the high scattering length has permitted astonishing pure science results such as the observation of the Quantum Hall Effect at room temperature[73].

Electronic-readout chemical sensing is one of the technological applications for which single-walled carbon nanotubes and graphene are well suited, for a variety of reasons that will be explored in section 1.5. The Johnson group has developed a variety of gas and fluid sensors based on these materials, with the sensitizing agents being single-stranded DNA (ssDNA)[42,47,48,60,89], RNA, antibodies to specific antigens[101], or human olfactory proteins[34]. The mechanisms for the sensing based on ssDNA and RNA is unknown and presumably complex. One of the focuses of this work is to investigate the form of systems closely analogous to such sensors.

Graphene film thickness determines its properties in ways that do not simply follow from scaling of intensive properties. One such property, which forms the second focus of this work, is the formation of metal nanoparticles on the surface. We show this dependence and give a theory accounting for it.

## 1.2 History of $sp^2$ carbon

Graphene, the conceptual component material of graphite and parent material of carbon nanotubes, is a two-dimensional crystal with a hexagonal lattice and a two carbon atom basis; all bonds are $sp^2$. Graphene was long used by theorists as a model system for derivations of properties of these structures and as a subject in its own right [94], even when its existence on is own was believed to be physically impossible. This belief in the impossibility of graphene was due to the seminal 1937 paper by Lev Landau [53], the first part of which gave the theory of second order phase transitions, and the second part of which declared the impossibility of free-standing two dimensional structures in three dimensional space. This paper was correct; the following deduction that graphene cannot exist was flawed. Even setting aside the issue that the paper clearly doesn't exclude the case of a thin film on a backing, which is how graphene is usually arranged in the laboratory, there is a more core issue. While graphene is most simply described as two dimensional, it remains graphene even if slightly deformed. The fluctuations required by Landau's theory do not destroy graphene, nor force it to immediately roll up into a spiral graphite structure or fold into graphite, nor even eliminate its essentially two dimensional character. While folding and rolling are ultimately the stable states, wrinkles provide sufficient metastability that graphene can exist in uncondensed states for extended timescales.

Graphite is composed of a stack of graphene sheets. When graphene stacks, there are two relevant stackings: Bernal stacking, as in graphite; and turbostratic, which is to say,

misaligned with the effect that they do not stack as closely. Few-layered graphene flakes are frequently turbostratic when they are formed by folding or misaligned growth [37]. Bernal stacking is an alternating A-B stacking. Adjacent layers are 'π-π' stacked, as the π electrons are those making contact. This stacking can be visualized as one layer being a copy of the next, rotated 180° around one of the atoms. Thus, half of the atoms are directly above atoms in the next plane; the other half are lined up with the centers of the hexagons of both neighboring sheets [92]. The Bernal stacking thus maximizes the difference between the two basis atoms by letting one be in contact on both sides while letting the other be free on both sides.

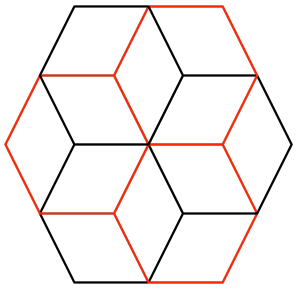

Figure 1: a top-down view of the bond network of a section of two layers of graphite. The bonds in the A layer are in black, the others are in B layer. The atom in the center and the atoms on the center of each side are directly stacked; the atoms at the 3-way intersections are not.

Buckminsterfullerenes (or more briefly, 'buckyballs' or 'fullerenes') were discovered in 1985 by Harry Kroto and the Smalley group at Rice University [52] and are the least related to graphene of these three. They are carbon cages, closed by the inclusion of pentagons. The most stable fullerene, and the first noticed[20], has a 'soccer ball' shape: $C_{60}$. Other shapes are less prominent, but have been catalogued[100]. For example, one form of $C_{80}$ has two pentagon-containing caps, and a roll of hexagons is laid between.

A particular class of fullerene takes this concept to extremes: carbon nanotubes. They were likely created for some time as a byproduct of carbon fiber production, then certainly created during fullerene production, and finally entered scientific awareness in 1991 with a paper by Ijima [39]. Within two years, reliable processes for production of single-walled tubes were available [9,40], and shortly thereafter could be grown in bulk[35]. Carbon nanotubes were found to have the highest tensile strength known[26,99], to grant their electrical charge carriers extraordinarily high mobility [26,27,99], to be sensitive to very small changes in their environment [30], and be extraordinarily black when grown as a 'forest'[67]. All but the last of these properties were later discovered to also apply to graphene, which exceeded carbon nanotubes in the first two categories.

Carbon nanotubes can be described nearly completely by a 'wrapping vector', which can be most easily understood by modeling the tube as rolled-up graphene. The wrapping vector is the smallest vector connecting two points on the graphene plane that map onto the same point of the nanotube. In Figure 2, the wrapping vector extends along the short side of the cut-out rectangle.

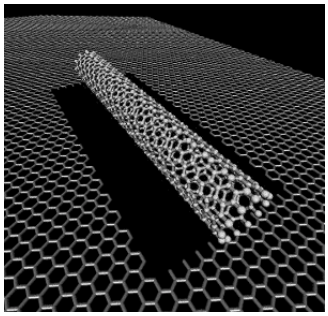

Figure 2: a cartoon of a carbon nanotube rolled up out of graphene. Image courtesy of Bob Johnson.

Graphene itself had been observed in peculiar and difficult-to-reproduce circumstances (reviewed in The Rise of Graphene [32]), but it was in 2004 that André Geim and Konstantin Novosëlov devised a method for isolating it and made a good start on experimentally determining its properties [71,72]. Known affectionately and derisively (often at the same time) as the 'scotch tape method', this easy technique quickly opened graphene itself up to investigation by a large number of groups.

Graphene was soon found to have many of the same exceptional properties as carbon nanotubes, with graphene slightly edging out carbon nanotubes in most respects such as tensile strength [54] and carrier mobility [15], and making claim to entirely new extremes such as the maximum elastic strain in a crystal.

A recent effort has been to develop methods for growing graphene over large areas rather than exfoliating it into tiny fragments. The older method, still with an element of promise, is to evaporate silicon out of a silicon carbide surface, producing epitaxial graphene [8]. More recently, catalyzed chemical vapor deposition has emerged[4,49,56,62,83]. In this family of processes, carbon feedstock is flowed over a metal film at 1000 K and dissolves into the surface layers of the metal. As the film cools, carbon falls out of suspension and forms graphene. The metal can be dissolved away, and the graphene transferred to any desired substrate.

## 1.3 Electronic theory of graphene

The bonds between the carbon atoms in graphene are hybridized as $sp^2$, leaving the $p_z$ orbital only weakly hybridized. This orbital protrudes out of the plane in both directions. The $sp^2$ hybridization forms σ bonds, which are sufficiently strong that the participating electrons' contributions to the electronic structure of graphene are buried deep in the valence bands. The remaining p electrons are substantially freer to move, so the $p_z$ orbitals combine into less local π orbitals. As the unit cell of graphene contains two atoms, there are two π electrons per cell.

The band structure of graphene is usually calculated with a simple tight binding model taking into account the crystal symmetries and nearest neighbor interactions, but not wave function overlap[85]. This approximation yields the relatively compact expression

$$E = \pm\gamma\sqrt{1+4\cos\left(\frac{\sqrt{3}}{2}ak_x\right)\cos\left(\frac{a}{2}k_y\right)+4\cos^2\left(\frac{a}{2}k_y\right)}$$

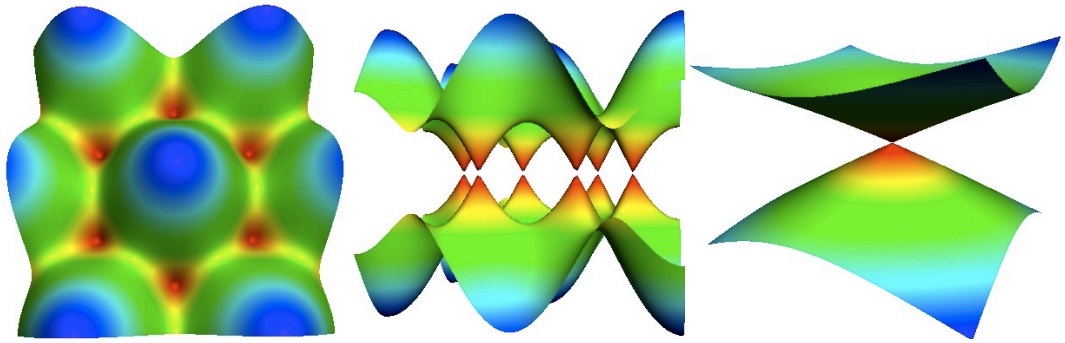

Figure 3: Three views of the approximate band structure of graphene, not cutting off at the edge of the Brillouin zone. A (left): A top-down (or bottom-up) view; the Brillouin zone boundary stretches directly between the small dark dots (red online), while the origin is on the central peak. B (center): A side view – the dark dots in A are the sharp meeting points here. C (right): a close-up of one of the K-points.

This simple tight binding approximation is very good around the K-points, where the bands meet. Away from them, and in particular around the Γ (origin) and M (directly between two adjacent K-points) points, this approximation is not so good. Accounting for the direct overlap between the $p_z$ states improves the approximation drastically for these ranges of momentum[81].

Each carbon atom contributes one $p_z$ electron, and there are two atoms per unit cell. Two electrons per unit cell are enough to fill one band with spin up and spin down states. Thus, the valence and conduction bands meet at the K-points. This is where the tight binding approximation is good, so the electronic structure as pertains to conduction is indeed well-described by this tight binding model. The most pertinent observation is that the band structure can be closely approximated as two double cones. This can be most clearly seen by shifting the Brillouin zone to be centered on an M-point.

Having a conical band structure has several consequences. First, all of the charge carriers, regardless of momentum, have the same speed, around 0.3% of c. This makes electron dynamics in graphene a highly unusual two-dimensional system of effectively massless charged fermions. Their behavior is best modeled using the Dirac equation. This in itself has many consequences, such as Klein tunnelling[98].

The most pertinent of these effects is an unusual form for electrical field screening. Because the energy of carriers varies linearly around the K-points, the density of states is only

$$\mathcal{D}(E) = \frac{2E}{\pi(\hbar V_F)^2}$$

This limits the charge concentration in a different way than in a massive carrier system, in which a large number of carriers become available at modest energies in excess of a band minimum. Rather than decaying exponentially, a voltage applied at one surface decays as the inverse square of the number of layers [24,77,84].

The Fermi velocity of the carriers is also important. Though this speed is not extraordinary for a metal, it is extraordinary for a semiconductor near charge neutrality. This contributes to one half of its most outstanding feature: its extremely large electron mobility.

$$\mu = (q_e n \rho)^{-1} = V_F L \frac{q_e}{E_F}$$

The other material factor in mobility is the scattering length, L. Graphene has an extremely long scattering length – in the best cases, the scattering length of 1.2 μm arose solely from the edges of the graphene and the electrical contacts [15]. This long scattering length arises from graphene's having little of two of the three principal sources of scattering: crystal defects, phonons, and external scatterers.

First, graphene frequently occurs with high crystal perfection. This is due in part to the strong in-plane bonds but also to the material's two-dimensional nature. Even holding the fraction of misplaced atoms constant, laying them out flat will give a larger distance between defects than arraying them in three dimensions.

Second, electron-phonon scattering is extremely weak. In contrast to 2-dimensional electron gases, which also have very high mobility at low temperatures, the temperature dependence of graphene-phonon scattering is only linear [78]. At low temperatures, GaAs 2 dimensional electron gases have much larger electron mobility than in graphene, but the phonon scattering rises exponentially; with its linear dependence, graphene tolerates temperature much better. Above 200 K, substrate phonons become a significant contributor to the phonon scattering on Silicon Oxide substrates.[18]

The third form of scattering is from external defects. This consists of charges in the substrate (if there is a substrate), and debris under and on top of the graphene. For graphene this is usually the strongest form of scattering, especially under commonly encountered circumstances. That is, electronic circuits involving graphene are usually backed by silicon oxide, which has an abundance of charge traps acting as scattering sites. Scum from lithography, and even adsorbed water, will aggravate this unless it is specially removed.[90]

The best efforts to maximize graphene's mobility required etching away the substrate, leaving graphene suspended, and baking it clean; this resulted in a mobility of $2 \cdot 10^5$ $cm^2/Vs$, which is the highest recorded room-temperature electron mobility in any material [15]. Even removing only the lithographic scum results in a substantial improvement in room temperature mobility, rising for instance from 1600 to 5500 $cm^2/Vs$[21]. This mobility is already well in excess of the mobility of Silicon, around 750 $cm^2/Vs$. Short of removing the substrate altogether, the scattering from charge traps and substrate phonons can be reduced by selection of a new substrate; this is an ongoing search[25,38].

Intriguingly, if graphene is layered turbostratically, it typically displays electronic behavior most similar to a single layer, with its massless behavior intact though modified. This can be understood through perturbation theory – the two sets of single-sheet K-points are misaligned, and so the mixing of their states with the states from the other sheet will face a large energy difference, inhibiting hybridization [59,65]. This is in contrast to A-B graphite, in which dispersion is quadratic, yielding an effective mass for low energies[66]. For moderate energies, even A-B-stacked few layer graphene has effectively massless carriers.

### 1.4 Graphene transistors

We have compared graphene to silicon above because it has the potential to be an excellent switching material for electronics. Its large mobility makes it a good material

for analog amplifiers; however, for digital electronics, the ability to turn off nearly completely is vital. Bulk graphene is semimetallic lacking a gap. The reduction in the density of states around the k-point has produced an on-off ratio as high as 30, though this is exceptional[71].

However, higher on-off ratios are possible in graphitic structures: long before this was an issue for graphene itself, carbon nanotubes had already achieved on-off ratios of $10^5$, with 1000 being common. Modeling a carbon nanotube with graphene this requires imposing a periodic boundary condition for separations equal to the wrapping vector; thus restricting the possible states. 1/3 of the wrapping vectors are lined up so as to include the K points as valid wavevectors. The other 2/3 of wrapping vectors exclude the K points; such tubes thus have a gap, and are semiconducting.

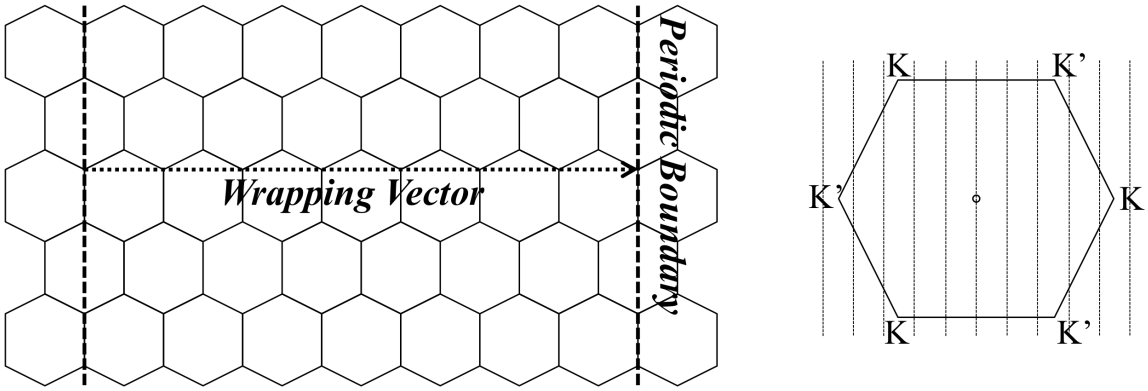


Figure 4: the correspondence between a real-space lattice (left) and a K-space Brillouin zone. This (7,0) zigzag tube is semiconducting because the valid k-states (dotted lines) do not intersect the K-points.

Carbon nanotubes with on/off ratios in excess of $10^5$ have been found[3]. A similar effect can be produced in flat graphene through a boundary effect. By constricting the lateral extent of the graphene, the lateral momentum is similarly quantized, permitting the development of FET with high on/off ratio[5,41,55,61]. A difficulty arises from the extreme narrowness of the constriction required to make a transistor that can turn off strongly: 10 nm is a rough maximum on this width. Standard lithographic techniques cannot produce such dimensions. The papers referenced earlier each used a different method for achieving such narrow strips; each method has distinct strengths and weaknesses.

For electronics applications, graphene is far more promising than carbon nanotubes, for several reasons. Principally, it is difficult to scale up production of carbon nanotube devices. 1/3 of carbon nanotubes are metallic, and separation techniques rely on solvation. This dirties the tubes, requiring further processing to re-clean them. An additional difficulty arises as the tubes must be deposited on the surface in precisely selected locations. In contrast, graphene can in principle be laid down as a single continuous layer, then the undesired parts etched away. Secondly, bulk graphene can be used for interconnects that seamlessly become the transistor material, minimizing carrier scattering.

## 1.5 DNA-decorated graphitic surfaces and gas sensing

Semiconducting carbon nanotubes make an excellent material for chemical sensing with optical or electrical readout. This is due to the confluence of several factors. First, these tubes have a high electron mobility. The sensitivity of the conductivity to changes in carrier density are given by

$$\frac{d\sigma}{dn} = e\mu_e$$

so a carbon nanotube's high mobility will let it produce a large change in current for a small change in carrier density. This directly amplifies the effect of charge injection mechanisms, which will be one of the principal sensing mechanisms. It also helps indirectly, as this high mobility is in large part due to low scattering. If the presence of an analyte causes additional scattering, it will be a comparatively large addition to the scattering of the tube.

Secondly, these tubes are all surface – there is no parasitic conduction through the bulk adding spurious current.

Thirdly, being made of carbon, the chemistry of making covalent modifications to incorporate functional chemical-selecting molecules are already available[10,46,96]; though this is not featured in this work, other members of the laboratory have used it[101].

Lastly, nanotubes in themselves are not strongly chemically active, so that chemical affinity created by the functionalization will not compete with a strong native response.

The Johnson group has been working with electrical readout carbon nanotube gas sensors, functionalized with short oligonucleotides of single-stranded DNA noncovalently bound to the nanotube sidewall [89]. These sensors are not specific to a single chemical target, but each sequence provides a distinct reproducible response to a wide variety of stimuli. The sensors are quite sensitive, being able to detect parts-per-billion concentrations of some analytes, and can distinguish between very similar molecules, such as octanal, nonanal, and decanal. Despite differing by only a single atom in chain length, decanal is sufficiently distinct that it elicits a response opposite in sign from the other two for some sequences[47].

By varying the sequences used for functionalization, it was possible to determine several things about the nature of the sensing response: it depends on the details of the sequence, not merely the bases present – alphabetizing sequences 1 and 2 changed their sensing responses totally[48].

The same sort of device can be made out of graphene instead of carbon nanotubes, to similar effect[60].

Even aside from the practical questions of how this can be best exploited, it would be very desirable to fully understand the details of the detection mechanism. Critical to understanding this is the issue of how the ssDNA conforms to the nanotubes.

The DNA backbone is a sugar. This, being polar, does not bind well to nanotubes, which are nonpolar. The bases, however, are not polar, and so do bind well. Further, they contain aromatic rings. These π orbitals form a specific π-π stacking on graphitic surfaces. Beyond this, at the start of this work, nothing was known.

The first two experiments in this work show more about the system's details, by direct imaging of ssDNA on Carbon Nanotubes in Transmission Electron Microscope,

and by X-ray Reflectivity measurements of ssDNA adsorbed on graphite. These experiments were complemented by a series of Molecular Dynamics simulations performed by Bob Johnson[43-45].

## 1.6 Summary

In the first four parts of this chapter, background material concerning graphene and the conceptually derivative material carbon nanotubes was presented. There are three notable consequences of graphene's band structure: its carrier mobility is extremely high; it screens weakly; and an energy gap can be opened by quantum confinement effects on scales ~ 10 nm and below. Many of the properties of carbon nanotubes can be understood by "zone-folding" arguments based on modeling the tube as having the band structure of graphene combined with wave vector quantization associated with a periodic boundary condition.

These properties are relevant in various ways: the mobility enters only as motivation to better understand the graphene-nanotube family of materials, and to understand why certain characteristics are desirable. The screening plays a mechanically important role in the details of the nanoparticle experiments. The means of rendering graphene semiconducting connects to possible uses of the nanoparticle phenomena observed.

The fifth part of this chapter described the system and circumstances leading to and motivating the two ssDNA experiments in this work.

Chapter 2 is a very brief summary of the results of the work.

Chapters 3 and 4 will explain the methods used in this work, divided into experimental techniques and analysis methods. The focus in these chapters is on the concepts; practical notes about some of the techniques can be found in the first appendix.

Chapter 5 shows an experiment probing the conformation of single-stranded DNA to carbon nanotubes, by transmission electron microscopy,

Chapter 6 concerns the X-ray reflectivity of films of single-stranded DNA applied to graphite.

Chapter 7 describes an experiment revealing that some species of metal nanoparticle on graphene experience limited growth, and some assume elongated shapes. Several metals are examined.

Chapter 8 provides a theoretical framework for understanding the observations of the previous chapter.

# 2: Summary of Results

There are two groups of experiments presented in this work. The first group pertains to single-stranded DNA(ssDNA) on graphitic surfaces, and contains two experiments: imaging single-stranded DNA adsorbed onto carbon nanotubes in Transmission Electron Microscope, and obtaining X-ray reflectivity curves for single-stranded DNA on graphite. The other group is imaging metal nanoparticles on few-layered graphene flakes in Scanning Electron Microscope. The particular metals imaaged were gold, silver, titanium, palladium, and ytterbium. The purposes of these experiments and their results follow.

## 2.1 Single-stranded DNA on graphitic surfaces

The Johnson group has for years been working on gas sensors based on a combination of single-stranded DNA (ssDNA) with single-walled carbon nanotubes. These combined the moderate chemical activity and outstanding diversity of ssDNA with the nanotubes' semiconductor properties. A similar device design uses ssDNA and graphene. These devices' properties are described in some detail in the introduction. This work proceeded initially on the supposition that there would be some sensing mechanism – charge injection or scattering changes. However, there was little information about the actual shape the ssDNA would assume as it adsorbed on the carbon nanotubes – only some AFM measurements, which did not address most of the critical questions about how the sensing worked.

These were:

Do strands chirally wrap around the tube, or assume some other shape?

Does the strand lie entirely close to the tube, or are there extensive excursions?

When multiple strands meet, do they remain separate, chain end-to-end, bundle with each other, or arrange in some other way?

How do things change at tube intersections?

These questions were addressed in multiple ways – by direct observation in Transmission Electron Microscope (TEM), by Molecular Dynamics simulations, and by X-Ray Reflectivity. I performed the experimental examinations, while Robert Johnson performed the simulations[43-45].

### 2.1.1 TEM of ssDNA on carbon nanotubes

We observed the strands tend to line up end to end, often giving the illusion of being a single longer strand. All of the strands in the chain remain in direct contact with the tube. These chained strands are not seen to wrap in a chiral fashion, confirming the simulation result. Some samples, especially those with a large amount of added material present, displayed strong clumping behavior, with clumps somewhat more frequently seen near the surface and at intersections.

The simulations indicated that individual strands would frequently curl tightly around the tube at the ends. We observe this on the chains and on small loops, not for the individual strands in a chain.

### 2.1.2 X-ray reflectivity of ssDNA on graphite

X-ray reflectivity is used to find patterns in the vertical density distribution of a surface layer. In this case, we used a thin ssDNA layer on graphite. We used graphite as a model material for graphene or carbon nanotubes to simplify the experiment and its interpretation. The system is of course far more similar to graphene than carbon nanotubes.

The X-ray reflectivity data of ssDNA on graphite produced two significant results. First, it agrees with the TEM data that the ssDNA tends to produce an inhomogeneous layer. Second, the general scale of this layer is 1.6 nm, which is consistent with a single layer of ssDNA lying mostly flat.

## 2.2 Metal nanoparticles on few-layered graphene flakes

For some time the Johnson group performed experiments on using metal nanoparticles to etch into graphene in order to create semiconducting graphene nanoribbons[23]. Some metals did not etch the graphene but instead balled up on it. During these attempts, a tendency for the size of the nanoparticles to vary sysematically by location was noticed. This prompted a more careful investigation of what sizes nanoparticles would assume under various conditions.

We found that some metals' equilibrium formations were circular droplets, with the size depending on graphene thickness; Ytterbium formed irregularly shaped filaments. Theories to explain these observations were developed, based on a competition between surface tension and electrostatic dipole repulsion brought on by work function difference.

The observations of gold particles yielded a layer-dependent mean radius function $3.23 \pm 0.34$ nm Layers$^{0.331 \pm 0.061}$. The theory predicts 3 nm Layers$^{1/3}$. Titanium particles followed $4.5 \pm 1.0$ nm Layers$^{0.51 \pm 0.19}$, with the theory predicting 5.7 nm Layers$^{1/3}$. Silver particles would not be expected to follow this pattern, and did not.

These observations are shown in chapter 7; the theories are given in chapter 8.

# 3: Experimental Methods

Diverse techniques were employed to achieve the three goals established in chapter 2. Each of the three sections of this chapter walks through one of the experiments, giving its requirements, a summary of the method of satisfying those requirements, and then a detailed description of each step.

## 3.1 TEM of ssDNA on carbon nanotubes

One experiment in the effort to determine the physical structure of ssDNA adsorbed onto the surface of carbon nanotubes was to look at such structures in a Transmission Electron Microscope.

Transmission Electron Microscopy only approaches atomic resolution when the object to be imaged is free-standing. Thus, we suspend our carbon nanotubes across a gap.

Additionally, the carbon nanotubes must be extremely clean to minimize non-ssDNA adsorbed species. This rules out tubes deposited from solution; such tubes are very difficult to completely clean. Tubes grown via catalyzed Chemical Vapor Deposition (CVD) and used in place are far cleaner, so we use them.

Now, the common substrate for TEM is a grid of copper. This interfered with the CVD process sufficiently that we elected to make a chip with holes in it by a preexisting recipe.

To summarize sample fabrication, we start with a silicon wafer coated with 100 nm of nonstoichiometric Silicon Nitride ($SiN_x$). We etch away lithographically defined areas of this chip to leave a thin membrane with holes in it. Carbon nanotubes are grown across these holes by CVD. ssDNA is applied from solution. The resulting sample is then imaged in TEM.

When imaging, there were two available choices of beam energy: 80 keV and 200 keV. The highest resolution is granted by 200 keV, both due to the intrinsic limits of such electrons and by the fact that the finer of the two available instruments (the JEOL 2010F, the other being a JEOL 2010) worked at 200 keV. However, electrons at 200 keV cause damage to carbon nanotube samples much more rapidly than electrons at 80 keV [87]. Additionally in favor of 80 keV, carbon is substantially more visible at lower beam energies. Therefore, we elected to operate at 80 keV, on the JEOL 2010.

The specific steps of this experiment are:

- Photolithography, to define the region of the membranes
- Sulfur Hexafluoride Reactive Ion Etching, to expose the chip body
- Potassium Hydroxide etching, to create the membranes
- Electron-Beam Lithography, to define the areas of the windows
- Sulfur Hexafluoride Reactive Ion Etching, to create the windows
- Oxygen Reactive Ion Etching, to clean the chip
- Chemical Vapor Deposition, to create the nanotubes
- Solution deposition of ssDNA, to decorate the nanotubes

• Transmission Electron Microscopy of the resulting sample

I describe each of these steps in detail below.

### 3.1.1 Lithography

The concept of lithography is to cover the substrate with a protective layer that can be selectively removed; then, any processing, such as deposition or etching, on sample will affect only the exposed areas. It is essentially a made-in-place disposable stencil.

In modern lithography, this stencil will be made of a variety of plastic sensitive to either high-energy photons, making it a photoresist; or electrons, making it an e-beam resist. If the photons or electrons weaken the resist and enable it to be dissolved away in a development step, it said to have positive tone, and be a positive resist; if the photons or electrons strengthen the resist so that only the unaffected regions are developed away, it has negative tone and is a negative resist. Figure 1 shows a positive lithography process, skipping over the technique-dependent exposure step.

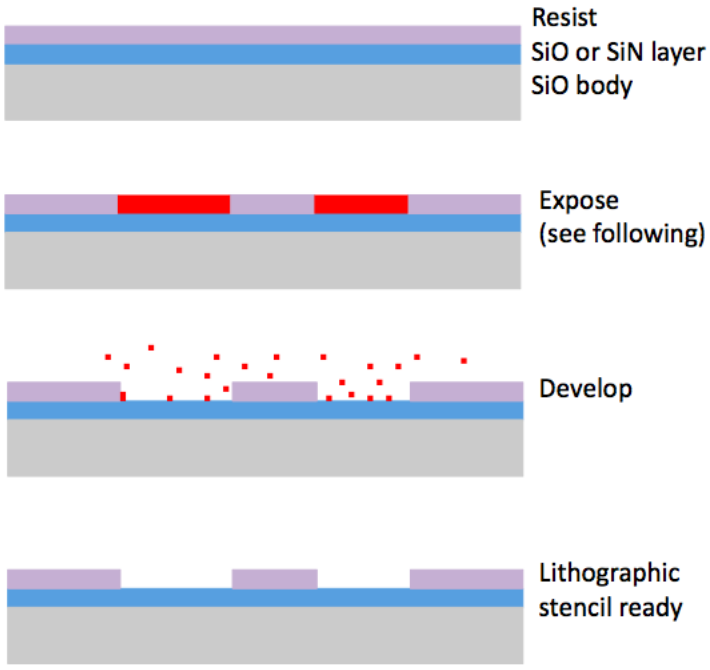


Figure 1: visual outline of a lithographic process, as described in the text. This lithographic process is positive.

The process is largely the same between the two techniques. The chip is cleaned as above. Resist is applied in a spin-coater, and the chip is spun at a controlled speed and duration to achieve a specific known thickness of resist. This is baked in dry air to remove the solvent, rendering the resist a solid rather than viscous fluid. The resist is then exposed, as will be described shortly. For some resists, a second bake is used after exposure to complete a chemical reaction that exposure began. In all cases, development is completed by immersion in a resist-specific chemical that dissolves the appropriate portions of the resist.

Exposure, the heart of the process, is now described. In E-beam lithography, the sample is placed in a Scanning Electron Microscope and the SEM beam controls are used

to control where the resist is exposed; each area is exposed one after another with a controlled charge dosage. In Photolithography, a prefabricated opaque mask is placed on the resist and ultraviolet light shone through with a controlled energy dosage.

Thus, Photolithography requires substantial preparatory time and effort to make the mask; however, with that mask, very large and complicated patterns can be generated in under an hour; the exposure time itself will typically be under a minute. E-beam lithography requires little preparatory time, but the time spent exposing depends on how much is being written; large patterns can take days.

The other notable difference between the techniques is the limiting factors on their resolution.

Photolithography is limited by the wavelength of the light used. Light passing through a slit smaller than its wavelength will not behave according to ray optics, but diffract and spread out. Shortening the wavelength permits smaller features, which is part of why UV light is used rather than visible light. While photons can be generated with arbitrarily short wavelength, as the wavelengths shorten, it becomes harder to make masks capable of blocking the light while still being thin enough to be fine-featured. Very high-end photolithography can be used to produce features of 50 nm. The lithography used in this work did not approach this performance, having no features smaller than 1 mm; the same equipment can be used to generate features of 1 μm.

The limiting factor for the precision of high-end e-beam lithography is of a fundamentally different nature. Even if the beam is focused to a single nanometer spot size, the beam electrons scatter with the resist and substrate, both making them deviate from their course and knocking other electrons free. These secondary electrons can also expose the resist, widening the features. These problems are exacerbated for thick resist layers. Extremely high-end E-beam lithography systems have been used to produce features of 15 nm[16]; we can produce 50 nm features even on thick resist chosen more for etch resistance than resolution. The work presented did not push this limit either, making features no smaller than 1 μm.

### 3.1.2 Reactive Ion Etching

Reactive Ion Etching (RIE), a form of plasma etching, works by the application of energetic ions to the surface; they there chemically interact with the target, and the resulting compounds are volatile. This opens up the lower layers to interaction and removal. Because the process is essentially chemical, there can be strong selectivity between target materials. Because the process involves bombardment by highly energetic radicals, polymers such as resists are going to be etched without as much regard to the specific chemistry involved.

The energetic ions are created by an intense radio frequency (RF) electrical field. Neutral molecules striking one of the highly charged plates absorb some of the charge and are repelled; they strike other molecules, creating a plasma. The RF electrical field accelerates ions back and forth; those near the sample will strike it energetically. An alternative method of creating the plasma is with a rapidly oscillating magnetic field; this creates a large enough electrical field in the gas to strike a plasma; but this does not accelerate the ions into the sample. This system is an Inductively Coupled Plasma (ICP), and the resulting plasma is said to be ‘soft’ because its components have less downward momentum than REI plasmas; more of the action on the surface is chemical rather than

mechanical bombardment. Most often, a recipe involving ICP will also use a simultaneously applied RIE component.

Plasma etches are slow, of the order 1 nm/s, and they attack the resist at a similar rate to the target compounds. This places a hard limit on the depth of patterned plasma etches, unless the extra step is taken of evaporating a metal etch stop to replace the resist. A less involved and somewhat effective method of improving the selectivity is to use ICP; we never needed to do either of these in the given recipes.

The specific plasma etches used above are Sulfur Hexafluoride ($SF_6$), and Oxygen ($O_2$). $SF_6$ plasma etches $SiN_x$, and can work its way through the nitride layer conventionally used in 90 s of etching at 150 W. It also slowly attacks SiO, and does not appreciably affect crystalline Si. If the etch is to expose a silicon surface, one can tell that the etch is finished when the exposed areas are bright silver in appearance. Brown appearance indicates that some $SiN_x$ remains on the surface; further etching will be required to expose the Si properly. If the etch is to punch holes in a $SiN_x$ membrane, one can tell that the etch is finished by examination in an optical microscope: topside illumination will show black holes where it has etched through.

$O_2$ does not affect these Si compounds at all at low energies, and so is ideal for cleaning organic matter (including leftover resist layers) from chips.

### 3.1.3 Potassium Hydroxide wet etch

A Potassium Hydroxide (KOH) wet etch is used to attack the crystalline Si. KOH attacks SiO only very slowly, and has no effect on $SiN_x$. The etch into Si has the very useful property that while it can attack a [100] or [110] plane, the etch rate of Si from a [111] plane is exceedingly slow, a few thousandths the rate of etching from other planes. This makes a stop-etch layer extend right through the middle of the crystal, wherever it occurs that the KOH cannot get around the plane to attack it from another angle. The $SiN_x$ shields most of the chip; the exposed region becomes the base of a pyramid of etching that extends up through the chip until stopped by the $SiN_x$ on the far side (figure 2). Etching for much longer than required will normally have no deleterious effect.

Figure 2: SiN membrane on Si chip, left by KOH wet etch. Not drawn to scale, though the angle is correct.

As the etch proceeds along the 55 degree angle of the [111] plane away from the chip's [100] orientation, the ending size of the membrane depends strongly on wafer thickness. Each wafer will produce different sizes of membrane, or, if too thick, no membrane at all.

KOH is gradually added to water in a 1:2 mass ratio. This is placed on a hot-plate with integral magnetic stirring. The temperature is raised to 120 C, and the stir-bar keeps the solution moving. The chip is hung into this solution, not letting any metals touch the liquid. This is covered and left for 8 or more hours, typically overnight.

Once the windows are completely etched, one can see them by looking at a bright light through the chip. The windows appear as points of light. When examined in the microscope with backside illumination, they are very easy to find. With topside illumination only, they may be hard to make out, especially if the surface is non-ideal.

This process is delicate; additional details are provided in appendix 1.

### 3.1.4 Window creation

Windows in the membrane are patterned in e-beam lithography (described above) and opened by $SF_6$ etch (described above). The particular shape of windows is important. Too small, and there is little area, giving less chance for tubes to stretch across. It is worse than simply proportional to the area: the tubes hanging across may be able to find a path around a small window even if they are initially draped across it. If the windows are made too large, the membrane itself is at risk of destruction. Initial experiments (those producing most of the results used) used 5 micron square windows separated by 15 microns. However, the yield of this pattern was quite low, frequently finding zero tubes over windows anywhere on the chip. Later experiments employed a brick-wall-like pattern of rectangular windows devised by Yangxin Zhou (shown in figure 3); this was found to be an excellent combination. A large area was exposed, tubes could not reach around, and the membrane remained strong.

### 3.1.5 Nanotube growth

Carbon nanotubes are commonly grown in the lab via Chemical Vapor Deposition (CVD). The physical process is that methane at 900 C dissociates upon contact with iron nanoparticles on the substrate; the hydrogen escapes, and the carbon dissolves into nanoparticle. When a fluctuation brings the iron's carbon concentration supersaturated, a carbon nanotube springs out of the particle.

The setup is shown in figure 5; it was made following developments by the Lieber and Dai groups[36,51]. The system is open; however, a flow of Argon prevents air from flowing down the exhaust tube or in any small leaks; a similar $H_2$ flow helps ensure that the chemistry in the chamber is reducing even with whatever $O_2$ may slip by the Ar.

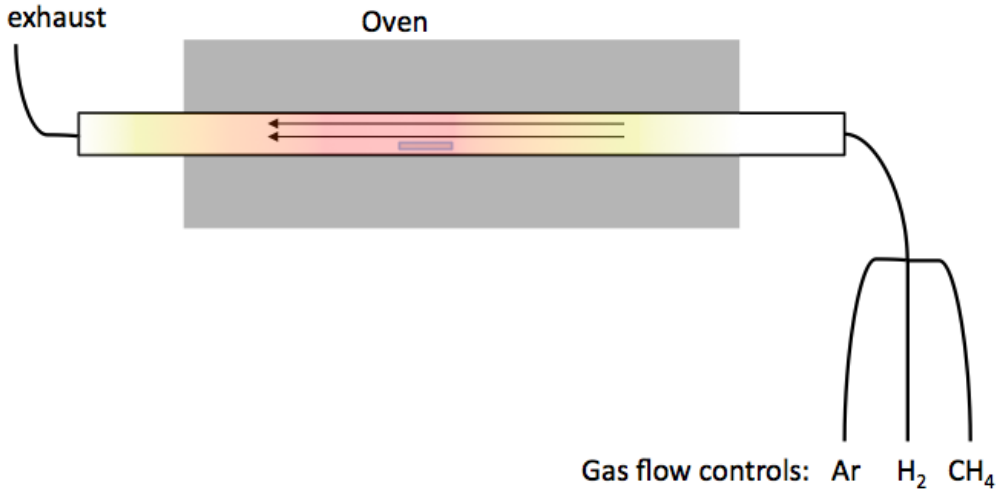


Figure 3: a schematic of the CVD oven setup.

The recipe used follows:

Clean the chip thoroughly. Dissolve 50 mg/L of Iron III nitrate in IPA and apply this to the substrate by dropping on, then spinning it dry at 3000 RPM.

Place the sample in a quartz tube. Optionally, place a retarder – something to force the gases to flow close to a surface – in the upstream end of the tube; this will greatly reduce the impact of chip placement (retarder design is addressed in appendix 1). Seal the tube in a cool oven and gas flow setup as shown in figure 3, placing it such that the sample will be in the middle or downstream end of the tube (downstream for more tubes).

Flush all gas lines for 10 minutes; then, while heating, flow Ar and $H_2$ over the sample at 600 and 300 standard cubic centimeters per minute (sccm) respectively. Once the furnace is at 900 C, let the system equilibrate for ten minutes. This reduces the catalyst completely.

Flow $CH_4$ at 2.5 standard liters per minute for five to ten minutes. Most carbon nanotube growth occurs in the first few minutes. Growing longer produces progressively less additional tubes, and increases the amount of amorphous carbon on the surface.

Once the $CH_4$ flow is shut off, allow the oven to cool. Once the oven temperature has dropped to 400 C, the nanotubes are relatively safe from combustion and the $H_2$ flow can be halted; this is also a good time to open the oven to speed cooling, as it is not so hot as to cause major thermal shock to the oven. At 100 C, the Ar and $H_2$ flows can be halted and the sample removed. The principal time-consuming element in this process is cooling, which can take several hours.

CVD does not proceed as prolifically on $SiN_x$ as it does on SiO. Several methods were used to encourage adequate growth. The more obvious but more involved method was to sputter SiO onto the $SiN_x$ surface, thereby protecting the catalyst from the SiN.

The second and more commonly used method was to bring the flowing gas to high temperature by forcing it through a retarder as it entered the oven (figure 4). This greatly increased tube growth without any visible disadvantages. This was also found to be a good way to homogenize the temperature in the oven, making the exact placement of the chip matter less. Lower tube densities can be obtained while using a retarder by reducing the iron nitrate concentration.

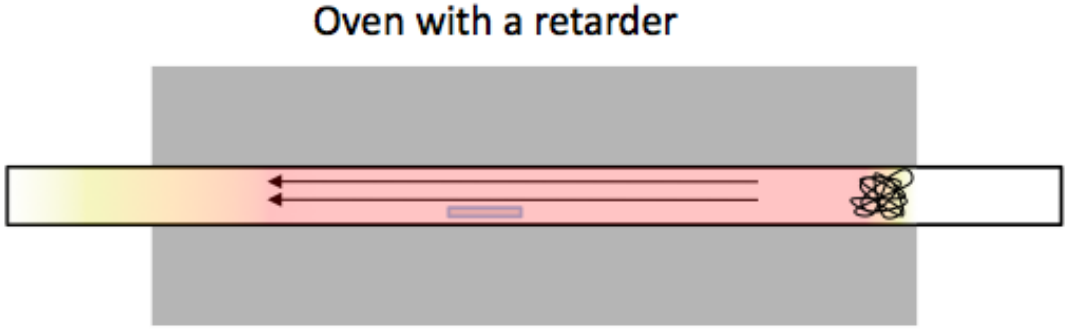


Figure 4: placing a partial obstruction in the inflow end of the oven ensures that any gases entering the oven are heated fully before reaching the sample.

### 3.1.6 ssDNA application

When applying ssDNA to samples, there are three principal concerns. First is to ensure that the DNA has time to find the carbon surface. Second is to prevent droplet drying -- a drying droplet leaves a thick inhomogeneous scum of DNA and buffer. Third is to remove the fluid cleanly.

The lab standard technique in this situation is to make a humid environment enclosed under a large petri dish by placing a small petri dish of hot water in it. The DNA solution is applied to the chip and placed in this humid environment for approximately forty minutes. Under these conditions, it does not dry. When done, we blow the sample dry with dry nitrogen.

When doing this with nitride membranes, it is important to blow across the sample so as not to blow out the nitride windows.

The sequence we applied is 5′ CTT CTG TCT TGA TGT TTG TCA AAC 3′; in the gas sensor papers[42,89], it is called ‘sequence 2’.

### 3.1.7 Transmission Electron Microscopy

Transmission Electron Microscopy was developed through the 1930s by a group led by Max Knoll[29]. It is analogous to standard transmission optical microscopy, using electrons instead of photons. The principal advantage of this approach over optical microscopy is that electrons have much shorter deBroglie wavelengths than visible light – an 80 kV electron has a wavelength of 0.04 angstroms, compared to 400 nm for violet light. This reduces the diffraction limit to a much smaller scale. In order to match its precision with electromagnetic radiation, short X-rays need to be used. These are so penetrating that forming a diffraction-limited image with them is impractical, both on account of the difficulty of making optics for them, and how little they would be affected by a nano-scale sample. Therefore, X-ray microscopes cannot match electron microscopes for resolving power. Electron optics, on the other hand, are made of inhomogeneous electric or magnetic fields; these cannot be penetrated in a comparable manner. The difficulty of weak coupling is alleviated as electrons interact more strongly with matter than photons.

The particular instrument used for this work is a JEOL 2010 TEM operating at 80 kV. Following is a description of its beam path, which is typical for TEM; figure 5 illustrates this. The beam is generated and accelerated at the top of the beam column by a hot cathode electron gun. The following lenses are inhomogeneous magnetic fields. The first condenser lens controls the brightness of the beam by spreading it across the condenser aperture. The second condenser lens concentrates the beam on the sample, focusing the beam down to the region of interest. After passing through the sample, the electron beam is magnified by two chained objective lenses, which produce a greatly amplified image at a specific image plane. The projector lens refocuses this image onto a distant point so that it will be similarly focused at the various positions of the detectors used to capture the image, minimizing the need to refocus for each one.

Additional coils control stigmation and beam deflection.

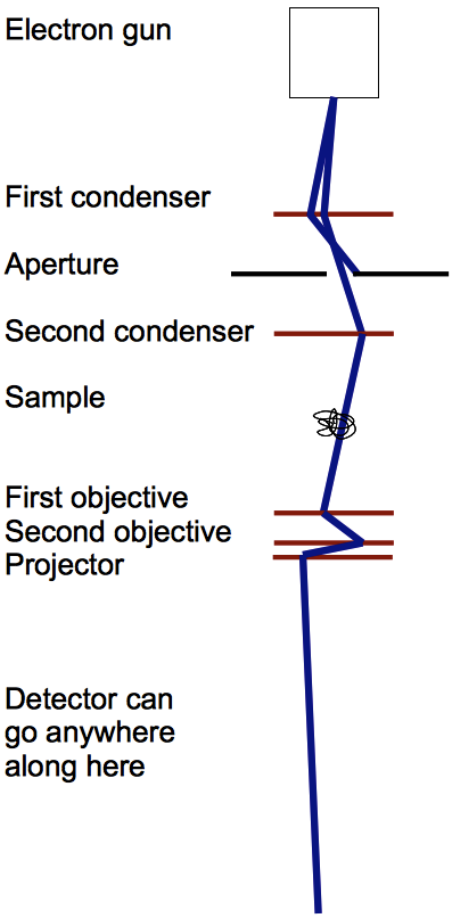


Figure 5: The TEM beam path, with two electron paths shown, one of which is blocked by the condenser aperture. Horizontal lines other than the aperture are magnetic lenses (redonline).

For very thin samples such as the carbon nanotubes examined in this work, the scattering and absorption are very slight, and would be insufficient to glean any useful information. Fortunately, thin samples do effect a significant phase change on the penetrating beam. If the image is perfectly well focused, this phase information is lost. If however the image is even slightly out of focus, then some reference electrons passing by the structure are mixed with those passing through the feature, and the phase difference between these produces intensity contrast the detector picks up [31]. Figure 6 shows the contrast enhancement from this effect.

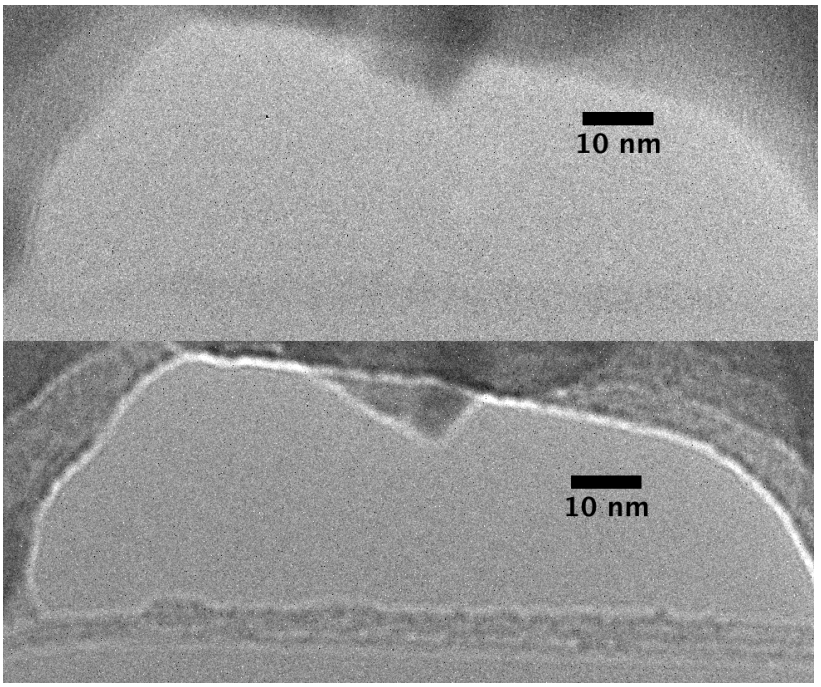


Figure 6: At top, a bundle of carbon filaments, likely nanotubes, in near-perfect Gaussian focus; it is barely visible. Below, the same view focused slightly below the sample.

If the image is more than slightly out of focus, multiple fringes will form as more of the diffraction pattern resulting from the edge is revealed. The best images will be focused well enough to have only a single sharp fringe.

### 3.2 X-ray reflection of ssDNA on graphite

The other experiment concerning the structure of ssDNA on graphitic structures was X-ray reflection of ssDNA adsorbed on graphite. This most directly models the ssDNA-graphene sensors previously described. We are interested in any structural information we can get about this system. X-ray reflection can provide plane-averaged electron density information, particularly being sensitive to the tallest frequently-occurring features. This would indicate whether the ssDNA has many excursions from the surface, stacking on itself rather than lying flat.

The experimental procedures were required to fulfil the following requirements:

The graphite must be solid and as co-oriented as possible to best reflect X-rays into the detector. It must be clean to minimize spurious signals, and of course it must fit into the X-ray system.

For best comparison to the graphene sensors, the ssDNA should be applied in a manner as similar as possible to the application method for those sensors. That process follows these principles: First is to ensure that the DNA has time to find the carbon surface. Second is to prevent droplet drying -- a drying droplet leaves a thick inhomogeneous scum of DNA and buffer. Third is that the fluid must be removed cleanly.

We need control runs to compare to the ssDNA: bare graphite, and graphite which has been wetted with just water.

X-rays must be produced in sufficient quantity to permit looking at large angles of reflectivity. The X-rays must be sufficiently monochromatic that the results are meaningful.

Air scatters X-rays to some extent; the background from this must be minimized.

The steps to satisfy thes requirements are:

- Cut a block of Highly Oriented Pyrolytic Graphite (HOPG) down to length
- Cleave it to reveal a fresh crystal plane
- Load it into the system and perform the X-ray reflectivity measurements
- Apply water to the graphite as if it were ssDNA
- Repeat the measurement step
- Apply the ssDNA
- Repeat the measurement step

### 3.2.1 Graphite block preparation

We started with a highly co-oriented block of HOPG. The measure of this co-orientation is the 'mosaic angle', the mean angle between two far apart surface planes. We used graphite with a mosaic angle around 0.4 degrees, provided by Prof. Jack Fischer; this is very good quality for graphite.

The graphite block can be cut with a sandblaster. The area that is to be kept must be protected by metal sheets clamped onto the graphite block. The clamping must be gentle and even in order to avoid disrupting the block. Shearing and bending forces are especially disruptive – in any direction. We cut the block to reduce its length so it would fit in the sample chamber.

Cleavage of graphite is not a cutting process as the name implies – instead, it is achieved by gently and evenly pressing scotch tape to a block of graphite, and slowly pulling it off. This takes off many layers at once, leaving a clean middle layer exposed. It also often leaves a ragged flap of graphite behind.

Ideally, this flap should be peeled back and entirely removed. This however runs the risk of ripping the surface and starting another ragged flap. If the flap is near the edge, facing such that one would need to peel back across the whole block, it may be left if and only if it is tight to the graphite block across the bulk of its length. If it has begun to peel away, it must be removed entirely so no cavities are left.

### 3.2.2 ssDNA application

The technique for applying ssDNA is similar to that mentioned before for nanotubes on windows. We make a humid environment enclosed under a large petri dish by placing a small petri dish of hot water in it. The DNA solution is applied to the chip in small droplets lying next to each other. Because graphite is highly hydrophobic, these droplets will be very tall and mobile. The graphite is left this humid environment for approximately forty minutes. Under these conditions, it does not dry. When done, we blow the sample dry with dry nitrogen. Then we repeat, applying a second set of small droplets in between the first set of locations, to improve coverage.

### 3.2.3 Sample mounting

The sample rod of the Triple-Axis Diffractometer System (TADS) was inserted, then the sample clamped on. Normally samples are clamped on and then the sample rod inserted; however, even the cut sample was too wide to go up the sample rod tube, so we clamped it on with the rod in place. If this can be avoided by using a normal-sized sample, do so; this operation is extremely awkward. Because the graphite block was thinner than normal samples, we used a custom clamp that took up the extra space. This was less secure than the usual washer-clamping.

In order to minimize the air-scattering, our measurements were made in a sealed humidified helium environment. This may not have been important due to the action of the analyzer assembly, described later.

### 3.2.4 X-ray reflectivity experiment

This section describes the X-ray reflectivity experiment, divided into the physical layout and the actions performed. The analysis is described in chapter 4, section 1.

#### *3.2.4.1 Design of X-ray reflectivity system*

The schematic of a TADS system is shown in figure 7.

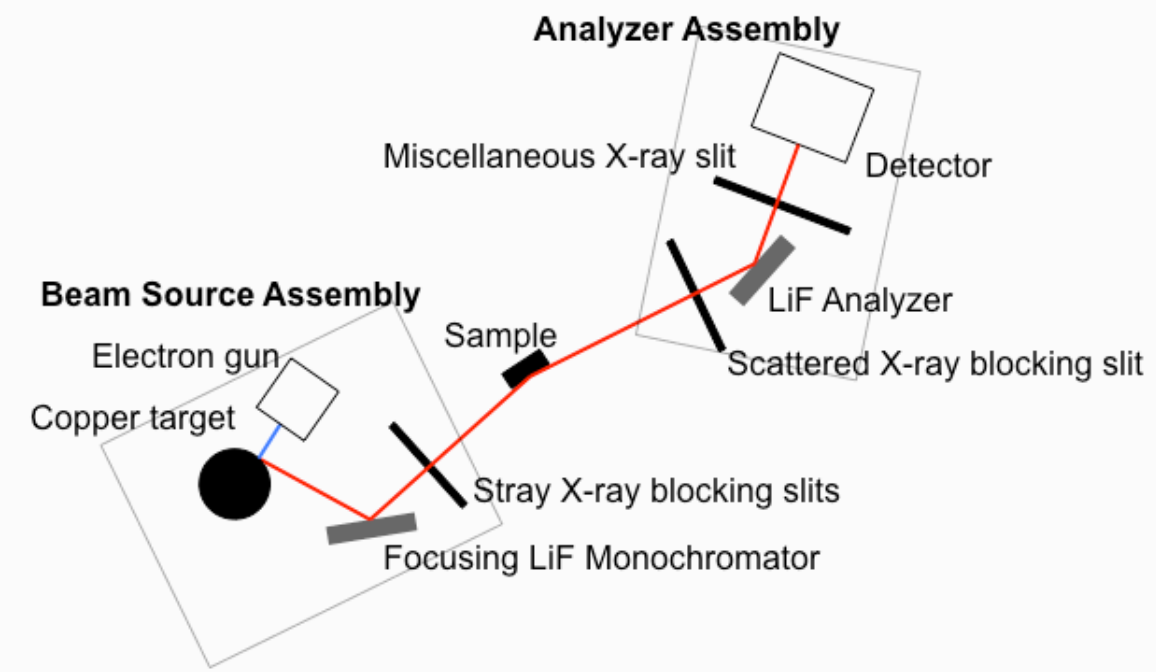


Figure 7: schematic of the X-ray reflectivity system.

The diffractometer beam source is a Bruker-Nonius FR591 rotating anode generator. This consists of a 0.2 by 2 mm wide, 40 keV, 85 mA electron beam striking a copper cylinder. The cylinder spins to spread out the irradiation, preventing its overheating. This produces primarily X-rays at the copper K-alpha emission line. The X-rays from this source are collimated in the horizontal transverse direction by passage through a slit.

A curved Lithium Fluoride reflector is placed just after this initial slit. The curve focuses the beam vertically.

LiF crystals are highly transparent to X-rays compared to other crystalline solids, which makes them suitable for X-ray optics. Their transparency helps prevent attenuation more than it impedes forming a reflection. Based on the plane separation, a peak of reflectivity occurs at one specific angle for each wavelength. Other wavelengths are not reflected, but instead pass through the reflector and are lost. This reflector is placed to reflect X-rays with an energy of 8.1 keV. Due to this process of selecting out one energy to reflect, this reflector is called a monochromator. The bandwidth permitted by LiF is intermediate between two other common monochromator materials, graphite (more bandwidth) and germanium (improved resolution). Lowering the bandwidth improves the precision of the measurement; however, this causes a loss of signal strength. This instrument normally uses LiF as a good medium. Due to the lack of sharp features in our results, graphite would have been a reasonable alternative, increasing the signal at large angles.

To recap: the beam has been collimated, focused, and monochromated. There is nonetheless some leakage of stray X-rays at other angles; these are excluded by a second slit.

The sample is placed in this beam and aligned to a known angle of incidence. These angles range from a grazing-incidence 0.07 degrees up to 20 degrees. Each angle probes a different Fourier component of the graphite-air interface (how it does so is described in the following section).

The detector aperture is placed so as to catch the specularly reflected beam. The detector has the reverse arrangement of the beam source: the beam is passed through an aperture, then after some distance reflected off a LiF crystal and through another aperture, into the NaI scintillation detector. This LiF crystal is flat as no focusing is needed. It is called an analyzer as it helps restrict the detected X-rays to those of interest.

This is achieved so: most spurious X-rays originate from a location not along the intended beam path. Those making it through the first slit must be going the wrong direction to be reflected from the monochromator, which instead absorbs them. A final slit catches a few stray X-rays, mainly diffuse scattering from edges, or from the air inside the analyzer.

#### *3.2.4.2 X-ray system operation*

The operation of the TADS is now described.

When using grazing incidence, the beam intensity reaching the detector is very high. Therefore, for low angles, we attenuate the beam just outside the beam source with calibrated aluminum foils. Each foil cuts the beam intensity by roughly a factor of 3; because the exact attenuation factor depends on the angle of the foil, we determine it after the fact by repeating a few data points before and after removing the foils.

Then the TADS is aligned for the sample. The two principal alignments done at experiment time are to find the angle the sample is lying in ($\varphi$), and find the angle it will reflect into ($2\theta$). Small errors can occur in these, producing distortions in the curve. These two alignments are performed at low angle – 0 for the first, and the critical angle, 1.6°, for the second.

For each angle to be measured, the sample is placed in the beam at that angle, and the detector at twice that angle. This catches the specular reflection. The number of incident

X-rays are counted over a period that varies from angle to angle, depending on the count rate. Low angles, as noted, have a very fast count rate, so the integration time is left short. Even so, these end up with the most counts. At large angles, the reflectivity drops off so much that the signal drops below the background level. Around this point, the integration time to yield useful data increases to unreasonable levels.

## 3.3 Metal nanoparticles on few-layer graphene

The experiments on metal nanoparticles on few-layer graphene were designed to determine the dependence of equilibrium particle size on graphene thickness and nanoparticle metal properties.

The graphene substrate must be clean and have known thicknesses spanning a wide range. This requires using AFM, which in turn requires that the graphene be readily located optically.

The metal nanoparticles must be dense enough to provide good statistics, yet well separated so each particle could be considered independently. They must also be in their thermal equilibrium configuration.

They must be imaged at a resolution sufficient to produce a reasonably accurate estimate of size.

The steps of the experiment were:

- Mechanically exfoliate graphene and apply it to a Silicon wafer
- Anneal to remove residue
- Optically identify graphene flakes
- Determine graphene thicknesses in Atomic Force Microscope
- Evaporate a thin metal film, typically 0.3 nm, of the desired metal
- Anneal to equilibrium
- Optionally, measure particle height in Atomic Force Microscope
- Image in SEM

### 3.3.1 Substrate

Start with a ~1 $cm^2$ silicon chip with a top layer of 300 nm of SiO. Decide on a canonical orientation for this chip, marking the back-side appropriately with a diamond scribe if required to maintain this convention.

This chip must be extremely clean. Wafers are shipped sufficiently clean. However, cutting the chips from the wafer will tend to introduce dust. A cleaning protocol is described in appendix 1.

### 3.3.2 Mechanical exfoliation

Mechanical exfoliation of graphene places several pieces of micron-scale flakes of high quality graphene of various thickness at random locations on a chip, interspersed among many much thicker flakes. Presently, this is the only way to make truly high quality graphene flakes; however, its clear limitations are spurring the search for another method such as CVD or epitaxial growth.

Mechanical exfoliation starts by placing a single small piece of kish graphite* on scotch tape. The tape is repeatedly folded and unfolded to press the flake in various places. Each time it is unfolded, the flake is split. This should be repeated such that near-continuous coverage is obtained (figure 8).

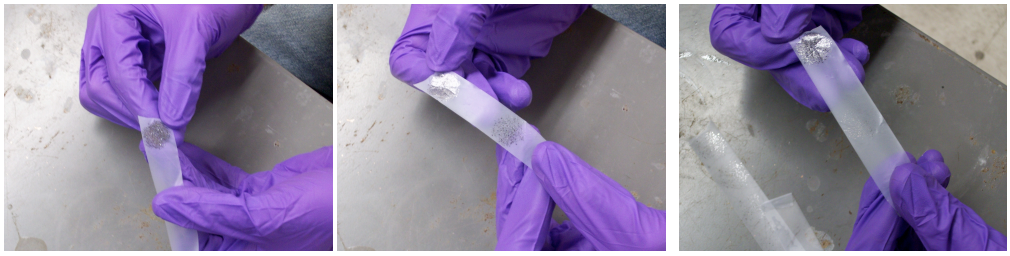

Figure 8: A (left): the graphite flake, after being split many times by pressing against tape, thus providing good coverage in that area. B (center): The second-generation transfer onto a clean tape surface for thinning. C (right): the thinned second generation is in the center of this tape. It is not so easy to see as it used to be. In particular, any parts of it remaining at all shiny are way too thick. The tape used to thin the second generation can be seen at left; the graphene pulled off onto it is successively thinner at each use (the upper left was the first press, and proceeding down out of the frame)

This thick layer on the tape is pressed to a fresh section of tape, creating a second-generation transfer onto that tape. This second generation does not include much if any graphene from the surface layers of the original graphite stock, which helps ensure the cleanliness of the final product. (Figure 8B) The first generation graphite smudge can produce multiple second generation transfers. It is important to check that coverage is adequately high on each transfer; later second generation transfers are more likely to have inadequate coverage.

This second generation transfer is thinned by repeatedly applying clean sections of tape and withdrawing them. When the surface of the second generation has entirely lost its shine, it is nearly ready for transfer to the final substrate. Thin this non-shiny surface 2 or 3 more times. (Figure 8C)

The second generation is then pressed smoothly onto the surface that is to bear the graphene. These should be gently pressed and rubbed together for some time. Rubbing too hard will cause folding, which is usually undesirable.

Then the tape is to be peeled back very slowly. Fast peeling takes graphene with it and leaves more residue. One can gauge the speed to pull back by watching the tape adhesive relax. If the adhesive snaps in the middle, the speed is likely too fast.

Once this is done, graphene and graphite flakes will be left on the sample, along with tape scum.

### 3.3.3 First anneal

By elevating the temperature in a reducing environment, volatiles such as tape scum can be removed, while leaving graphene behind. The sample is placed in a quartz tube in a cool oven set up as for CVD, described above. Ar and $H_2$ are set to flow at 300 sccm. The system is flushed for 10 minutes, then heat is applied to reach 400°C. After 1 hour,

* A similar procedure can also be performed with Highly Oriented Pyrolitic Graphite, omitting the first spreading-around step; however, the resulting graphene tends to be of lower quality.

the tape scum should be gone, and the heat may be released. $H_2$ must continue to flow until the temperature has reached 100 C, at which point the sample can be removed. The temperature of halting $H_2$ flow differs from the corresponding step in the recipe for growing carbon nanotubes due to graphene having edges which oxygen can attack more readily.

### 3.3.4 Optical identification

In order to optically identify the graphene, specific multilayer surfaces must be chosen; the typical selection is a Si wafer with 300 nm of SiO. Graphene can be optically identified on these. It appears to be a light violet against the background. Thicker graphene multilayers will be progressively darker purples, then blueish green, then yellow for thick flakes (figure 9). The graphene flakes will typically be several microns across or smaller, so high magnification must be used. Often, some graphene will be detectable before annealing, as the tape scum does not obscure all areas. The only reason to do this search before annealing away the tape is if there is doubt whether any graphene was produced; cataloguing should wait until the scum has been removed so all pieces can be found.

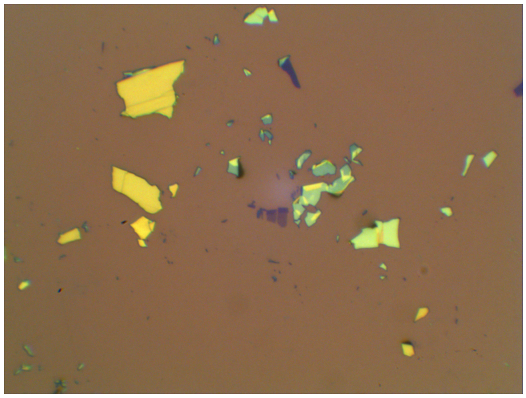

Figure 9: The dim flakes in the center are 2 layer except at their folded tops. 2 layer graphene can be difficult to visually distinguish from single layer graphene. Also visible are thicker flakes showing the whole thickness progression of lavender - violet – aqua – yellow (color online).

The location of a flake was stored by photographing the flake at high and low magnifications, then taking additional low magnification pictures to show the path to the nearest edge of the chip. The location along the edge may be noted. With the chip orientation known, locating a flake is quick. Metal alignment marks were generally not used as they frequently sustained damage from the scotch tape, rendering them worse than useless for our purposes.

### 3.3.5 Thickness determination

The thickness of the graphene is determined by Atomic Force Microscope (AFM). In order to understand how AFM works, some background is in order.

Atomic Force Microscopy (AFM) is a ‘scanning probe’ technique for surface characterization, meaning that a sharp tip on a cantilever – the probe – is moved to a sequence of positions on or above the surface – the scan – and makes a measurement at each one. Scanning probe techniques developed rapidly starting in the 1980s, with Scanning Tunnelling Microscopy appearing in 1981[12], the first AFM techniques being

developed by 1985[11], with significant developments continuing even to the present. The basis of these continuing developments is the wide array of possible interactions testable by a scanning probe. In AFM, the measurement is deflection of the probe by surface forces, measured by the subsequent deflection of a laser reflected off the back of the probe.

In general, the surface forces come in two types: attractive forces which get stronger gradually as the tip approaches the surface, and repulsive forces which are zero until the tip is in hard contact with the surface and then get very strong within a short distance. The attractive forces will typically be van der Waals attractions, but can also include fluid surface tension and electromagnetic forces, depending on tip and substrate properties. The dominant repulsive force, if contact is made, will be the contact force; repulsive variants of the other two forces can also contribute, again depending on tip and sample properties.

AFM can proceed in several ways – contact mode, in which the tip is dragged directly over the surface; noncontact mode, in which the tip oscillates in space above the surface and is deflected by attractive forces; and tapping mode, in which the tip oscillations periodically bring it in contact with the surface. Tapping Mode is the technique employed in this work; it is a trademark of Digital Instruments, which built the instruments used – a Dimension 3000 and a Dimension 3100.

The cantilever is vibrated at or slightly below its sharpest vertically vibrating resonance, with a fixed power applied. If the tip comes in contact with the surface, the amplitude drops as it bounces off the surface above where it would turn around otherwise. This amplitude is monitored by a laser directed onto the reflective back of the cantilever; the laser beam is deflected in proportion to the tip deflection. The laser is directed onto two distant photodiodes such that the laser spot straddles the two; as the laser beam is deflected up and down, the balance of light impinging on each changes. For small deflections, the voltage difference between the two is linearly related to the tip deflection.

This photodiode signal is fed through analog electronics to extract the wave characteristics such as amplitude and phase. The amplitude is the input to a feedback control system operating the tip height controls. If the amplitude is below the setpoint, the tip is withdrawn; if the amplitude rises above the setpoint, the tip is lowered. The whole loop is shown in figure 10.

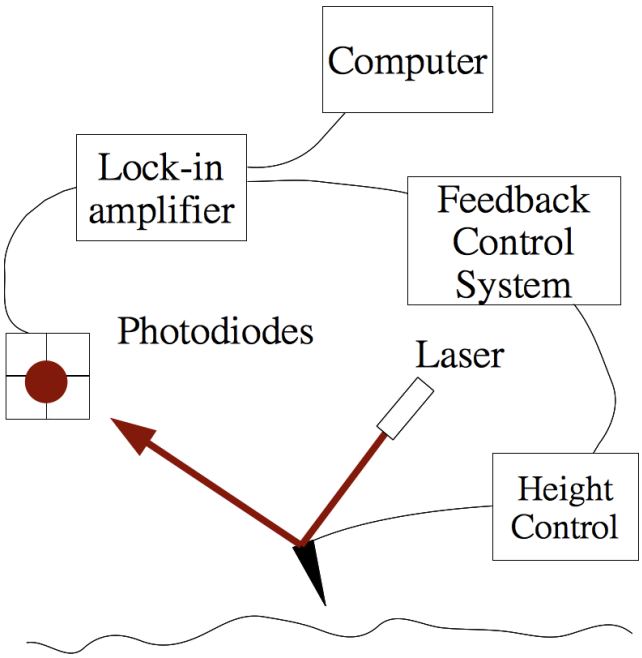


Figure 10: Schematic of the AFM system, presently at an unrealistically large depressed angle.

This tip assembly height control motor is the one feature of the system that is persistently calibrated: it does not depend on tip characteristics that can change from session to session. Therefore it is this feedback response that is the source of the regular 'height' data channel. Amplitude and phase channels are also available, and show information about the surface slope and roughness (figure 11). The scales of these data channels depend on details of the tip and cantilever, so in regular AFM imaging applications they are typically only used for qualitative characterization. They are in particular useful for avoiding many of the artifacts often plaguing the height data. However, high-frequency noise impacts the amplitude and especially the phase channels more strongly than the height channel.

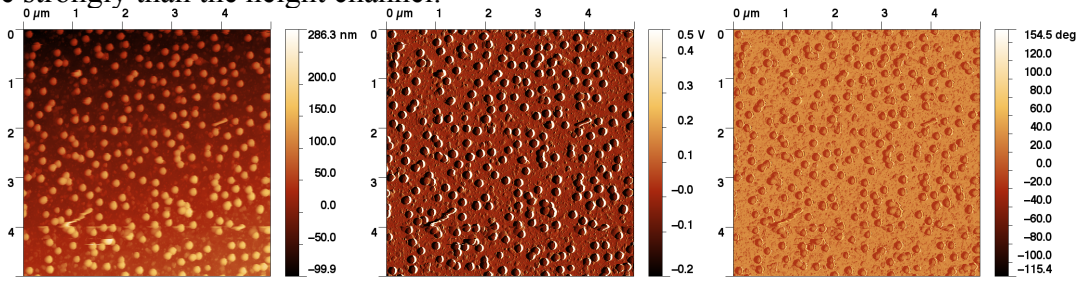


Figure 11: Height, amplitude, and phase channels, unprocessed. The shown amplitude scale is the difference from the amplitude setpoint. Images courtesy of Raghuman Govindan Karunakaran

Data samples are taken periodically; ideally, they are not taken more frequently than several relaxation times of the feedback loop so that the data actually represent the underlying substrate. This requirement places a limit on the scanning rate. On rough samples and for large-area scans, a much harder limit is encountered that during fast scans, the feedback loop may not withdraw the tip from an obstacle quickly enough, and the tip can be damaged or pick up debris. These factors limit the tip speed to around 10

micrometers per second under normal circumstances. A single image of 512 rows of 512 pixels per row, taken on a 10 micrometer field, will take around 17 minutes.

Appendix 1 contains practical information for taking and interpreting AFM images.

### 3.3.6 Metal evaporation

Now that the graphene flakes have been identified, we evaporate a controlled amount of metal onto the surface. The metal evaporator has three essential parts: a vacuum system, a large current source, and a crystal thickness monitor. The current source runs current through a 'boat' in which rests a lump of the metal to be evaporated. The Joule heating of the boat in turn heats the metal until it has a nontrivial vapor pressure. The vacuum is kept low enough that the evaporating metal atoms fly mostly ballistically. All other surfaces are cool enough that the metal instantly condense upon contact.

The metal flux and accumulation are measured using a quartz microbalance crystal thickness monitor. This must be given the density of the metal being evaporated for proper calibration. A shutter blocks the metal from reaching the sample except when desired.

When depositing metals with low evaporation temperatures, such as ytterbium, special care must be taken not to contaminate the chamber. This concern is especially acute for ytterbium as toxicity data is sparse; while there is no direct evidence of toxicity, it is safest to assume that it is acutely toxic[95]. Even aside from such concerns, it would contaminate ensuing metallization processes. Therefore, for Ytterbium samples, aluminum foil was spread about to catch as much as possible of the metal not deposited on the sample.

One rough-pumps the chamber down to around 50 mTorr, then switches to the cryo-pump to lower the chamber pressure below 1 μTorr. Warm up the current source not faster than 20 amps/minute, until evaporation has begun. Allow the deposition rate to stabilize, and adjust current to achieve the desired rate.

There are two sources of error in the film thickness – error in the exposure time, and error in the deposition rate. We found that for very low evaporation rates the error in the rate could become more significant than the error in the exposure time. Therefore, we kept the evaporation rate at or above 0.05 nm/s. At this moderate rate, the exposure time for the thinnest layers was 3 seconds, which was long enough that we could ensure that the error in exposure duration was small.

Once the rate has stabilized, open the shutter and simultaneously instruct the thickness monitor to begin accumulating. When the thickness has reached the target, close the shutter. Close the high vacuum valve and allow the chamber to cool for ten minutes. Vent and remove the sample.

### 3.3.7 Second anneal

The anneal to relax the metal nanoparticles is nearly identical to the tape scum removal step. The only changes are the temperature and the duration.

For gold particles, we annealed at 400°C, either for 5 hours, or for 3 hours. Silver particles were annealed at 400°C for 3 hours. Some metals have a higher surface tension, and these we elevated the temperature: for titanium, we annealed at 800°C for 3 hours; for palladium, we annealed at 800°C for 5 hours. Ytterbium is known to form an oxide,

but the rate is not well known, so we elevated the temperature to help reduce it. This anneal was at 600°C for 5 hours on one sample, and for 4 hours on another.

### 3.3.8 Particle height measurement

Particle heights are examined in AFM. The same methods are used as before, except that the field size must be much smaller in order to properly resolve the particles. We used a field of approximately 1μm, as opposed to the 10 μm field used for measuring graphene thickness.

### 3.3.9 Particle imaging

The SEM used for imaging in this work is a high resolution SEM and Focused Ion Beam device used in SEM mode – an FEI strata db235 FIB.

The beam column in an SEM is arranged very similarly to the top half of a TEM (described above, in section 3.1.7). The first condenser is the same, determining beam brightness by spreading over an aperture. The next lens, though mechanically the same as a TEM's second condenser, plays the role of objective lens: it focuses the beam down to a point. Deflector plates raster this point across the sample. The sample blocks the beam entirely, and emits electrons and photons in response. In normal SEM operation, only the emitted electrons are detected.

There are two mechanisms producing such emitted electrons.

First, the beam electrons can be scattered directly back out of the sample. These backscattered electrons typically have a large fraction of the beam energy. This process is greatly enhanced if the sample has electrons with high kinetic energy – core electrons deep in the shells of heavy atoms. Light elements such as carbon do not backscatter significantly.

Secondly, the beam electrons can knock free electrons from the sample. These secondary electrons have much lower energy than the beam, and the process is far less suppressed for light elements than backscattering is.

The exact backscattering and secondary electron emission depends on beam energy, atomic number, and incidence angle. The strongest and most consistent feature is the suppression of backscattering at low atomic number. For example, across a variety of beam energies, solid carbon (Z=6) has roughly half the backscattering of solid aluminum (Z=13). Beyond this, the beam energy dependence becomes significant.[82]

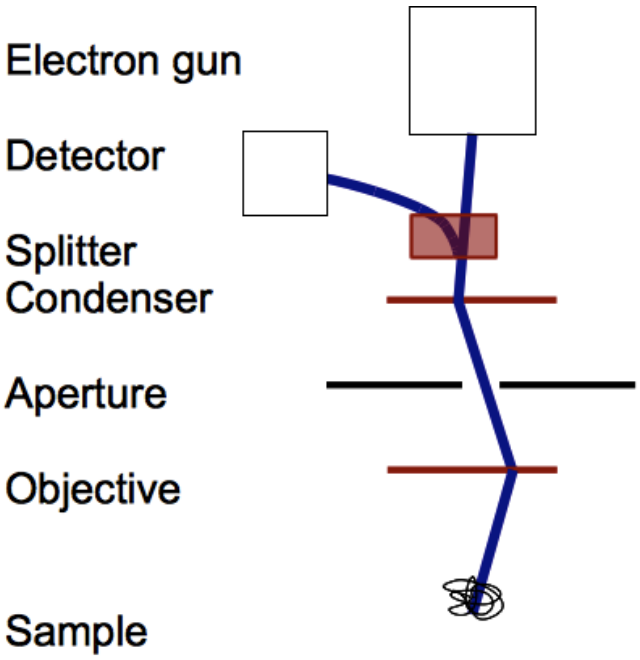


Figure 12: SEM beam path for an SEM with a through-lens detector.

Modern detectors can measure either of these two signals (or in some cases both independently) and use them to distinguish light from heavy elements. Furthermore, certain modern detectors only accept those electrons that managed to navigate their way back up through the beam optics – a 'Through Lens Detector'. In this case, electrostatic lenses are used for most beam optics so it affects both directions of beam the same; the one magnetic element separates the two just below the electron gun (figure 12). This technique rejects electrons emitted from points not in the beam spot. The secondary electrons, and, to a much lesser extent, the backscattered electrons, tend to spread out over a significant distance in the sample before being reemitted. Rejecting these less relevant electrons sharpens the image (figure 13).

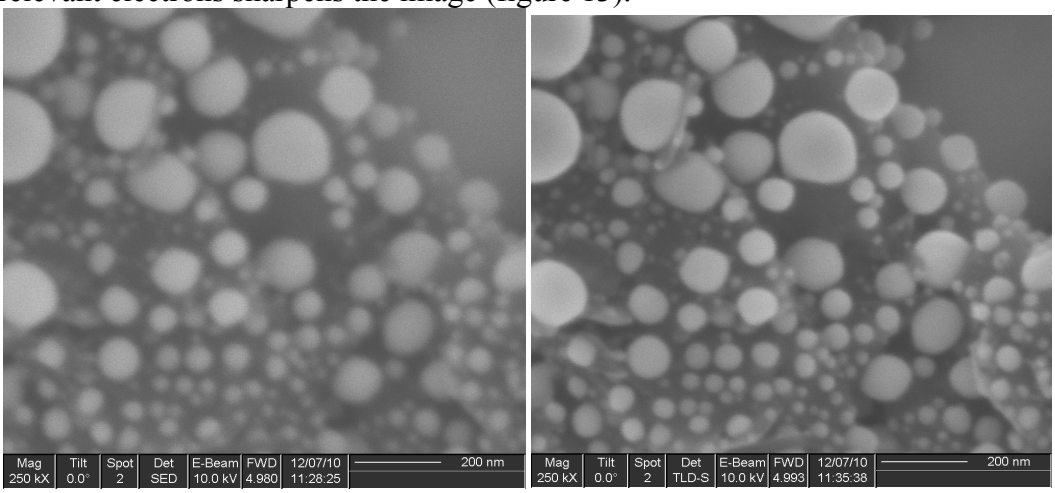


Figure 13: comparison of a sample image taken by an in-chamber SED (left) and through-lens detector (right). Image courtesy of Lolita Rotkina

## 3.4 Summary

The procedures of three experiments are described. The first involves creating ssDNA-decorated suspended single-walled Carbon Nanotubes and examining them in Transmission Electron Microscope. The second is to decorate graphite with ssDNA and take an X-ray Reflectivity profile of the surface. The third is to grow metal nanoparticles on exfoliated graphene and examine them in AFM and SEM.

# 4: Methods of Data Analysis

There are several analyses of the data presented which warrant independent explanation. The interpretation of X-ray reflectivity data breaks down into two relevant parts – normalization and interpretation. The method for extracting nanoparticle size from SEM images is described next. The simpler method of determining filament width from an SEM image with filaments is briefly described.

## 4.1 Interpretation of X-Ray reflectivity data

The X-ray reflectivity experiment in this work is part of the investigation into ssDNA's attachment to graphene and carbon nanotubes. X-ray reflectivity is an established technique for examining surfaces, and in particular thin layers of one material on another[75]. In the case where the films are perfectly homogeneous, the reflectivity measurement yields the modulus of the Fourier transform of the normal component of the mean electron density gradient. A number of ways of extracting profile information from this measurement exist, from the original sharp edge approximation, up to including accounting for roughness-broadening of features[14] and gradual transitions in medium[13].

We use graphite as a model system in order to simplify the interpretation by reducing the number of interfaces. This is most analogous to the sensing case of ssDNA on graphene, but the results may pertain to some extent to the ssDNA on nanotube system. Based on the simulations, we do not expect there to be a single conformation, but an ensemble. We are thus not expecting to find sharp peaks or valleys that could result from a highly structured system, but a distortion of the spectrum that indicates the overall length scale, and the vertical scale of any highly common features, if any.

An X-ray reflectivity experiment's raw data is a series of count rates at a variety of angles. The count rate varies depending on four factors: the portion of the incoming beam which is incident on the sample, the overall reflectivity of a flat surface at that angle, a background count rate, and surface interference – the signal. The first stage of the analysis is to isolate the surface interference portion from the others. The second stage of the analysis is to extract information about the surface from the interference signal.

### 4.1.1 Extraction of interference data

The confounding factors can in principle be calculated directly, and their effects normalized away. For angles small enough that parts of the beam are not incident on the sample, the reflected beam cross-section has a term proportional to the angle of incidence. The overall reflectivity is predicted by the Fresnel reflectivity law, and background counts can be measured separately. Additional spurious counts are also found due to air scattering and ambient background radiation. However, these calculations are not always practical: the geometric effects are sensitively dependent on alignment, and the mosaicity (non-flatness) of the sample complicates the reflectivity law.

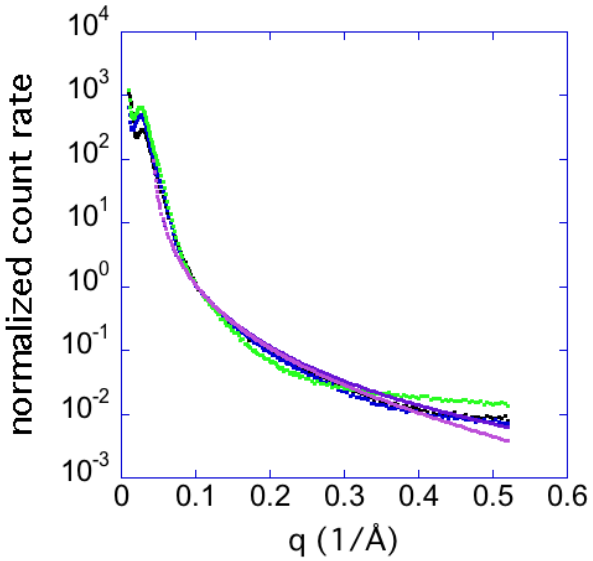


Figure 1: Actual X-ray reflectivity data (black, blue, green) and theoretical curves based on geometry and the Fresnel law (color online)

In particular, the observations drop off far less than the Fresnel reflectivity curve around 0.07/Å (figure 1). As all of the curves behave similarly in this regime, it is not due to the features of interest. We have no specific basis for adding theoretical elements to the model to account for this discrepancy. But we do have the multiple curves taken on bare and decorated graphite.

This leads to the simpler and more reliable method: divide the data taken with the ssDNA present by data from graphite only treated with water (figure 2). This naturally accounts for the mosaicity, and substantially reduces the magnitude of the geometric effects. We used the wet graphite data rather than the bare graphite data because the two curves only differed in that the wet graphite data was cleaner.

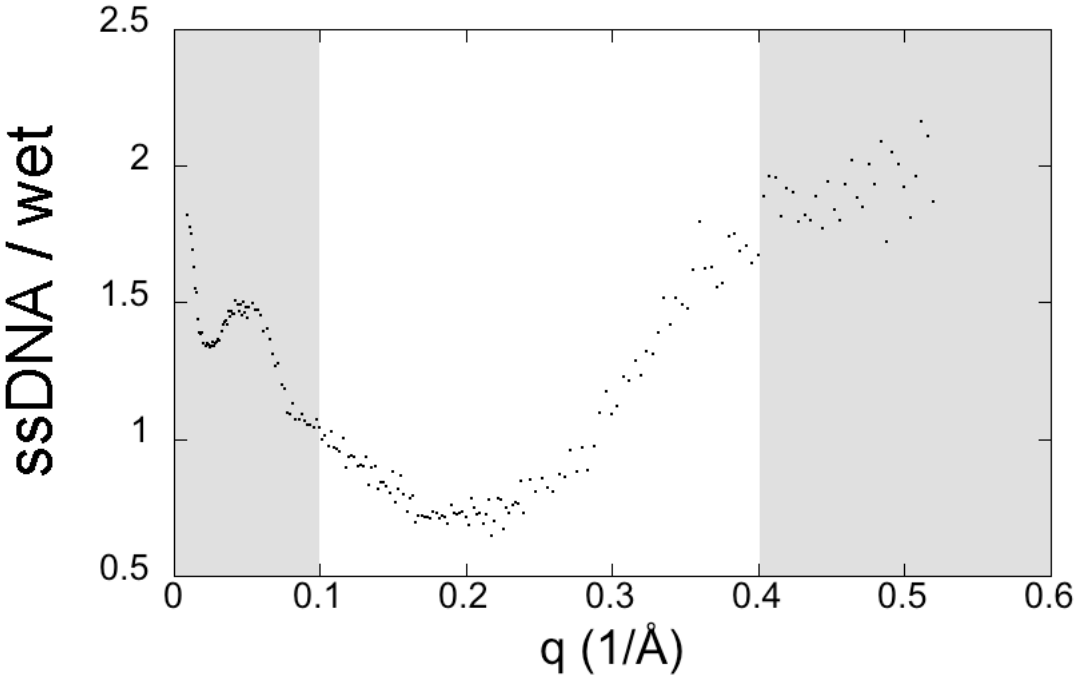


Figure 2: the reflectivity curves of ssDNA, normalized to the wet conditions.

Of this data, the two ends are less meaningful. The low-q end is dominated by geometric effects – how much of the sample catches X-rays and reflects them into the detector. This is clearly seen in the maximum around q = 0.05/Å. Therefore, we cut off the data below 0.1/Å. At high q, both terms of the ratio begin descend into the background. This also leads these measurements to be the noisiest. Therefore, we cut off the measurements above q = 0.4/Å.

### 4.1.2 Interpretation of interference data

Reflection of X-rays from surfaces differs from ordinary optics in several ways. The simplest of these factors is that the index of refraction of ordinary objects in these wavelengths is *less* than 1, leading to, at low angles, total *external* reflection[2] up to the critical angle. This is not an incidental aspect of the technique; the condition is required for access to low angles of incidence in the material's interior.

Secondly, the wavelength of X-rays – 1.5 angstroms, allows scattering off of atomic-scaled features.

From a single surface, the usual optical Fresnel reflectivity equations apply straightforwardly (the Fresnel reflectivity laws, based on continuity across a change of index of refraction, are reviewed in refs [75] and [13]). From multiple surfaces, though, the first approach that occurs is to consider multiple reflections and their interference.

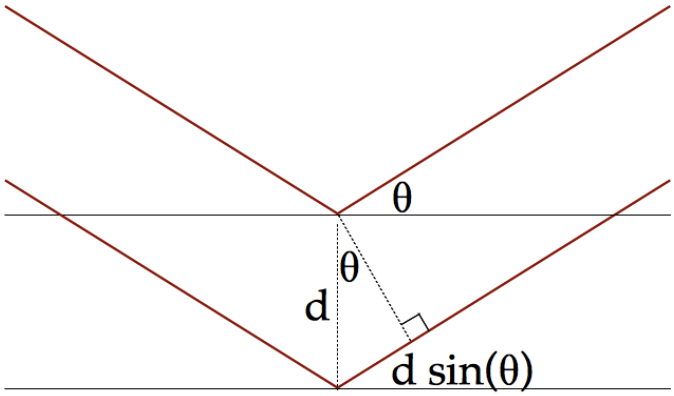

Figure 3: a ray diagram of the reflection, in the simplified case of two exact planes.

Figure 3 shows a reflection. The total path length difference is the twice the side of the right triangle shown, 2 d sin(θ). If the number of wavelengths of path length difference (n) is an integer, the reflections will interfere constructively and produce a peak. If n is a half-integer, there is maximal destructive interference resulting in a minimum. This is the central relation

$$n\lambda = 2\ d\ \sin(\theta)$$

It is convenient to combine the two beam parameters and multiply by 2π, yielding

$$2n\pi/d = 4\pi\ \sin(\theta) / \lambda = q \qquad (1)$$

This q is the difference between the incoming and outgoing wavevectors. As the angle is necessarily below a right angle, there is a 1:1 mapping between angle and q.

While this remains important, it is also useful to consider what is going on at other angles of reflection besides the interference extrema. I briefly summarize a more full treatment[2]. Rather than considering two reflections from two discrete changes of medium, we consider a continuum of medium changes, from the electron density gradient. The magnitude of each reflection is proportional to the electron density gradient at that point. The position of this reflection determines its phase to be the wavevector difference times the path length from a reference height: q d. Integrating over all possible reflecting planes yields the total reflected wave amplitude. It is formally identical to the Fourier transform of the electron density gradient. Labelling the Fourier transform of the electron density gradient $\phi(q)$, we get

$$R(q) = R_F(q)\ |\ \phi(q)\ |^2$$

With $R_F$ the Fresnel reflectivity function.

As the reflectivity assumption in this requires q to be normal to the surface, the only Fourier modes sampled are those also normal to the surface. This leaves the other two dimensions of integration in that Fourier Transform to always be simple averages.

X-ray reflectivity and diffraction setups are incapable of extracting phase information, so the exact shape of the surface cannot be reconstructed with perfect

certainty. The references given above describe strategies for working within this limitation.

Among these is the simple application of equation 1. The minimum in the normalized reflectivity with the lowest q corresponds to the lowest half-integer n (that is, ½), and the largest relevant d (that is, the film thickness). As our data only have only one minimum, this is the interpretation we will use.

### 4.2 Diameter determination of metal nanoparticles

One of the primary features of interest for a circular metal nanoparticle is its diameter. We acquired images of the nanoparticles by SEM; in these images, the nanoparticles are sometimes only a few pixels in extent. In this case, simply drawing a line across the particles is not nearly as precise as possible. Instead, we determine the area of each particle, and use that area to calculate the diameter as if the particle were perfectly circular.

This necessitates determining the extent of each particle. The essential method was to apply a brightness threshold and treat contiguous pixels above the threshold as a single particle. Several preparatory steps were required for this to be meaningful: removal large-scale inhomogeneity in the image which would favor some particles over others, remove small-scale noise which would create and destroy particles randomly, and select the region of interest. The analysis was performed in ImageJ version 1.37v[1].

First, we performed background subtraction with a radius of 5-10 pixels, depending on the image roughness – larger was preferred except when there was medium-scale background inhomogeneity. This background subtraction used the ‘parabolic’ option.

After subtraction, we smoothed the image once to reduce the number of stray pixels both in and out of the particles.

Then, all areas outside the region of interest are manually painted black, as are any obvious image artifacts. In particular, there were several vertical lines on each image produced by an electronics error in the SEM; these needed to be blotted out.

The threshold was then set to the lowest value such that there were no spurious multiple-pixel islands above the threshold.

The stages of this process are shown in figure 4.

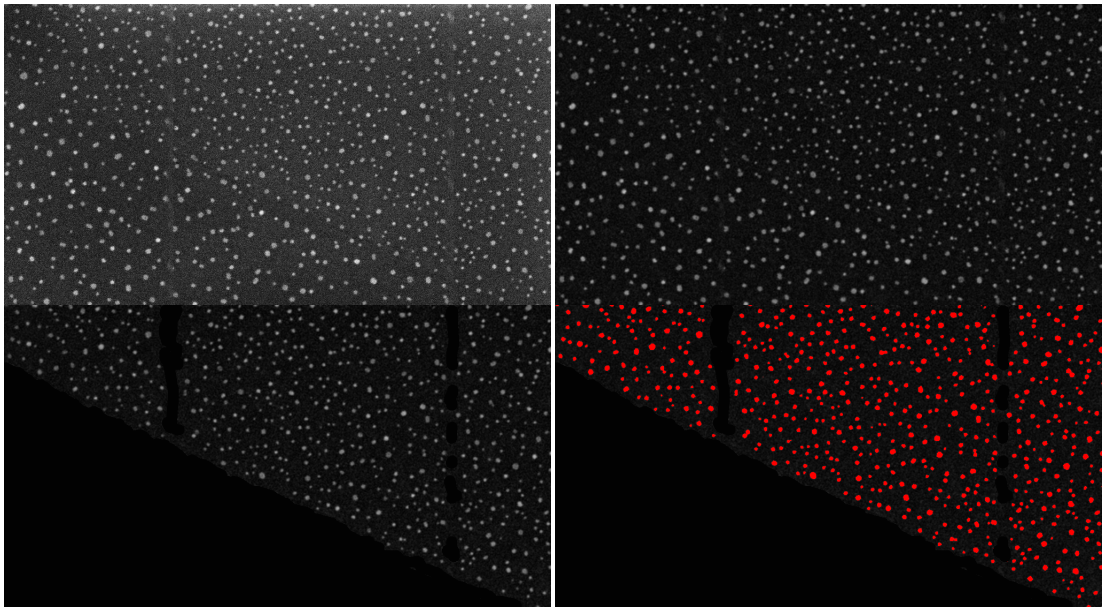

Figure 4: upper left: original unmodified image section. The field is 1506x829 nm. Upper right: background subtracted. Lower left: manual removal of non-region-of-interest and artifacts. Lower right: a threshold of 43 has been applied; contiguous pixels are taken as particles.

Particles on the edge of the image are ignored when counting, as they could extend off the side further; their size is simply unknown.

### 4.3 Determination of filament width

In some cases, metal formed into filaments. When determining the width of the filaments, no automated system could be devised due to the irregularities in the strands. A copy of the image was opened with ImageJ, and selections were manually placed as directly as possible across the filaments. The length was measured and the location marked by setting the selection to be white. The selections were placed at regular frequent intervals along several randomly selected strands, in relatively uniform regions.

# 5: TEM Measurement of ssDNA on Carbon Nanotubes

In our work with carbon nanotubes decorated with single-stranded DNA (ssDNA) [42,89], it is important to determine the physical conformation of the ssDNA with the tube. We took a multi-pronged approach, employing AFM, molecular dynamics (MD) simulations [43-45], and TEM. This chapter concerns this TEM experiment.

Transmission Electron Microscopy (TEM) is a powerful imaging technique suitable for nanoscale thin free-standing structures. Against a blank background, small clusters of atoms are visible, even for light elements such as carbon. This makes it appropriate for use on suspended carbon nanotubes; this technique was the characterization method of choice for all manner of structural studies of carbon nanotubes: demonstrating the identity of the first single-walled tubes[9,40], characterizing tubes even before their outstanding properties were recognized (figure 1)[74] and for examinations of structures involving them, such as bundles of peapods [88]. The method in itself is described in more detail in chapter 3.

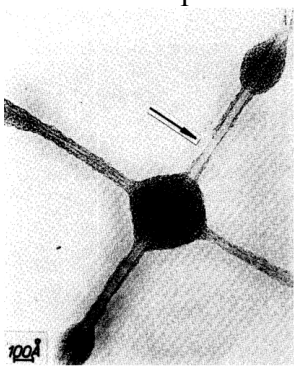

Figure 1: the first TEM image of a carbon nanotube, by Endo. This is a multi-walled tube.

Three principal TEM measurements answered questions relatively easily that would have been utterly impractical to address via MD, and were not answered definitively by AFM: determining the ssDNA distribution over the length of the tubes, the relative prevalence of multiple ssDNA strand stacking, and the behavior at tube intersections.

Even on the crucial question of the conformations of individual ssDNA strands, while MD has the advantage of knowing the positions of all atoms, TEM measurements provide a significant check.

Similar experiments have been performed on double-stranded DNA by Kappes et al [64] and Cathcart et al [17]. However, the dynamics of single-stranded DNA and double-stranded DNA are quite different. The nonpolar bases are much more exposed in a single-stranded system. The results of those experiments have no bearing on this experiment or vice versa.

Here I present the TEM results and when applicable compare them to the results of the simulations.

## 5.1 Carbon nanotubes in Transmission Electron Microscope

As described in Chapter 3, the carbon nanotubes are grown across a window in a silicon nitride membrane so no other materials are in the beam path. A typical low-magnification view is shown in figure 2.

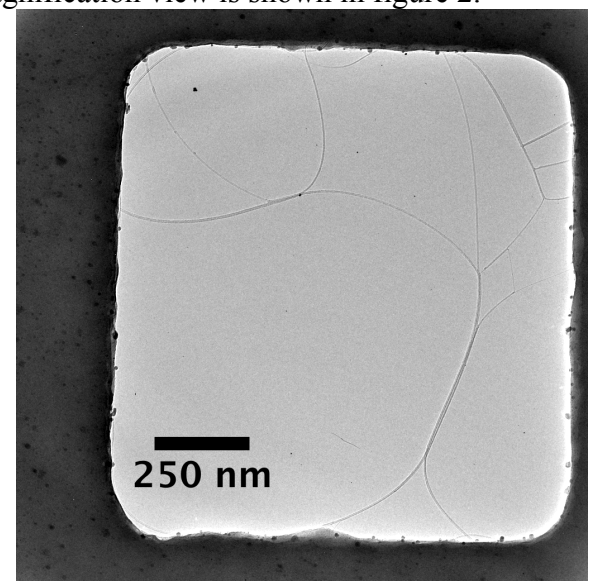


Figure 2: a carbon nanotube network typical of those used in experiments discussed in this thesis

The tubes are largely free of major scum deposits. Looking closer at the northwest section, we see that the single tube present is similarly mostly free of small-scale scum, but where tubes have bundled up, other dirt accretes around them.

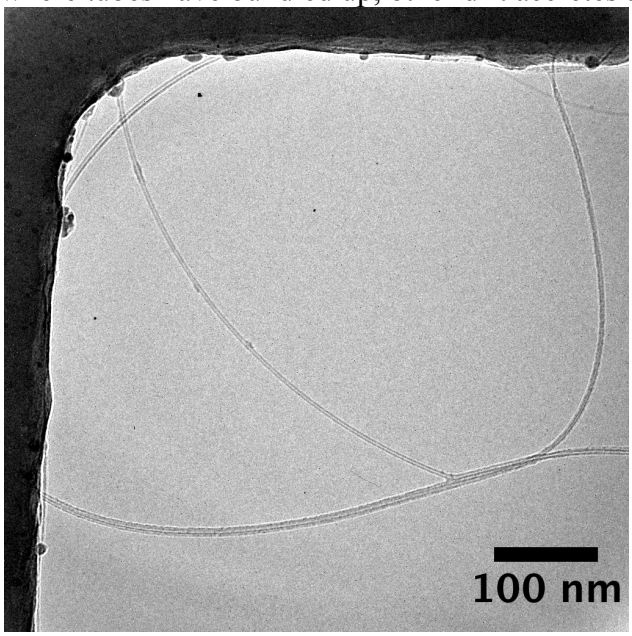


Figure 3: the northwest segment of figure 2. The tube bundles (south, east) are much dirtier than the single tubes (northwest, diagonally across center)

This is sensible as a dirt particle can be twice as coordinated between two tubes as it can be hanging off one; and as the dirt becomes concentrated along the line between the tubes, the particles can also more easily coordinate with each other.

Two alternative explanations can be discarded. The first is that having two tubes in close proximity, we are seeing an interference effect; even without going into details of diffraction from two tubes, the occasional existence of clean-looking pairs of tubes contradicts this explanation. The second is that that two tubes with different wrapping vectors line up differently at different points along the tube, like an atomic-scale Vernier pair exhibiting a 'beat' wavelength. This would be periodic, but the observed roughness is not.

Now, while there are some fragments of dirt attached to the single tube, they are quite infrequent, while if the roughness seen on pairs is due to dirt, then it is essentially continuous. Therefore, we will prefer single tubes, but in all cases require that the tube or tubes in question be primarily clean.

With good fabrication and handling techniques, as described in chapter 3, the single tubes were reliably clean before DNA application.

## 5.2 Single-stranded DNA distribution across single-walled carbon nanotubes

As described in chapter 3, we applied to these suspended tube networks single-stranded DNA, 'sequence 2' which we have been using for gas sensing since the beginning of our gas-sensing experiments[89]. As each 3 base pairs are 1 nm long, this sequence is 8 nm long.

Through this section the question of whether what we see on the tubes is actually ssDNA can be answered simply – there is enough material that the basic cleanliness of fresh tubes compels most of the material added to be from the ssDNA solution.

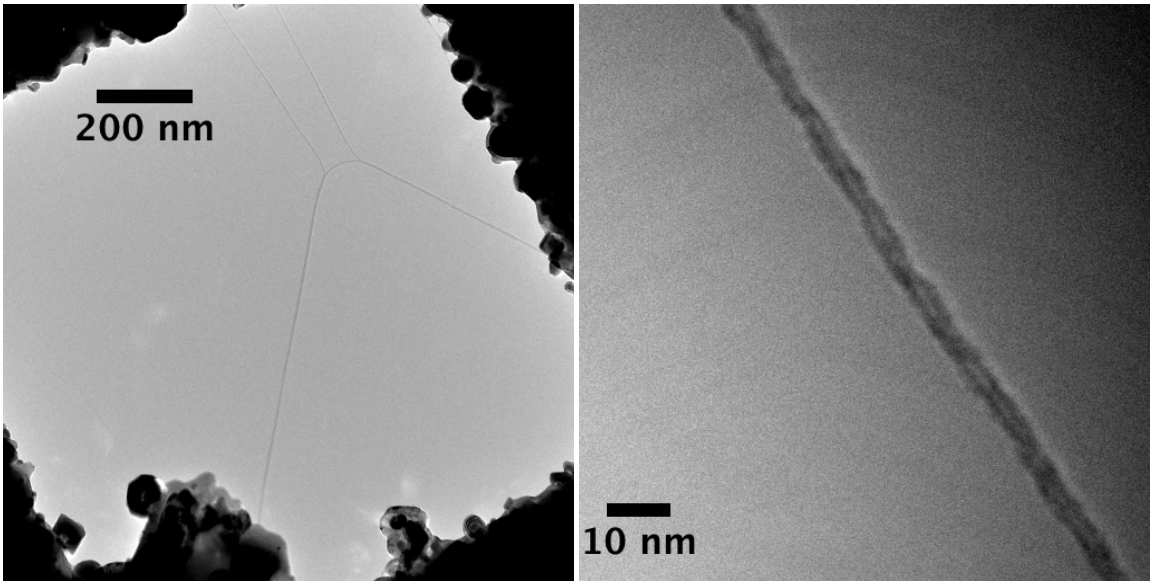


Figure 4: (a) a carbon nanotube netowork with ssDNA applied (b) a high-magnification view of the upper right strand, a single tube.

In some cases, when ssDNA was applied, it attached to the carbon nanotubes evenly from the nanometer scale up to the micron scale, as shown in figure 4.

The close view also addresses the individual strand structure, showing signs of strands crossing the tube in various ways. The precise significance of this aspect of the observation will be discussed later.

DNA did not always evenly coat the tube, as shown in figure 5.

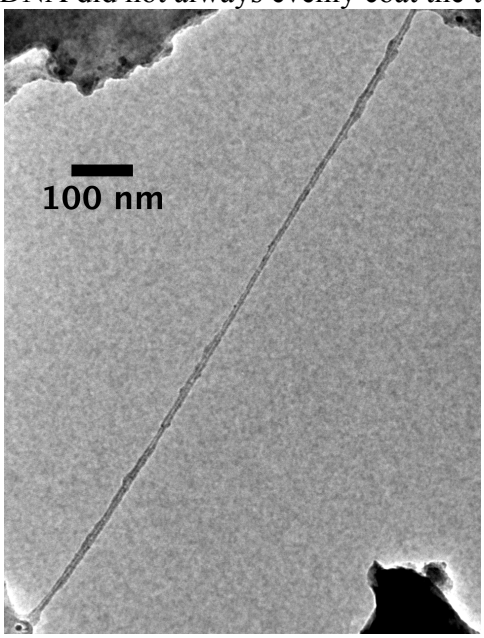


Figure 5: a single unevenly covered tube

The ends have thick layers, which thin down to bare tube in the center. Even more so, on samples where DNA is inhomogeneously distributed, we can see a massive tangle of DNA at the junction of two tubes. Similar jumbles have been seen in many cases without DNA applied, such as in figure 1. However, at junctions of the unadorned tubes in this work, the accumulations of dirt were of an entirely smaller order.

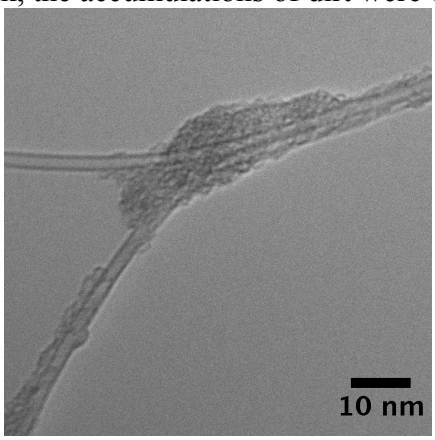


Figure 6: the junction of two tubes is frequently an especially strong accretion center.

So we see that ssDNA applied to suspended carbon nanotubes by our lab-standard technique can be evenly distributed or not, presumably depending on details of the

application. One obvious-seeming candidate explanation is that existing bundles of the ssDNA that formed while it was in storage must be broken up for ssDNA to be evenly distributed; however, those samples with large clumps seemed to have more ssDNA altogether, not simply deviating in that it was unevenly distributed.

One effect that could cause this would be insufficiently clean removal of the ssDNA solution. When preparing these suspended tube samples, there is a second requirement that conflicts with clean removal of the ssDNA solution: the dry nitrogen blast must not break the nitride membrane. This compels using a less vigorous nitrogen blast, which can at times completely fail to cleanly remove the solution. In the worst cases, this produces windows occluded with muck.

It is not a stretch to suppose that even in the cases where the result was usable that sometimes the solution was not entirely cleanly removed. This would deposit a moderately thick scum. This is an especially attractive explanation because it explains why the center of the tube in figure 5 was less covered than its ends: the meniscus of fluid would be thicker around the tube near the surface, and as it dried in place would deposit more material there.

This removal problem is less severe for gas sensing chips, as there is no significant limit on the nitrogen blast used to remove the solution.

### 5.3 Conformation of individual ssDNA strands to carbon nanotubes

As noted earlier, we have some information about wrapping behavior. There in figure 4b, it is alternating streaks across a uniformly coated tube. Their alternating direction precludes explaining them away as uncorrected astigmatism. For a single-strand example, on a different sample, we see what appears to be a single strand wrapped around the tube (figure 7).

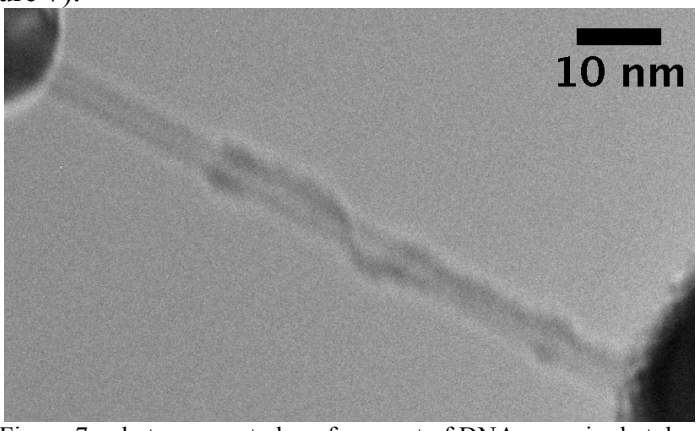


Figure 7: what appears to be a fragment of DNA on a single tube

Its length suggests that it is actually several strands next to each other, if it is DNA at all.

Pure chiral wrapping over a length great enough for it to be apparent was actually quite uncommon. Aside from the not conclusively chiral image in figure 4b, the only instance we saw at all was on a pair of tubes:

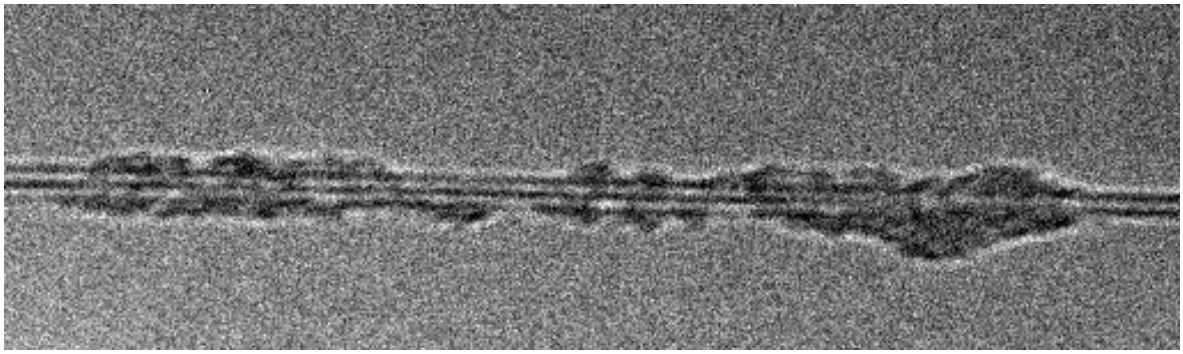

Figure 8: the middle alternating figure on this double tube appears to be chiral wrapping.

In order to gain more information about such structures, we rotated the stage in the TEM to view the same structure from a different angle. By turning the stage between images, we acquired the pair of images in figure 9.

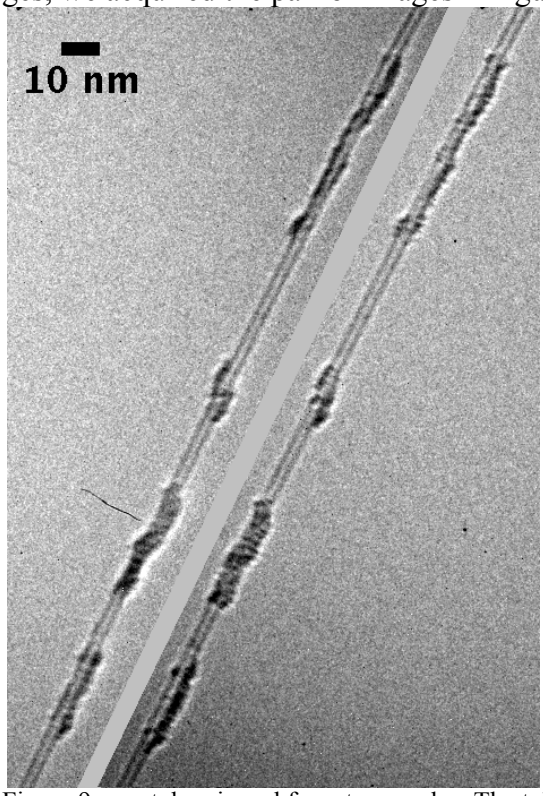


Figure 9: one tube viewed from two angles. The tube was rotated from the horizontal by 30° along the axis down the page for the left frame, and for the right frame, 30° in the opposite direction.

From this we can deduce some information about the vertical position of some of the features. In particular, the northernmost wrapping does indeed make at least ¾ of a turn with northward chirality at its northern end, and then at least ½ of a turn with southward chirality at its southern end. As this feature is 50 nm long, it cannot be one strand, but could be many strands end to end. The small feature in the center could be two strands next to each other.

## 5.4 Comparison to simulation

One principal result of the molecular dynamics simulations was an ensemble of configurations that the ssDNA would be expected to assume (Figure 10) [45].

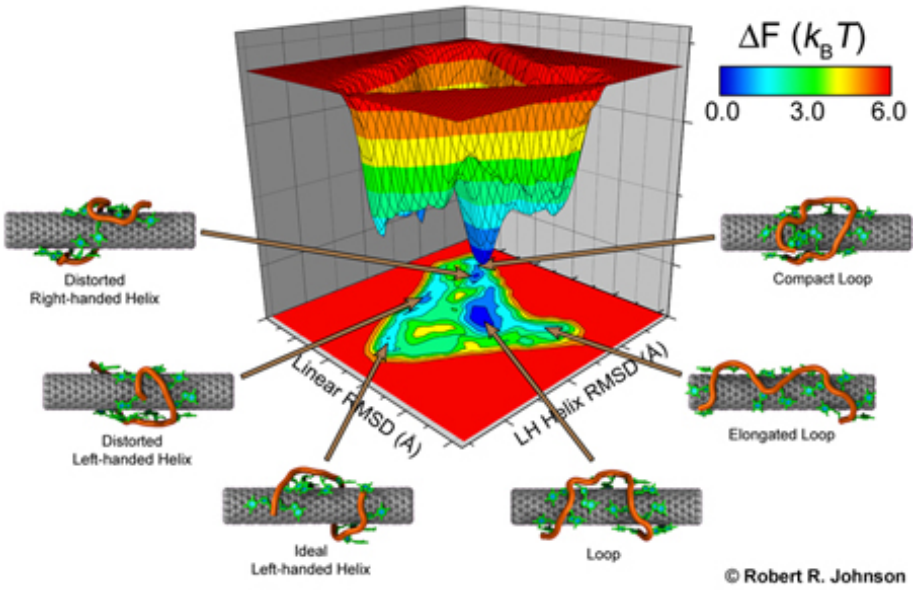


Figure 10: the free energy landscape of single-stranded DNA wrapped around a carbon nanotube

The most common arrangement according to this set of simulations is the 'Loop', which resembles a capital omega wrapped around the tube. Even among different arrangements, one common feature is that at the ends of the strand, the strand curls directly around the tube. This can be observed clearly at the ends of the strands in figure 7, and to a lesser extent in figure 9.

More specific comparisons are not readily made, except to note that strands this short do not appear to frequently form pure spiral structures on single tubes in either simulation or experiment.

The simulation gave the reason for not frequently forming as a spiral as the bases becoming crowded. ssDNA wrapped around two tubes would have a very different set of constraints, and crowding could easily not be the most significant, thus permitting the extended spiral seen above.

## 5.5 Conclusion

We have seen that after the application of ssDNA solution to suspended carbon nanotube networks, the density of string-like features along the carbon nanotubes increases from very sparse to very common, giving us some confidence that these features are in fact ssDNA. Some samples, especially those with a large amount of added material present, displayed strong clumping behavior, with clumps somewhat more frequently seen near the surface and at intersections.

The individual strands were observed to conform in a fashion consistent with the simulations: they do not generally prefer when on single tubes to spiral neatly around the tube, and at their ends they do tend to curl in more markedly.

Lastly, the frequency of strands of length greater than a single applied ssDNA strand is too great to be explained solely by random dirt specks. We can conclude that the strands tend to line up next to each other, often giving the illusion of being a single longer strand, even when all of the strands remain in direct contact with the tube. This is also consistent with the observation of continuous even coverage seen in figure 4b. This suggests that strand end to strand end interactions may be significant to the sensing behavior.

## 6: X-ray Reflectivity of ssDNA on Graphite

The final prong of the investigation into the physical structure of single-stranded DNA (ssDNA) on graphitic surfaces was to measure the X-ray reflectivity of a film of ssDNA on graphite. This system parallels most closely the ssDNA-on graphene system used for gas sensors.

We took X-ray reflectivity measurements under three conditions: bare, freshly cleaved graphite, graphite with pure water was applied as if it were DNA solution, and graphite to which sequence 1 DNA was applied. The raw results are shown in figure 1.

The data renormalized as described in chapter 4 are shown in figure 2.

The lowest-q minimum is associated with the size of the largest commonly occurring features. The sharpness or lack thereof indicates the regularity of this size. The minimum here occurs at 0.2/Å. This corresponds to a typical film thickness of 16Å.

The peak itself is very broad – so broad that it is impossible to identify the background against which a Full Width Half Maximum should be compared. This indicates great surface inhomogeneity.

These observations are consistent with the expectation of a partial monolayer of ssDNA lying approximately flat. The experiment was somewhat of a fishing expedition, checking in case consistent structures arose. Such structures would have produced a sharper first minimum at the least, and quite possibly additional maxima and minima. Their absence confirms that the ssDNA arrangement on graphene is as chaotic as we had suspected.

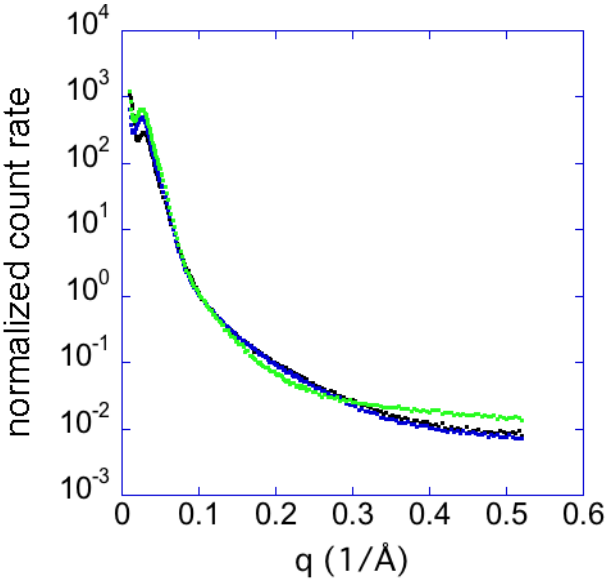


Figure 1: the normalized counts per second. Bare graphite is seen in black; it almost entirely coincides with the wetted graphite data, in blue. Green indicates graphite with DNA applied (color online)

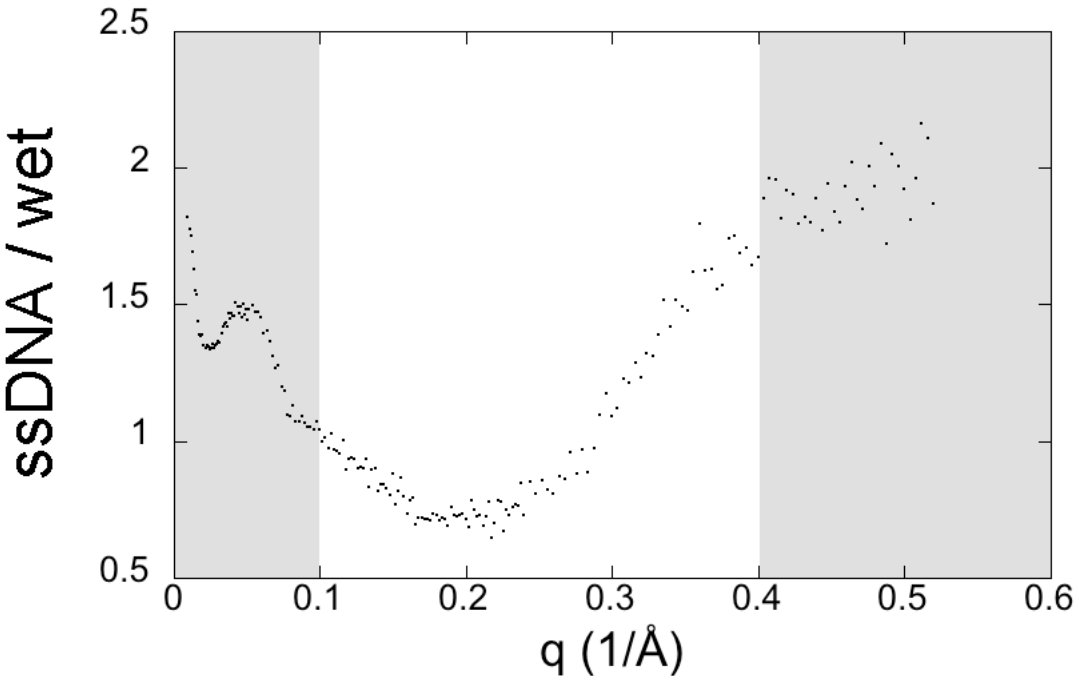


Figure 2: The ratios of the DNA and wet signals, as a function of q. The grayed-out regions are those most subject to confounding effects and noise.

# 7: Observation of Metal Nanoparticles on Graphene

Portions of the next two chapters were previously published as *Z. Luo, L.A. Somers, Y. Dan, T. Ly, N.J. Kybert, E.J. Mele and A.T. Charlie Johnson, Nanoletters 10, 777 (2010)* and *Nanoparticle Shape Selection by Repulsive Interactions: Metal Islands on Few Layer Graphenes*. L. A. Somers, N. A. Zimbovskaya, A. T. Johnson, and E. J. Mele. Physical Review B **82**, 115430, 2010.

The behavior of micro- and nano-particles on surfaces has been studied for a long time [28] and is of continuing interest [19,22,68,76,80]. This system is of general interest because the surface alters the properties of the deposited particles and vice versa. It is of practical interest for a wide variety of reasons: many conventionally made technological structures are (and more will be) nano-scale themselves, and a part of their behavior can be learned through study of these standalone nanoparticles; these particles can be anchors for chemical components; they can be catalyst centers, either for further fabrication or as part of a device; they can be used as etch masks.

In particular, on graphene, we used them to etch into the graphene itself in an effort to make graphene nanoribbons. During these attempts, a tendency for the size to vary by location was noticed. This prompted a more careful investigation, which is the subject of the remainder of the chapter. These investigations led to the formation of theories of nanoparticle growth, which are described in the next chapter.

These two theories depend on the balance between a surface tension force, and an electrostatic force created by charge transfer between the metal and the graphene substrate. The electrostatic force depends on the work function difference between the metal and the graphene, and also the thickness of the graphene. Graphene's work function is taken to be 4.66 eV, based on first principles calculations[86].

For reference, the relevant properties of surface tension at melting point and work function are shown for selected metals in figure 1.

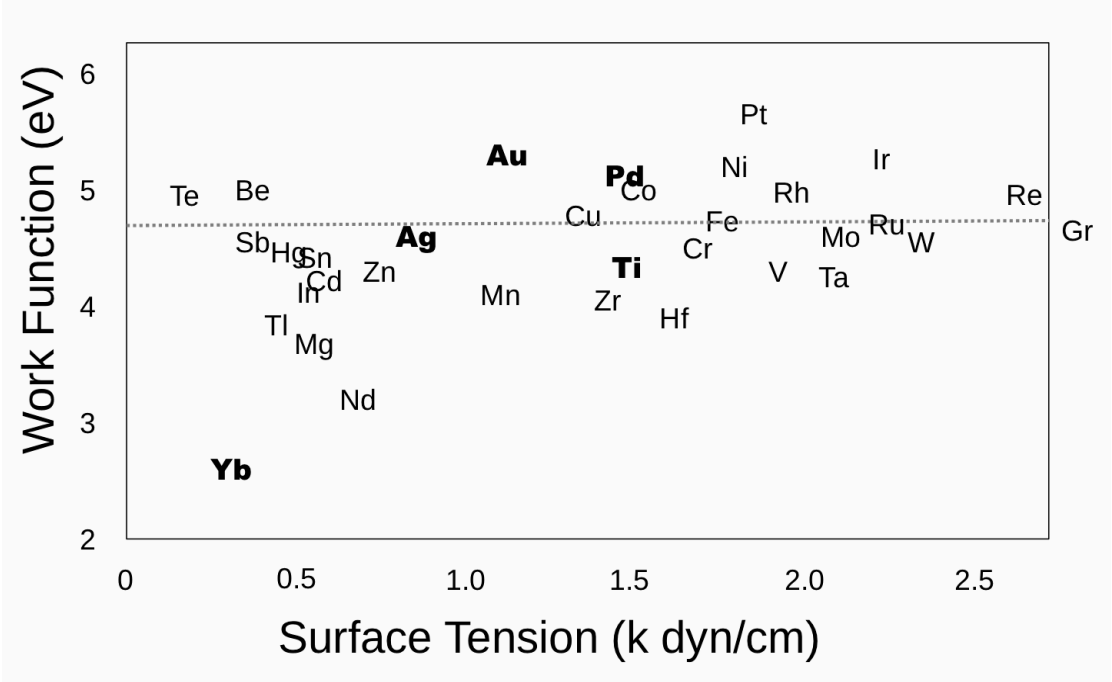


Figure 1: Materials parameters for a variety of metals; experimentally examined symbols bolded. Symbols are centered on the best known values [69], [57]. The dotted line marked Gr is the work function of graphene.

### 7.1 Gold

Having prepared the Au nanoparticle samples as described in chapter 3, we examined them in SEM. We found multiple regions with thicknesses of graphene from 1 to 3 layers, two of 6 layers, and one each of 4, 5, 8, and 13.

Figure 2 shows Au nanoparticles on several thicknesses of graphene.

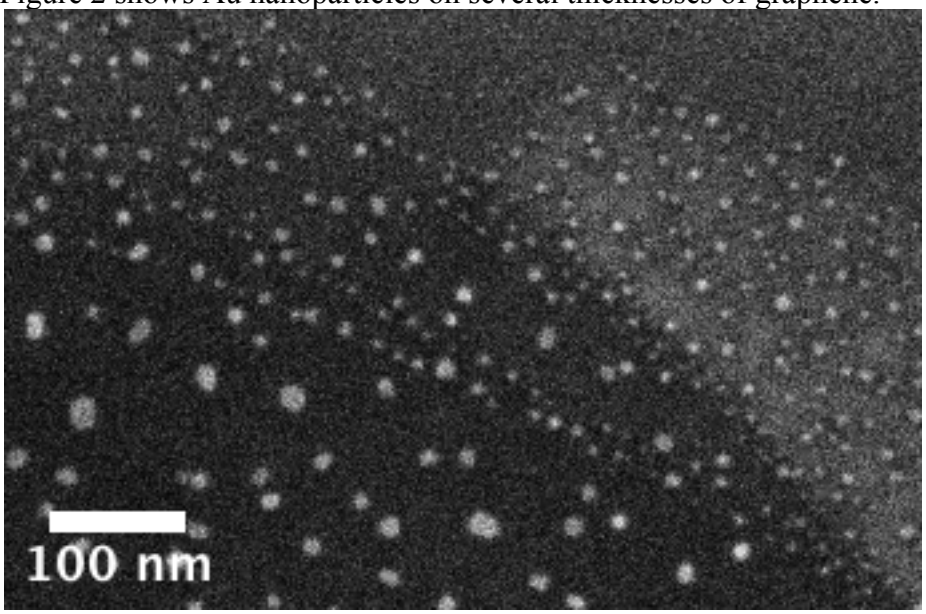


Figure 2: Au nanoparticles on 2 (right), 3 (upper left), and 13 (lower left) layer graphene. Some faceting can be seen on the 13 layer graphene.

The principal features of note are that these nanoparticles are circular, and the diameter distribution depends on the number of layers of graphene. Figure 3 shows some typical several-layer flakes

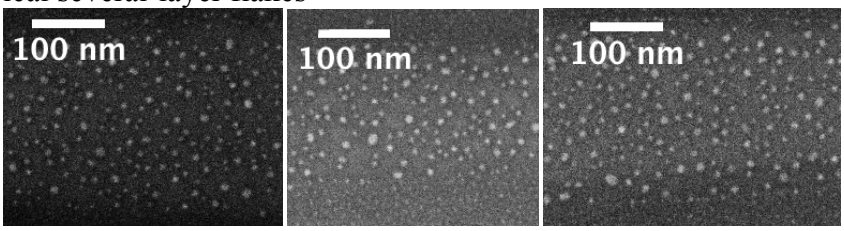


Figure 3: Circular particles on single-layer graphene (left), two-layer graphene (center), three-layer graphene (right). These are small sections of the full images used for analysis.

As seen in figure 4, on thick graphite flakes, particles are observed to form facets. These particles are sufficiently large and their size varies so broadly that there is no clear size-limiting effect aside from the availability of gold.

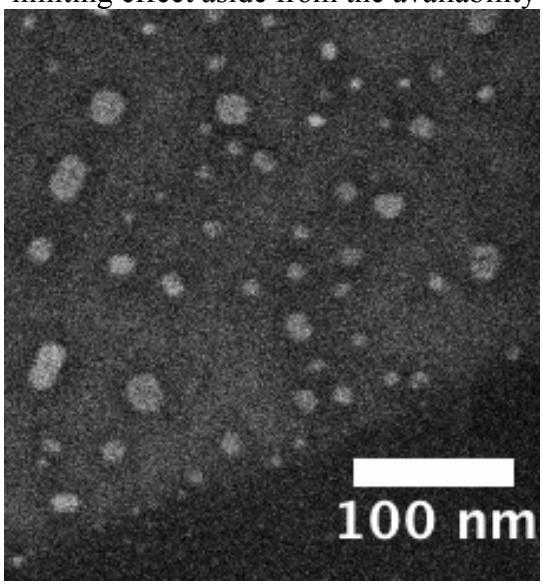


Figure 4: Au nanoparticles on a section of bulk graphite. These faceted particles fall outside the domain of the theory presented in this chapter.

In a few cases, faceting was observed on single-layer graphene (shown in figure 5).

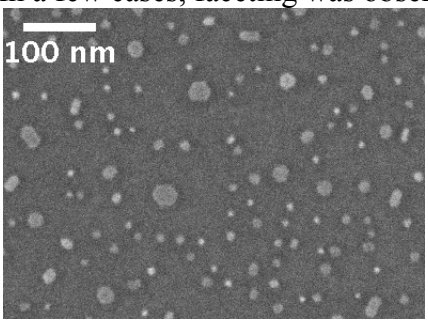


Figure 5: faceted gold particles on single-layer graphene, intermingled among round particles.

Faceting was extremely rare on intermediate thicknesses – one faceted particle each on four and five layer graphene, with all other particles from two to six layers being round. On 8 and 13-layer graphene, there were again a substantial number of faceted particles.

Figure 6 shows an AFM measurement of a region of nanoparticles also imaged in SEM. By matching up particles measured in the two, we constructed a plot comparing particle height to diameter (figure 7). We found that the height was proportional to the diameter, at a ratio of roughly 1:6. The particles on thin graphene are thus seen to be fairly flat.

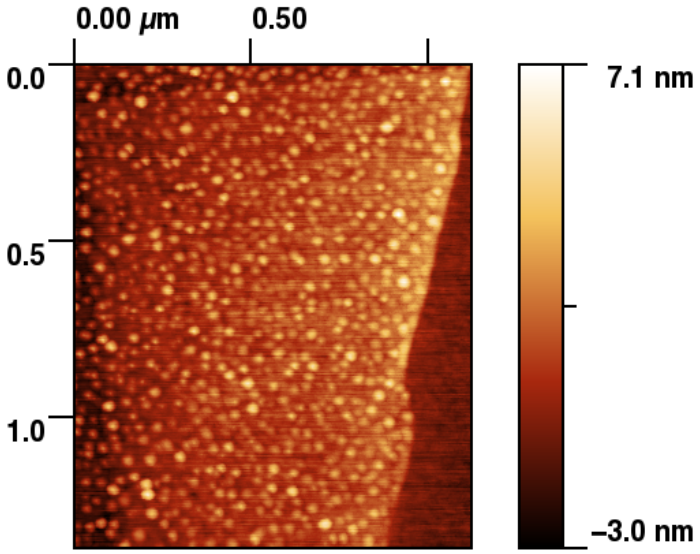


Figure 6: These particles were used to evaluate the height of the gold nanoparticles as a function of their width.

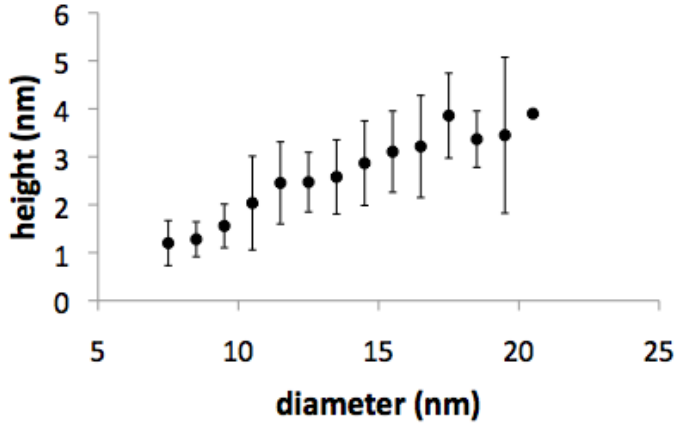


Figure 7: mean particle height as a function of diameter bin; the distributional standard deviation for each bin is shown. Diameter determined by SEM investigation of the same particles

We extracted the particle diameters using the method described in the analysis methods chapter. We considered these as histograms of particle diameter. The particles smaller than These are shown in figure 8, along with a gaussian fits to each histogram.

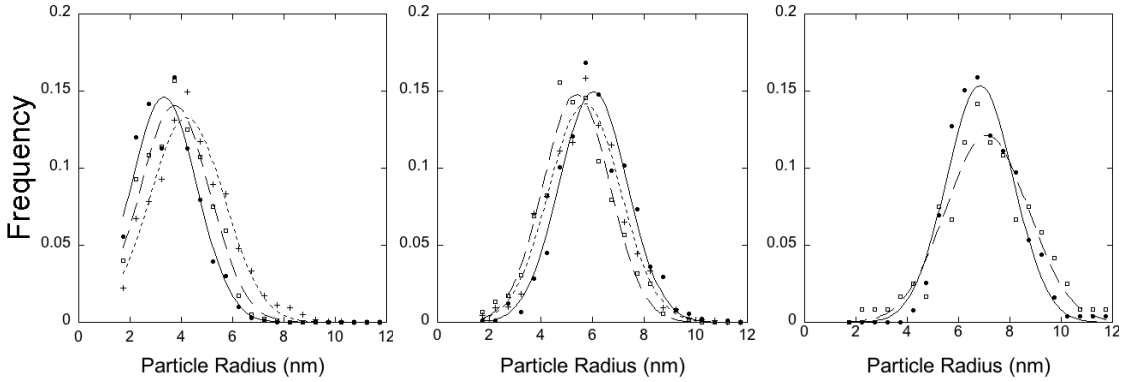


Figure 8: Radius frequency distributions, ignoring the very smallest particles, which could be spurious. 1-3 layer in left plot, 4-6 layer in center plot, 8 and 13 in right plot. In each plot, thinnest gets the solid marker and lines, next gets hollow and dashed, and next gets + and dotted. The middle plot is not in error – particles on 4 layer graphene were larger than on 5 and 6 layers.

Each of these gaussian fits has two features of interest – the radius distribution center, and the standard deviation of the particle radii.

To determine the pattern of behavior between the layers, we plot these curvefit parameters against layer count (figure 9). We use the center and width of the distribution instead of the actual mean and standard deviation of the particle sizes for two reasons. First, they are less sensitive to outliers. Second, they come with error estimates. These error estimates had to be adjusted as described in the analysis methods chapter.

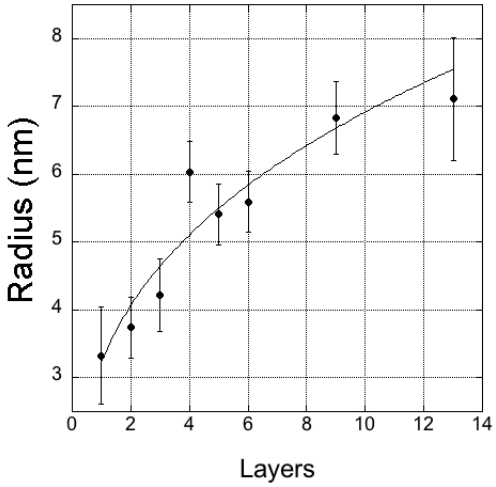


Figure 9: Radius distribution center versus layer count, fit to a power law

The power law fit is $3.23 \pm 0.34$ nm Layers$^{0.331 \pm 0.061}$. There is no procedural reason to suspect the point at 4 layers, but as an outlier it is natural to consider what happens if it is discounted. Repeating the whole fitting procedure without it, the error estimates for the other points shrink by a factor of 2.7, yielding a similar fit result is $3.00 \pm 0.13$ nm Layers$^{0.355 \pm 0.025}$. Both fits are consistent with the theory. This scenario is provided only to show that the one data point did not skew the results excessively.

### 7.1.1 Unannealed gold

One of the most striking results was for Gold that had not been annealed. While the gold on most surfaces had not separated into particles, the gold on exactly single layer graphene did separate into distinct particles (figure 10). These particles were 5.96 nm in diameter on average, with a standard deviation of 5.28 nm. This is essentially the same as the size of the annealed particles (6.65 nm), though they differ in that the standard deviation for annealed particles was merely 3.64 nm.


Figure 10: The gold on exactly one layer of graphene formed large particles; on the other layer counts, the particles are indistinct. On the substrate (top of image and similar regions), there were small distinct particles.

One last feature of note is that there are dead zones in which large particles did not form, but smaller particles did. These are principally around the edges of the single-layer graphene flake but also at other places spreading through its bulk. We do not know of an explanation for this phenomenon.

### 7.2 Silver

Various measurements of Silver's work function have produced results from 4.52 to 4.74 eV; this is nearly centered on the work function of graphite. This suggests that the electrostatic force will be very small. Silver thus comprises a negative control: if the same effects were seen for silver, we would rule out an electrostatic origin.

Silver particles were seen to be circular on thin graphene. However, the diameter distribution did not depend to a significant degree on graphene thickness. Most 1-4 layer regions had silver particles with a radius of approximately 3.3 nm. (figures 11, 12)

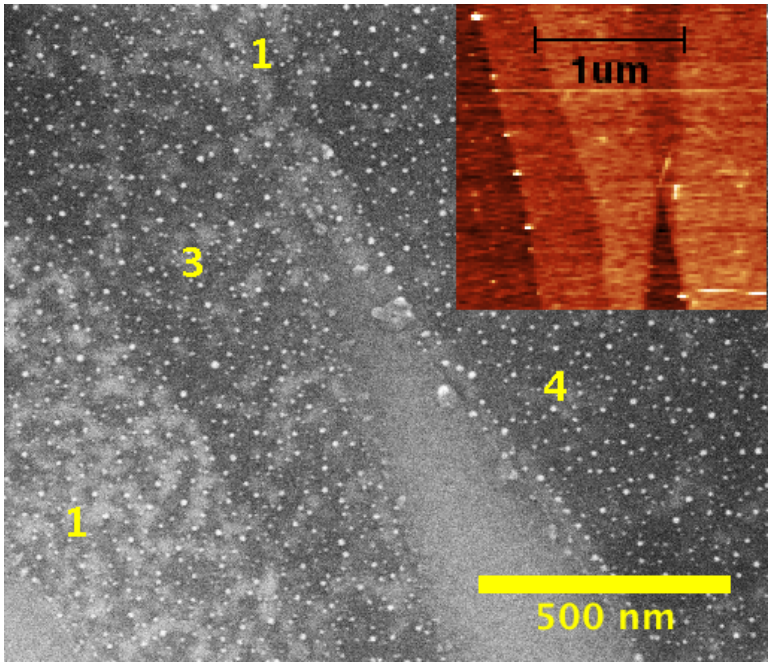


Figure 11: SEM image of silver on various thickness of graphene. Inset: AFM image of the same region prior to addition of silver.

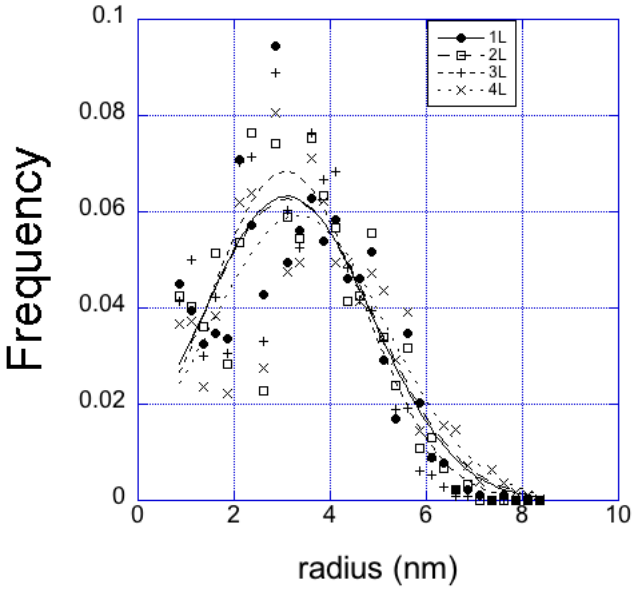


Figure 12: Histograms for silver nanoparticle radii on those regions of 1-4 layer graphene displaying constrained growth.

A 7-layer graphene flake was also identified; on this surface, the silver nanoparticles were slightly larger than these (mean radius 4.5 nm), with a similarly narrow distribution.

Two regions of thin graphene – one single layer and one two layer region – were observed to have very large particles (figure 13)

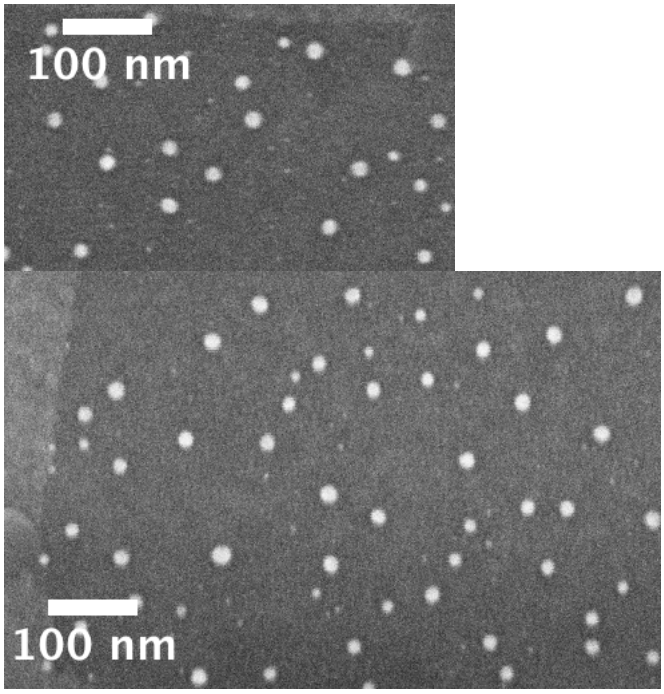


Figure 13: Silver nanoparticles which have grown to a larger size. Top: single layer. Bottom: two-layer graphene.

The distributions of sizes on these flakes are bimodal, unlike the distributions on other flakes.

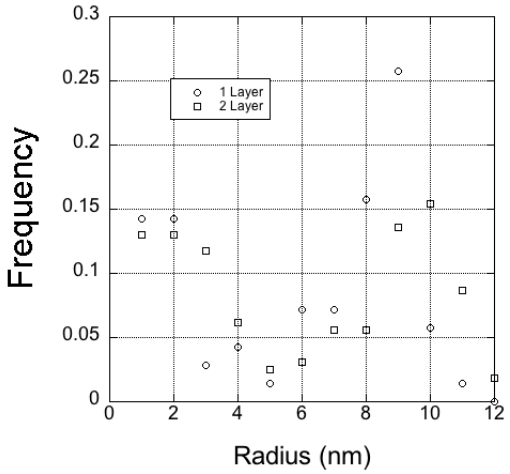


Figure 14: particle radius frequency distributions for the two regions with much larger particles.

These observations indicate that particle size is not strongly limited for silver in the way it was for gold. This supports the general notion that the limiting effect is electrostatic in nature, as the electrostatic limiting effect would be very weak for silver. Further testing with additional metals is required to show that it is not actually a gold-specific phenomenon.

### 7.3 Titanium

Titanium has a work function of 4.33 eV, for a difference of around 1/3 eV. This is smaller than the half-eV difference from gold, but significantly larger than the difference for silver.

Titanium particles formed in two ways – either small nanoparticles like gold produces (figure 15), or clusters. These clusters dominated on all flakes thicker than 4 layers, and on some thin flakes as well (figure 16).

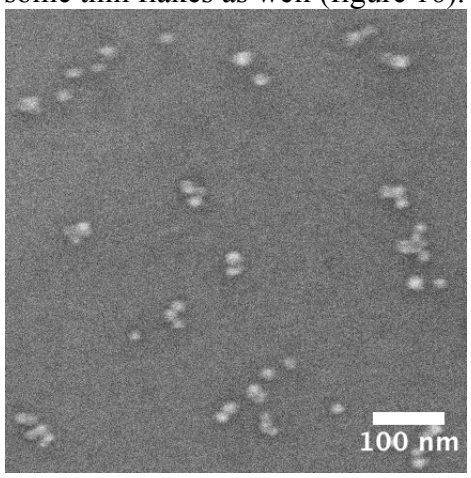


Figure 15: Titanium particles on 4 layers of graphene. Clustering is already becoming a significant effect, but it is still possible to measure the size of some individual particles.

As a result of this, we were unable to get any particle size distributions for thicknesses beyond four layers.

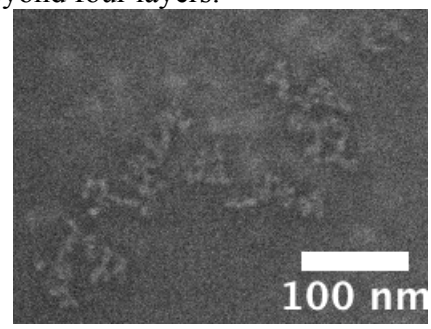


Figure 16: Ti particles on 7 layer graphene. This bundling was typical; there were no free-standing particles.

The particles were in some cases not brighter than the background in the SEM images, only being detectable by their dark haloes (figure 17). This interfered with the particle detection methods – background subtraction tended to bring the true background level up to the particle detection threshold, making the results strongly depend on particular image processing choices.

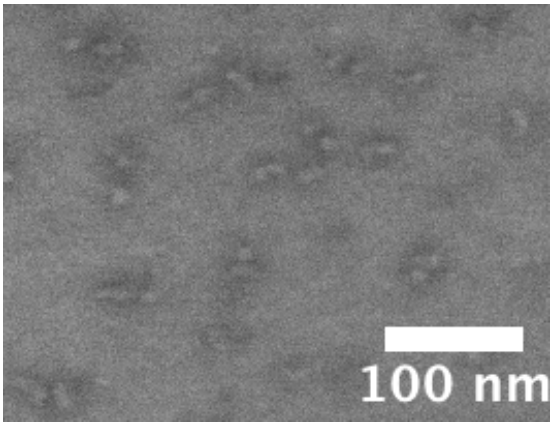


Figure 17: these Ti particles on 2 layer graphene are principally identifiable by the darkening around them. The size attributed to them is highly dependent on image processing choices.

The particle size data were compiled only for the individual particles. The distribution centers are plotted versus layer count in figure 18. The 2 layer point is most suspect, based on the halo issue and the commonness of clustering; that it is also the outlier is thus not surprising.

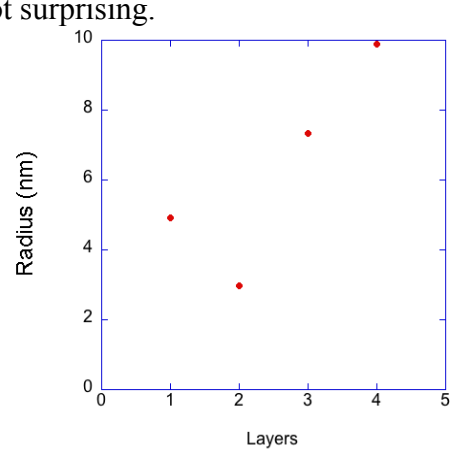


Figure 18: Mean radii of Ti nanoparticles on various thicknesses of graphene.

With only 4 layers of single particles, there is insufficient data to produce a strong determination of functional form as we did with the gold particles. Therefore we reuse the power law form. Taking reasonable error estimates based on the degree of replication and reliability (partially discounting the 2 layer data point with its shadows), the fit yields a power law of $4.5 \pm 1.0$ nm Layers$^{0.51 \pm 0.19}$.

## 7.4 Palladium

Palladium nanoparticles were seen to form a pear shape. Each particle had a fat bright end, and a dim thin end (figure 19). This phenomenon was reproduced through several experiments.

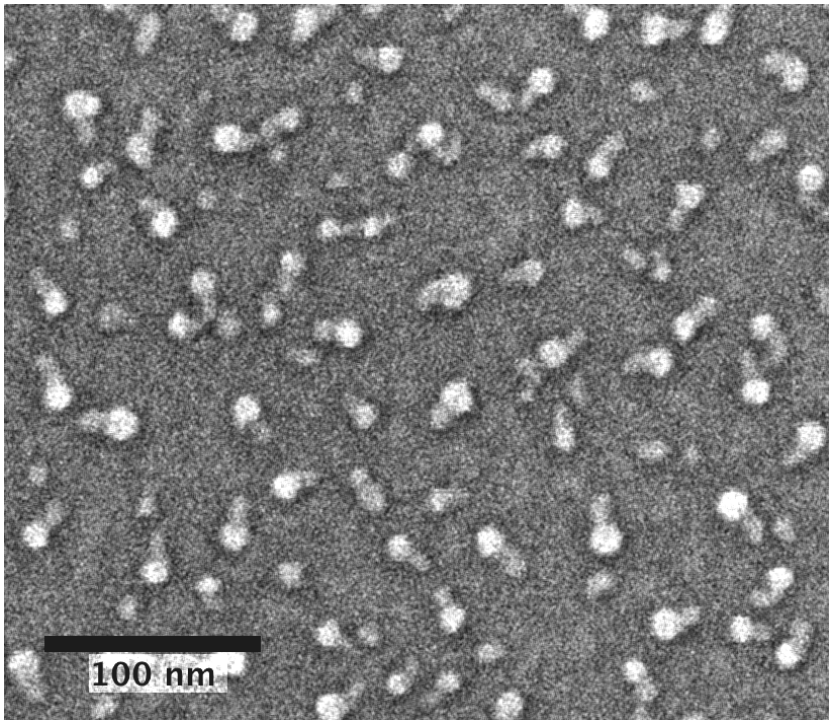

Figure 19: Palladium nanoparticles on single-layer graphene. The dimmer sections may be regions in which the Palladium etched the graphene.

Two explanations are offered: first, that our Palladium was impure, and the metals separated upon annealing; second, that the Palladium etched into the graphene and produced a short trench before getting stuck. In this case, the dim ends are not particle at all, but gaps in the graphene.

Palladium carbide is a known compound, so the latter explanation seems reasonable. I had been avoiding metals that formed carbides to avoid this effect, but did not check for Palladium. Palladium carbide has been seen to form at temperatures as low as 500 C under comparable circumstances[63], so the process temperature would have to be drastically lowered to have a chance of not producing this effect.

On the other hand, if this deformation effect was etching of the graphene, one might expect to find that the orientation is highly lattice-dependent[23]. Therefore, we also measured the angle distribution of the nanoparticle tails. This found that there was no particular pattern when we took the angle modulo $\pi/3$. This argues against the etching interpretation.

### 7.5 Ytterbium

Ytterbium has a work function difference from graphene of 2.06 eV. We deposited 0.12 nm of Yb on the surface and annealed, as described in Chapter 2. The resulting particles were not at all circular, instead forming elongated structures (figure 20).

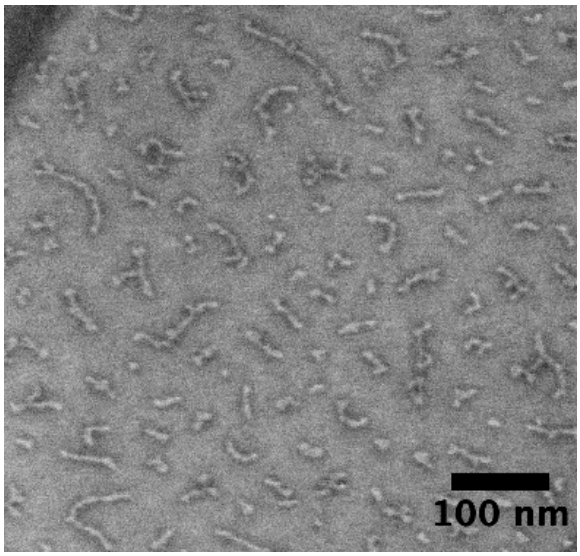

Figure 20: nanoparticles formed from 0.12 nm Yb deposition and 5 hour 600° C anneal, on 1 layer graphene

These strands are well isolated on single layer graphene, but on thicker graphene this is not reliably so (figure 21). Here is a mix of clusters of small particles, and longer chains than seen on single layer graphene.

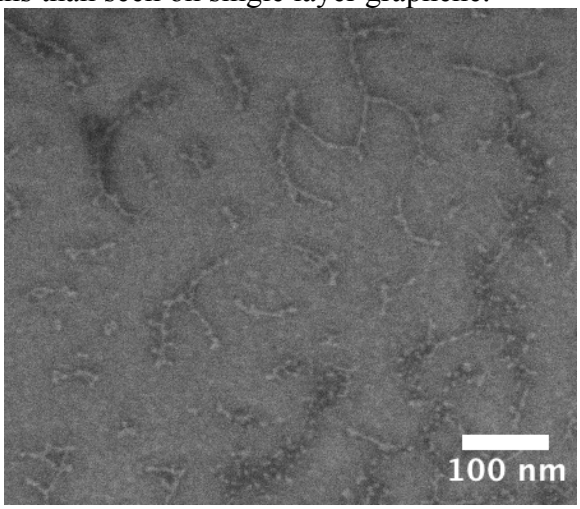

figure 21: an SEM image of Yb nanoparticles formed from 0.12 nm deposition and 5 hour 600° C anneal, on 3 layer graphene

By repeating the experiment with 0.5 nm of Yb deposited, we were able to form networks of Yb (figure 22).

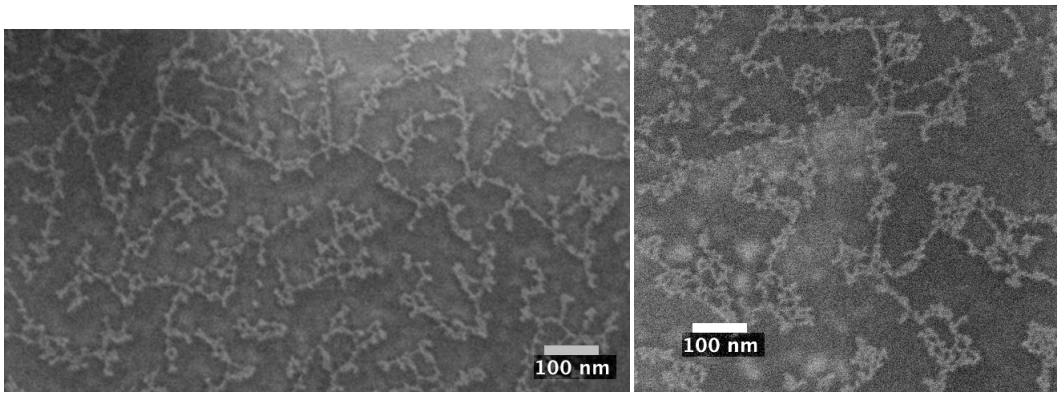


Figure 22: Yb formed from deposition of a layer 0.5 nm thick, on single layer graphene (left image), 2 layer graphene (right image, upper left), and 4 layer graphene (right image, remainder).

The networks produced by the 0.5 nm Yb experiment have many more bumps, knobs, and tight loops than the well-separated strands seen in the 0.12 nm Yb experiment. Also observed were mottled masses, seen principally on the 4 layer graphene (Figure 22, on the right side). These are nearly solid but are pierced with frequent holes.

The simple strands formed the focus of further investigation. The width of the strands was measured manually on two samples (details described in chapter 4); the width of the strands was essentially independent of the amount of Yb deposited (figure 23).

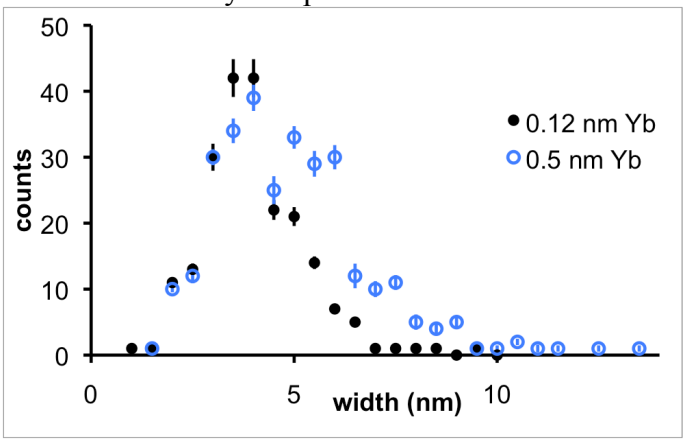


Figure 23: The width distribution of the Yb filaments for the two deposition thicknesses (color online)

The orientation of the strands is evenly distributed. We measured the particles in AFM; their heights are around 2 nm (Fig 24).

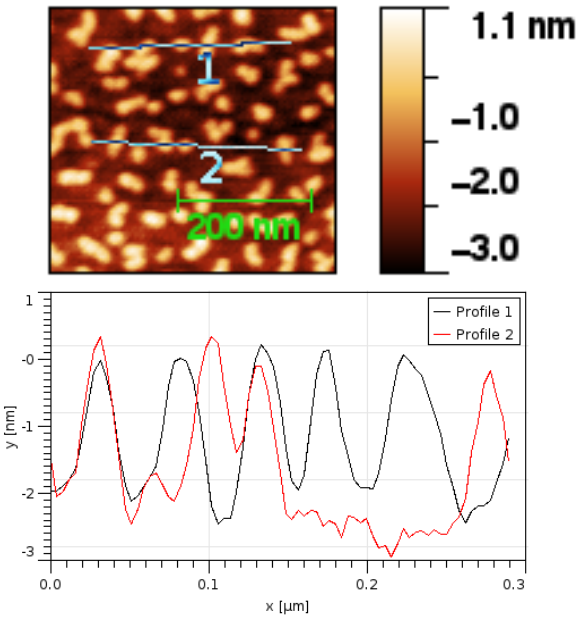


Figure 24: two sections of AFM data showing particle heights. 2 nm height is typical. (color online)

## 7.6 Summary

We prepared metal nanoparticles by evaporation and annealing, on graphene flakes of various thicknesses. Gold particles prepared in this way were primarily circular; these circular particles exhibited a narrow range of sizes. The size selected depended on the number of layers of graphene; the particle radius distribution center is given by 3.23 ± 0.34 nm layers$^{0.331 \pm 0.061}$, up to 13 layers. Similar though not as clear behavior was seen for Titanium, the distribution centers being 4.5 ± 1.0 nm Layers$^{0.51 \pm 0.19}$, up to 4 layers.

Silver had no layer dependence, but rather had regions of constrained and unconstrained growth.

Ytterbium, in contrast, formed filaments and filament networks.

# 8: Theories of Nanoparticle Growth and Deformation

Two theories are presented that account for most of the phenomena described in chapter 7. They have the same fundamental basis: the work function difference of the materials produces a dipole density along the interface. This introduces a repulsive interaction which dramatically changes the energy landscape of metal particles.

The two theories share this beginning, but go on quite differently. The first theory assumes the particles remain circular, and finds their equlibrium size. This covers the gold and, to a lesser extent, titanium particles. The second theory explores dipole-driven deformations of the nanoparticles. This provides a mechanism by which the existence of the irregular ytterbium particles can be understood.

These changes occur on the background of nanoparticle growth in the absence of any long-ranged electrostatic forces, which is generally described by Ostwald Ripening[7,58]. This framework can also be adapted to address the growth of voids in irradiated materials[91] and the stimulated growth of dendrites[33]. The most primitive of these theories (covered in the introduction of [93]) is sufficient for our purposes, as we are interested in the static final state, which is trivial in conventional Ostwald Ripening. We will describe this background first, then the common elements of the two modified theories, then each theory in detail.

## 8.1 Ostwald Ripening

The concept is to consider a quasistationary solution of the diffusion field of free particles (hereafter 'components') around larger agglomerated particles (hereafter 'particles'). It is quasistationary in that it exists on two time scales: on the short, diffusion time scale, it is a static solution of the diffusion equations; but it has a parameter, the particle radius, that is permitted to vary over longer time scales. The boundary conditions are: at infinity, the concentration of components is some bulk value; near a particle, it is inversely proportional to the particle radius.

This latter is justified by considering the chemical potential of a component on the surface of a particle,

$$\mu = \mu_0 + V_m \gamma K$$

with $\gamma$ the surface tension, K the surface curvature, and $V_m$ the volume of one mole of the components. The K is reciprocally related to the radius, yielding a $\mu$ with this dependence. The surface flux will be proportional to the difference in free energies; the chemical potential is a term of that free energy, and it contains the inverse radius term. As we are considering the quasistationary solution, any transient relaxation between the particle and the components in its neighborhood will already be finished; thus, the only net contributors to flux are change in the inverse radius term or changes in the chemical potential of the bath.

The concentration of the bath is a solution of Laplace's equation. If the local concentration is lower than the bulk concentration, diffusion will be inward. This occurs when the particle has higher radius and thus a lower chemical potential than the mean –

it grows further. Oppositely, if the local concentration is higher than the bulk concentration, diffusion is outward, drawing material away from smaller particles.

While the larger particles can continue to absorb material as long as the bath can supply it, the smaller particles must eventually run out. In the long term, this leads the bath concentration to drop. An unstable equilibrium is found at the point of no net flux. As the component density drops, particles can fall behind from growth to shrinkage. This continues until the system is quenched (in the case of an anneal), or there is one large particle.

## 8.2 Common elements of the modified theories

In a normal metal, the screening of these boundary charges would follow an exponential decay. However, graphene is not a normal metal. Its charge carriers have a constant high velocity around c/300 [70]. A direct consequence of this is a low density of states, especially near the charge neutrality point. This difficulty creating a strong charge concentration weakens screening in graphene. Due to the unusual dispersion relation, this weakness is not expressed as a long penetration depth, but rather as the decay following a rational function form rather than an exponential[24]. This decays over a far greater distance.

So, for any few layered graphene flakes, it is a better approximation to consider these charges spread evenly through the thickness of the graphene than to consider them bunched up at the interface. We take this to actually be the full film thickness; as the film thickness is generally taken to a low power, this does not strongly impact the results. While this greater separation increases the dipole strength at large ranges, having a larger dipole widens the near-field region where the point dipole approximation fails. Thus, thicker graphene produces a lower dipole energy. The particular way this is handled differs between the two derivations.

A second pertinent observation is that the particles are around 5 times wider than they are tall. We approximate that the top-side surface energy is constant per surface area, and that the height is a constant, h. This permits us to treat the bulk volume energy as a part of the per-area energy. This approximation also allows us to ignore this per-area energy when considering where gold will go: under this approximation, rearranging the gold cannot change the top-side area. The remaining unbalanced tension is the tension around the edge of the particle. We take this tension to be the surface tension of the metal at its melting point, $\gamma$, times h, times the perimeter of the particle. The tension per length is labelled t.

## 8.3 Theory of circular nanoparticles

Most of the particles seen are circular. In this case, the electrostatic energy will merely limit the particle's size. We can assume that the particle is circular and proceed from there.

When computing the energy due to these dipoles, because the short-ranged component can be considered to be a part of the bulk energy of the droplet, we do not need to worry about its exact functional form. We cut off our energy integral so as to not include any contribution from lengths shorter than the dipole length, and then use the far-field approximation. This captures the relevant interactions while leaving the mathematics tractable.

We integrate the dipole repulsion arising from a work function difference ΔΦ, over the bottom area of a circular nanoparticle with radius R, cutting off the near term at the dipole width, which is proportional to the number of graphene layers L. The resultant electrostatic energy is

$$U_e = G\Delta\Phi^{\frac{7}{2}}R^4L^{-1}$$

G is material-independent and equal to 0.178 $nm^{-4}$ $eV^{-5/2}$.

The line tension is a simple multiplication, yielding a combined energy of

$$U = 2\pi Rt + G\Delta\Phi^{\frac{7}{2}}R^4L^{-1}$$

To find the preferred particle size, we evaluate the chemical potential μ. Rather than isolating the free energy per particle, we find the most convenient expression is for energy per area – so we multiply by the number density n and the droplet height h.

$$nh\mu = \left(\frac{1}{2\pi R}\right)\frac{\partial U}{\partial R} = \frac{t}{R} + \frac{2G\Delta\Phi^{\frac{7}{2}}R^2}{\pi L}$$

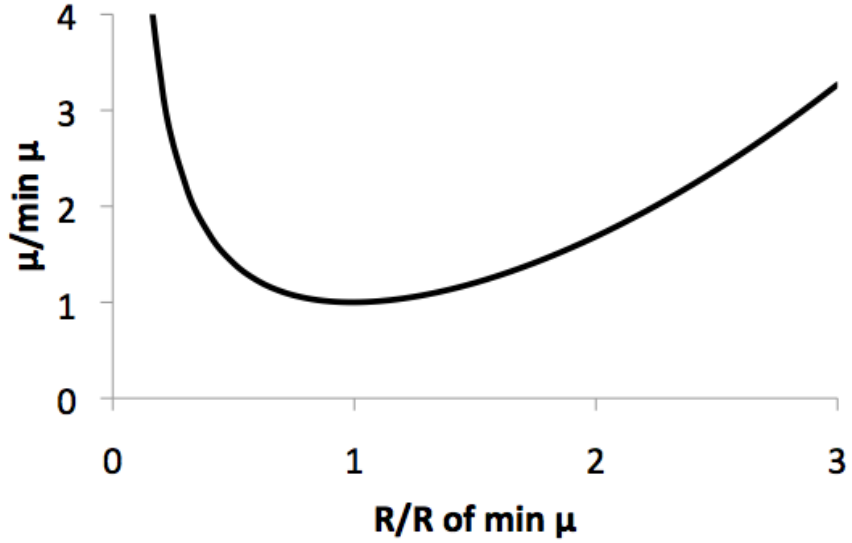


Figure 1: The form of μ as a function of radius. Different materials will have a different work function difference and surface tension, which will change the scale that this curve would take.

By minimizing this free energy (plotted in figure 1), we find the preferred radius for droplets. Principally, we note that the power of the layer count dependence is 1/3. This is in very good agreement with the observed 0.331 ± 0.061 power law observed for gold, and also consistent with the.

$$R^3 = \frac{\pi L t}{4 G \Delta\Phi^{\frac{7}{2}}}$$

As to the particular values, if we use the work function difference and tension of gold, we get a mean radius of 3.0 nm at one layer. The measurement yielded 3.23 ± 0.34 nm for this figure; this is very good agreement, within one sigma. For titanium, the prediction is 5.7 nm. This 1.2 sigma from the observed radius of 4.5 ± 1.0 nm.

The width of these energy minima will yield the expected dispersion in particle size. To be more specific, the variance in particle radius will be kT at the time of droplet

formation (during the anneal) divided by the second derivative of the free energy in respect to radius.

$$\frac{\partial^2 \mu}{\partial^2 R} = \frac{2\,t}{R^3} + \frac{4\,G\,\Delta\Phi^{\frac{7}{2}}}{\pi\,L}$$

At the mean radius given above, this simplifies to 12 G $\Delta\Phi^{7/2}$ / π L; this results in a variance of

$$\sigma^2 = \frac{L}{4\,G\,\Delta\Phi^{\frac{7}{2}}}\,kT$$

The single L in the numerator produces the linear form used in the fit above. For gold and the annealing temperature we used, this coefficient evaluates as 1.6 $nm^2$ per layer; we observed 0.71 $nm^2$ per layer. It is quite possible that the particles finally quenched only as they were cooling, which would lower the effective temperature. Also, if the observed variance on *thin* graphene was inflated due to greater observational error on smaller particles, that would lower the observed per-layer difference below the actual value. Both of these effects would tend to improve the agreement.

## 8.4 Theory of irregular nanoparticles

The earlier derivation assumed that the nanoparticles are circular. Here, we focus on the energy of small area-conserving deviations from circularity. If the energy to deviate from circularity is negative, particles will be required to deform; as we will see, there will not be a continuous variation in most stable shape but rather, once the conditions are set for a particle to deviate from circularity at all, the deformation will grow until limited by some other effect.

### 8.4.1 Mathematical description of nanoparticle shape

As the particles are nearly circular, we can consider the particle shape to be a function taking an angle and yielding a radius. This must be periodic as we return to the same angle. Thus, it is natural to take the function's Fourier series,

$$r(\theta) = Re\left(\sum_{n=0}^{\infty} r_n e^{in2\pi}\right)$$

with complex $r_n$.

$r_0$ is simply the mean radius. $r_1$ is a figure of how off-center the particle is in whichever coordinate system we have devised. By moving our coordinate system around, we can set $r_1$ to be 0. It will be convenient to fix it in this way as it does not add any features of interest and would interfere with the area measure.

$r_2$ is the narrowing-lengthening mode. Its phase relative to $r_0$ indicates the direction of lengthening. As examples, 1 is E-W, i is NE-SW, -1 is N-S. $r_2$ is thus of prime interest (figure 2).

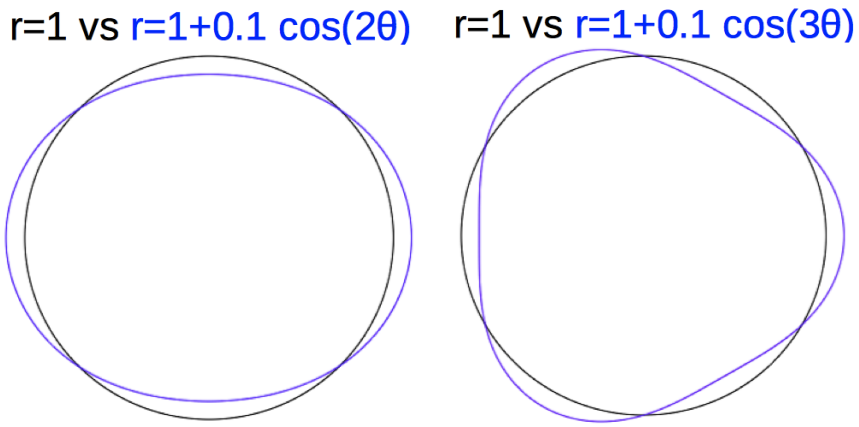


Figure 2: illustrations of the two most significant modes. The cosine is equivalent to the given exponential form, with a real amplitude; sine equivalent with an imaginary amplitude.

$r_3$, similarly, produces the observed Y shaped branches; but $r_4$ would be + shaped branches, which are very uncommon. No higher modes have been recognized in the observations, so we expect the theory not to produce them.

The area of a droplet is a function of its Fourier components. Splitting the particle up into differentially thin right triangles with length r and width r dθ, we get

$$A = \int_0^{2\pi} dr \frac{r}{2} (r d\theta)$$

This is very convenient – integrating this brings out the orthonormality of the Fourier modes, so we then get

$$A = \pi \sum_{m=0}^{\infty} |r_m|^2$$

This seems at first counterintuitive, as each anisotropic mode both fattens and pinches the particle; but the squaring makes the fattening portions outweigh the pinching portions. In short, if we consider the various Fourier modes as a vector, we must keep the vector normalized in order to preserve area.

### 8.4.2 The line tension energy

There are no new forces in this system, still only surface tension and electrostatic dipole repulsion. However, their functional forms have grown a substantial new layer of complexity.

The line tension energy is still of the same fundamental form, $U_L = Pt$ with P the perimeter, but each Fourier mode adds length to this perimeter. To simplify matters, we will unroll the edge of the particle and consider it as a rectilinear function, by letting p be the proper distance around the circle, $r_0\theta$. The full path length P will include this and the additional length from deformations. So long as the narrowest radius is not substantially smaller than the largest, this will be a valid approximation. The more pertinent approximation is that of a power of a number close to 1; taking this approximation turns the square root into a ½.

$$P = \int_0^C \sqrt{1 + \left(\frac{dr}{dp}\right)^2}\, dp \approx \int_0^C 1 + \frac{1}{2}\left(\frac{dr}{dp}\right)^2 dp$$

Changing back to θ so we can integrate the Fourier series, this then simplifies to

$$P \approx 2\pi\, r_0 + \frac{1}{2}\int_0^{2\pi} \left(\frac{dr}{dp}\right)^2 d\theta$$

$$= 2\pi \left( r_0 + \frac{1}{2r_0} \sum_{m=2}^{\infty} m^2 \left|r_m\right|^2 \right)$$

Since the deformations are of primary interest, we will transform this to conserve mass. As we are assuming that the particle has vertical walls, this can be achieved by arranging that these deformations will leave the area constant. This is accomplished by making any amplitude in a nonzero mode come at the cost of a corresponding amount of $r_0$. This reduction in $r_0$ induces a reduction in path length; this comes out as as subtracting 1 off of the $m^2$ coefficient. We also refactor to render the summed terms dimensionless. So,

$$U_L = 2\pi\, r_0\, t \left( 1 + \frac{1}{2} \sum_{m=2}^{\infty} \left(m^2 - 1\right) \left| \frac{r_m}{r_0} \right|^2 \right)$$

Thus, adding any other modes adds some line tension; the 3rd, Y-branching mode twice as much as the lengthening mode per amplitude, and the 4th, X-branching mode nearly twice as much as that.

### 8.4.3 The dipole energy

The full formula for the dipole energy is substantially more complicated. It is the two dipole interaction, integrated over the interior of the droplet. All of the difference from the last chapter is in the newly complicated edge.

The form of this energy will be the integral over pairs of dipoles,

$$U_d = \frac{\Delta\Phi^2}{32\pi^2 e^2} \iint d^2\vec{r'} d^2\vec{r''} f\left(\vec{r'}\right) f\left(\vec{r''}\right) g\left(\left|\vec{r'} - \vec{r''}\right|\right)$$

with g being the form of the interaction, and f being the support function: 1 wherever the metal is present and 0 otherwise.

So we can avoid having to use piecewise functions, we need to pick a near-field form for the dipole interaction. A convenient form is

$$g\left(r\right) = \left(r^2 + d^2\right)^{-\frac{3}{2}}$$

with d the dipole length. This is chosen as it gives the correct far-field form, a non-divergent near-field form, and is simple to work with.

Evaluating this integral is complicated; the most relevant portions are shown in appendix 2. The result is

$$U_{shape} = \sum_m g_m(r_0) |r_m|^2$$

with m > 1 and

$$g_m(r_0) = \frac{\Delta\Phi^2}{8e^2 d}\left(\frac{r_0}{d}\right)\int \zeta \frac{e^{-\zeta d}}{d} J_m^2\left(\frac{\zeta r_0}{d}\right)\, d\zeta$$

That expression is in the simple case where $r_0$ is held constant.

However, just as with the line tension calculation above, in order to preserve area – which is the relevant comparison for the deformation of particles – $r_0$ must decrease with increases in the various $r_m$ to hold the particle area constant.

To compensate for this, we split the energy into an area-independent part, U', and an area-dependent part, $\gamma_{int}$. This latter symbol is selected as this is conceptually related to a tension, though an 'internal' one. Stating the energy properly, we have

$$U' = U - \gamma_{int} A = \sum_m (g_m - \pi\gamma_{int}) |r_m|^2$$

with the sum over m including 0. The various g are functions of $r_0$, and $\gamma_{int}$ is a function of the full set of $r_m$. To determine $\gamma_{int}$, we can apply our requirement that U' is area-independent.

$$\frac{dU'}{dA} = \frac{1}{2\pi}\sum_m \left[\frac{1}{r_m}\frac{dU'}{dr_m}\right] - \gamma_{int}$$

We can isolate these parts by taking partial derivatives of the full energy in respect to each mode.

$$\gamma_{int} = \frac{1}{2\pi}\sum_{mn}\frac{1}{r_m}\frac{\partial}{\partial r_m}\left[g_n |r_n|^2\right]$$

Upon taking this derivative, we find that the area-independent shape dependence is the integral of a fourth-order polynomial in $r_0/d$, with the coefficients being Bessel functions of $r_0/d$ times the integration parameter. The integral of these Bessel functions will have inverse dependence on $r_0/d$. There is one constant term in the cubed term that avoids these inverse dependences, while the quartic terms are fully affected, lowering them to a net quadratic dependence. As $r_0/d$ is large for these particles, we keep only the leading term. For a more complete discussion of this, see appendix B in the paper given at the start of the chapter. This surviving term is

$$U_{shape} = -\left(\frac{\Delta\Phi^2}{16e^2 r_0}\right)\left(\frac{r_0}{d}\right)\sum_{m\neq 0}|r_m|^2$$

### 8.4.4 Island shape

The amplitudes of the modes are, to this quadratic order, independent. Therefore we can consider a combined line and electrostatic energy for each mode; this determines separately for each mode the energy difference for deforming in that mode. For m = 2 and larger, we have

$$U_m = \pi\gamma h \frac{|r_m|^2}{r_0}\left(m^2 - 1 - \frac{\Lambda}{2}\left(\frac{r_0}{d}\right)^3\right)$$

with

$$\Lambda = \frac{\Delta\Phi^2}{8\pi e^2 \gamma h}$$

The order m appears in one place in this equation, and larger m leads to larger $U_m$. Thus, we can see that higher modes are, within the bounds of this approximation, less energetically favorable. On the other hand, a larger overall particle radius, $r_0$, lowers the energy of deformation. Above a critical radius, the barrier to deforming in a given mode is negative. This tipping point is

$$r_c = d\left(\frac{2\left(m^2-1\right)}{\Lambda}\right)^{\frac{1}{3}}$$

If a particle grows from a small size it will deform into a linear (m = 2) shape first, which will prevent reaching the conditions under which higher-mode deformations occur. Based on this, we expect to see primarily linear structures. Furthermore, the width of these structures should be no more than twice the critical radius, as any material added above that radius would go to lengthening the particle. The width can end up less than this limit, as the theory given does not address the issue of how far the deformation will proceed.

As we were perturbing a circular particle, this derivation does not touch on the existence of lower-energy elongated shapes with smaller amounts of material. Thermal fluctuations or accident of dynamics - or an as yet undiscovered systematic effect in dynamics - could lead particles to be elongated at smaller radius than this. $r_c$ is the radius at which circularity is not even metastable.

Bringing in some numbers, for Yb, $\Delta\Phi$ = 2.2 eV and $\gamma$ = 320 erg/cm$^2$ ([69] [57]); additionally, we have measured these nanoparticles to be around 2 nm in height. Combined, this yields a critical radius of 5.7 nm. The observed width of the strands is 4.01 ± 0.08 nm, which is in the range dictated by the theoretical value.

By comparison, Au has a critical radius of 26 nm, far larger than the largest particles produced (7 nm radius).

## 8.5 Comparison between circular and deformed particles

The reverse comparison is interesting. The equation given for circular particles in the previous section,

$$R^3 = \frac{\pi L t}{4 G \Delta\Phi^{\frac{7}{2}}}$$

produces a particle radius of 0.31 nm for Yb, far smaller than the critical radius or the observed half-width. This is shorter than the assumed height of the particle, suggesting that particles of this size are not in the domain of the theory.

Regardless, the expected radius of a circular particle has little to do with the expected width of a straight strand. Recall that one term of the free energy of a surface is inversely proportional to its curvature. If a strand were only as wide as the mean R given above, the straight sides could be thus be lower in free energy than the sharp ends. This would create a flux from the ends to the middle. The strand thickens, raising its internal repulsive energy and lowering the curvature of the ends, until the two balance.

For very small particles, the curvature term will be large, creating a large discrepancy between the width of a stable elongated particle and the diameter of a circular particle.

## 8.6 Conclusions

Two theories are presented, both based on electrostatic dipole repulsion arising from charge transfer between metal nanoparticle and the graphene substrate.

The first predicts the mean radius of a set of circular nanoparticles grown on graphene; the central, material-independent prediction is a radius dependence of layers$^{1/3}$, which is consistent with experiment. It also provides specific radius predictions for individual metals; the prediction for gold is also consistent with experiment. The theory is also consistent with the observations of silver and titanium, though these observations provide weaker constraints on the theory.

The second theory shows that this electrostatic force can destabilize circular droplets, leading the particles to assume elongated shapes. A critical radius at which circularity is unstable is determined. Experimentally, deformation of circular particles is seen for particles much smaller than this size. The theory shows that the mechanism of dipole repulsion is sufficient to overcome line tension; it does not make any claim as to the balance between circular and irregular particles, except that above the critical radius, circular particles will tend to deform even without the ability to explore large regions of configuration space.

# 9: Conclusions and Future Directions

Carbon nanotubes and graphene are materials with great promise that is only beginning to be realized. Their interaction with other materials is only second to their intrinsic properties as being fundamental to harnessing this potential. As a particular instance, our lab has studied chemical decoration with single-stranded DNA for gas sensing; as there are an exponentially large number of DNA strands, they provide many conceptually similar but effectively different functionalizations. With this application in mind, we launched investigations of how single-stranded DNA behaves on and around graphitic surfaces.

Most graphene manufactured, however, will never touch DNA. A more substantial portion of it will come in contact with metals, and as leads and contacts get smaller and smaller, the behavior of small metal particles on graphene demands . And there are indeed distinct opportunities to take advantage of: metal filaments on graphene have several potential applications, ranging from nanofluidics to a tool for wafer-scale nanofabrication.

The following section summarizes the thesis and its findings, and suggestions for future work.

## 9.1 Summary

Two experiments dealt with single-stranded DNA on graphitic surfaces. This system was in use for electronic transport experiments on the ability of this hybridization to provide a current response to the presence of various chemical analytes[60,89]. The outstanding diversity and surprisingly strong selectivity of this hybrid urgently raised the question of how it worked and what it was actually like on the atomic scale.

In response to this need, we examined systems as closely related as possible in different ways. Each provided quite different information. The TEM measurements directly imaged ssDNA strands wrapped around suspended tubes. We observed that the ssDNA clumped up wherever it could not form a continuous layer, as evidenced by the prevalence of features longer than the length of the strand applied. In some cases the clumping led to extension of the clump away from the tube. The wrapping of the tubes was seen not to be, for the most part, helical. In the one case where a comparatively neat chiral wrapping occurred, it was around two adjacent tubes rather than one.

A second means of examination provided complementary and agreeing results – the models of Bob Johnson [44], which indicate that the wrapping is generally not helical due to crowding. Having space available between the tubes could alleviate the congestion, permitting helicity.

Examination through X-ray reflectivity provided complementary information concerning ssDNA on flat graphite surfaces. This would be somewhat related to ssDNA on carbon nanotubes, but of course is far more similar to graphene surfaces. The ssDNA on graphene surface is also directly useful for sensing applications similar to those with nanotubes[60].

The X-ray data indicate that the most typical height of the ssDNA film on the surface is 1.6 nm. However, there is a great deal of roughness to the distribution. This is not at all surprising, as the just-discussed TEM images of ssDNA on swCNT produced similar results.

The experiments dealing with metal nanoparticles on graphene and few-layer graphene led to the development of a simple but powerful theory governing droplet growth. Though developed based on graphene and metals, the theory applies to any two materials bearing the following qualities:

The substrate and particles must both be conducting. One (or both) of them must be a poor screener of electrical fields. They must have differing work functions. The particle components must preferentially bind to each other rather than smoothly wetting the surface. They must also not chemically react, alloy, or otherwise mix with the surface. The particles must consist of parts that are mobile, and small enough to be effectively continuous (e.g. not conducting polymer macromolecules). The particles must be small enough that bulk crystallization has not set in.

Other candidate material pairs besides graphene with non-carbide-forming metals would be any metal substrate with organic conductors with low carrier densities, and metals on weakly doped semiconductors. Other materials currently under investigation that may exhibit this family of effects would be the semiconductors $MoS_2$[79] and hexagonal Boron Nitride.

The theory given is that the electrical dipoles induced by the work function difference of the materials press outward on one another. This energy competes with the boundary energy of the particles. If the work function difference is small or the screening lengths are both short, this is insignificant. If it is strong enough, it limits the growth of the particles and prevents runaway Ostwald ripening. If it is stronger still, it deforms the particles into elongated structures so that the ends of the particles can get further from each other. At high concentrations, extensive networks can form.

In the case of circular particles, there was excellent quantitative agreement between the theory and observations - the theory states that the preferred radius of particles varies as the 1/3 power of the number of graphene layers, while the measurement was a $0.331 \pm 0.061$ power law. Based on materials properties for Gold, the theory gives single-graphene-layer preferred radius of 5.9 nm, while the measurement yielded $6.48 \pm 0.68$ nm.

In the case of the filamentary structures, the prediction was not an equality, but an inequality – the width of the structures should not exceed twice the critical radius. Using parameters for ytterbium, the critical radius was 5.7 nm. The actual mean filament width was $4.01 \pm 0.08$ nm.

This explanation may also help explain aspects of an experiment on carbon nanotubes which observed that silver formed larger nanoparticles on the tubes than gold did[97].

## 9.2 Future directions of inquiry

The first experiment to pursue would be to see whether an analogous effect occurs with other material combinations. In particular, we have already seen hints that gold on hexagonal boron nitride (BN) may produce filamentary structures similar to those in the ytterbium on graphene system. Molybdenum disulfide, a layered semiconductor, also

seems promising. Leads could be grown by this method, then selectively severed by lithography to create devices.

If BN does produce filamentary structures in metals with high work function, then the wider variety of surface tensions among such metals opens the possibility arises of making metal networks with selectable width.

Whether that particular technique works or not, these filaments may be useful as negative moulds for fluid systems. This application would distinctly benefit from width control.

Lastly, the filamentation mechanism offers an entirely new route to fabricating nanoscale graphene constrictions. As discussed in the introduction, constricting graphene can lead to its being semiconducting, opening it up to application in digital electronics.

We have already demonstrated[61] that narrow metal strips can be used as a mask in an oxygen plasma etch to create graphene transistors with a high on/off ratio. The width of Yb ribbons on graphene is 4 nm, well into the range for usefully narrow semiconducting channels (< 10 nm).

Thus, a future direction of research is to explore whether dense, smooth Yb networks can be grown on CVD graphene and used to fabricate dense arrays of graphene semiconducting nanochannels. The processes outlined below would involve only gentle sample treatment, with no sonication, powerful acids, high temperatures, or other such disturbances. It also does not rely on electron-beam lithography, special alignment, or indeed any steps smaller than wafer-scale.

To get a dense network that can be expected with high reliability to cross any given gap, all that is required is to deposit a lot of Yb (figure 1). The issue that this creates in return is greater irregularity of features, which will increase scattering. The very high-coverage end seems promising, for rather than forming networks of Yb strands, the densest regions of Yb were uniformly mottled surfaces that could be a template for semiconducting graphene nanomesh[6]. If that works, it would simplify things greatly, rendering many of the following issues irrelevant. Therefore, the following is written supposing that narrow channels are the only possible option or only option applicable to the given purpose.

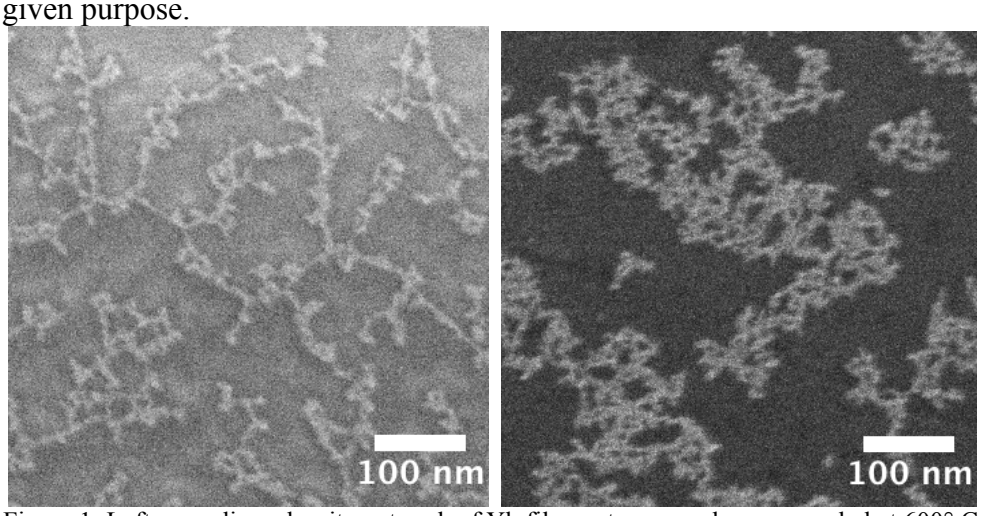


Figure 1: Left: a medium-density network of Yb filaments on graphene, annealed at 600° C. Right: a dense network of same, showing mottling.

To resolve the issue of irregularity, a lower annealing temperature should be used than the 600 C used for these experiments: lower temperatures will permit less thickness variation from the energy optimum. As the annealing temperature used to produce the networks in these experiments was quite high, and the melting point of Yb is quite low, there is a lot of room to explore. The only experimental lower bound we have on the annealing temperature of Yb is that it didn't anneal at room temperature.

A simple starting try would be to put 0.75 nm deposition of Yb on a hot plate at temperatures graphene can survive, such as 150 C, for extended periods – days may be needed. This could fail either because it is not hot enough to anneal Yb or it is hot enough to cause Yb or graphene to oxidize. Either way, moving back into the controlled-atmosphere oven would be necessary.

Rather than take up the oven for days on end while flowing $H_2$ and Ar, one distinct possibility is a two-part anneal, with a brief spike to elevated temperatures to get most of the work done, with the remainder of a more conventional several hour anneal time at a lower temperature.

Once the networks of Yb have been made, contacts should be laid down by conventional lithography. Designing the proper width of contacts to guarantee a semiconducting connection will depend on the deposition choice made earlier.

Once the contacts are in place, the extraneous graphene must be removed. Finding the proper recipe for an oxygen etch is crucial; we can start with the parameters used earlier in [61], though the dimensions of that constriction may differ from these.

Removal of the Yb is required to leave the transistors non-shorted. Dilute sulfuric acid is suggested[95]. Alternately, the Yb could be oxidized to render it insulating, simply by dunking it in warm water. This latter method is of course not suited to the fabrication of chemical sensors, which require the graphene to be exposed to the environment.

As these semiconducting channels would not require individual attention to create, and occur with controllable density, if the technique can be made to work for a few devices it can just as easily be made to work for many. While there is a great deal of room for problems to occur in this proposed process, limiting its application, there is similarly room for such problems to be resolved.

# Appendix 1: Experimental Details

The main text covered the main parts of the recipes – what is done under normal circumstances. This is useful, but to be able to reliably conduct the process requires being ready for a variety of contingencies and less interesting preparatory steps. For instance, it is reasonable for a recipe to assume a clean chip, even though getting it that way is not at all trivial. There are, similarly, several strategies in growing carbon nanotubes that were not directly relevant to this work, but may be useful. Potassium Hydroxide etching is simple if done correctly, but there are many simple errors one can make that will destroy the sample; a section describes how to avoid these.

Lastly, Atomic Force Microscopy is given a full practical examination, from beginning to end.

## A1.1 Chip cleaning

Starting with a clean surface is vital to a successful experiment. If a box of wafers is freshly opened, the wafers will be clean. If the box is immediately resealed, the wafers inside will remain quite clean. If the box is left open or, to a lesser extent, incompletely sealed, the surface will require cleaning to be used. Note, when cutting a clean wafer, silicon dust will get on the surfaces; it can be removed with a vigorous nitrogen blast. What follows is for the case in which the chip requires a full cleaning.

The first few steps dissolve anything on the chip, then rinse with 'clean' solvents that do not drop much material on the surface when drying (acetone in particular leaves scum behind). Even the cleanest solvent is removed vigorously rather than allowing it to dry in place.

First, soak the chip in Dichloroethane (DCE) for half an hour. If there is anything caked onto the surface, consider using a different chip without this problem; if it really is necessary to use this specific chip, sonicate it in DCE until the caked material has been removed.

Upon taking the chip out of the DCE, mount it in the spinner. Prepare acetone and isopropanol (IPA) squirt bottles, and a nitrogen gun with the regulator set so it produces a vigorous blast. While spinning at 2000 rpm, spray the chip with acetone for five seconds. Start spraying IPA, then stop the acetone. After five more seconds, spray the dry nitrogen and stop the IPA. Continue nitrogen until the chip is dry, and for 5 seconds after -- it will not be perfectly clear when the chip, spinning 33 times a second, is entirely dry.

Last, we remove any organic residue. At this point, there are two options. One is to put it on a hot plate at 300 C in air for half an hour. The other is to oxygen plasma etch at 40 W for 5 minutes.

This cleaning must be performed as immediately as possible before any following processing steps.

## A1.2 Carbon nanotube growth

This section discusses some details and alternatives to the method described in the main text.

The main choices arise when applying the Iron Nitrate in IPA solution to the substrate. Commonly, the iron nitrate is dropped onto the surface, and then spun off at

3000 rpm; in other variations, the drop is blown off with a dry nitrogen burst, or applied to the chip while it is spinning at a lower speed. This last was commonly used in the lab but abandoned, as the results are strongly dependent on the exact placement of the drops in respect to the center of rotation. The extreme method of dipping the chip in the solution and blowing dry was successfully used to produce large numbers of tubes, but the resulting samples were dirty.

The benefits of using a retarder are described in the main text. An improvised but effective retarder can be a wad of quartz fiber insulation jammed into the tube just inside the inlet end of the furnace. Once set up this is clean and effective; but when first used, it can be prone to flying down the tube, and tiny parts can flake off and dirty any samples placed too close. An alternate retarder with neither of these drawbacks is a modified quartz tube - indentations are made in the tube, partially obstructing flow and bringing about mixing.

There is an additional method of increasing nanotube growth: oxygen plasma etching away as-grown tubes and growing a second time on the same chip produces a huge number of tubes. These were highly defective, to the extent that tubes lying across windows had multiple kinks.

### A1.3 KOH etching

Before KOH etching, it is vital to ensure that the back-side nitride has been completely etched. The etched areas must look silver rather than brown or yellow. This indicates remaining nitride, and a slight tinge is enough to completely prevent etching. Start the KOH etch soon after the $SF_6$ etch to avoid forming a thick native oxide, which can also interfere with even etching. Within half an hour is fine.

As stated in chapter 3, KOH is gradually added to water in a 1:2 mass ratio. Select the smallest beaker possible while leaving room at the top for the hanger and room at the bottom for the magnetic stir-bar. Keeping it small does not only minimize waste, it also makes the temperature reading more accurate. The stirring should be sufficiently vigorous to form a bit of a vortex. Suspend the chip firmly, holding it by a sacrificial (unetched) corner; do not let the clip touch the solution. Metal contamination, especially copper, greatly reduces the etching ability of KOH. When removing, promptly dunk in DI water to avoid allowing the KOH to become highly concentrated.

The non-ideal surface referred to in the main text may arise from the presence of pinholes in the topside $SiN_x$. Each pinhole would grant KOH access to the Si underneath. Etching starting from the front side would finish off the 100 and 110 planes quickly, and the rest of the time would be slowly eating away at the 111 planes. This would excavate under the nitride, and thus make additional membranes centered on the pinhole, not opening up in the rear. These membranes can lead to confusion. In the worst cases, the membranes meet and merge; this can lead to most of the chip being etched away.

### A1.4 AFM image optimization, recovery, and processing

Atomic Force Microscopy is a very powerful technique, but the basic training given to users does not usually include much information on how to handle a variety of issues

that commonly arise. Especially in the direction of handling and processing AFM data, much is left unsaid. This section is intended to help fill in this gap.

This section should be broadly applicable, but in some cases, specific instructions are most relevant to particular equipment or software. We used Veeco's Nanoscope III system with the vendor-supplied control software, and analysis was principally performed in Gwyddion[50].

There are three software issues worth mentioning directly.

Two Nanoscope commands that save a great deal of time are 'jump to line' and 'capture last'. The capture last command acquires the previous frame, and is a lifesaver if one forgot to instruct the instrument to acquire the image in advance. The jump to line command is primarily useful while searching for the region of interest. Use it to create a 'jailbar' image, including only, say, lines 1-5, 50-55, 100-105, etc. Multiple lines should be taken for several reasons. First, it allows the after-jump transient to subside. Secondly, it allows one to see the slope of any boundaries. This usually gives enough information to tell where the frame lies in respect to the sample, while taking a tiny fraction of the time of capturing a full frame.

The second issue is that Nanoscope files are binary, but include a text header. This confuses many email clients and file transfer programs, which will examine the header and begin transferring in ASCII mode. The binary data are then corrupted by the various transformations that are applied to ASCII data (principally, line endings can change from CR-LF to LF or CR, but others are possible). In some versions of Windows, the default file compression utility accessible via contextual menu (right-click) also corrupts the data in this fashion. Several ways of avoiding this are possible. The most convenient is to save the data directly to a portable storage device, but any binary-mode transfer will suffice. The Gwyddion software described later can detect files corrupted in this way.

### A1.4.1 Image optimization

The controls we have available to optimize the image are the cantilever tuning, scanning speed, amplitude setpoint, the proportional and integral gains, and as last resorts changing the tip or something about the sample.

When tuning the cantilever, for tapping mode, choose a frequency slightly below resonance. Choosing a frequency with amplitude 5-10% below the resonance peak is ideal. This maximizes the amplitude difference between in-contact and nearly-in-contact.

The scanning speed should generally not go above 10 microns per second, the exceptions being for very clean and smooth samples, and when ruining the tip is not a problem.

The amplitude setpoint controls how much the tip oscillations should be smaller than they would be if it were oscillating freely. An amplitude setpoint too high will not cause the tip to engage, or even if engaged will not cause the feedback loop to respond strongly to a momentary disengagement. An amplitude setpoint too low will jam the tip into the surface roughly, damaging both.

To find a good starting value, go to the 'scope' line trace view, and raise the amplitude setpoint until the tip disengages. This can be determined by the line traces becoming smooth waves. Then slowly lower the amplitude setpoint until it reengages, then lower the amplitude setpoint by 15-25% of its value. In the Nanoscope software, this

can be achieved by pressing the left arrow 3 times with amplitude setpoint selected. This will generally be in the right neighborhood. Further lowering may be needed.

For the proportional and integral gains, there are three matters to consider. First, these gains make the feedback loop respond to stimuli more promptly. This becomes an issue if there are tall features. When the tip comes off the side of them, it may take some distance to reach the surface again. If this occurs, the primary way to let the tip track the surface is to raise the proportional gain; it may also be necessary to lower the amplitude setpoint, and in extreme cases, lower the scan rate.

The gains only need to be large enough to make the tip follow the surface closely enough for the scan to be accurate. The second factor is that if the gains get to be too large and have the wrong phase, the feedback loop can become a positive feedback loop. If this occurs, a large oscillating signal will be superposed over the data. Try lowering the gains, and changing the balance between them. Normally, reducing integral gain restores a favorable feedback phase, but experimentation is sometimes key.

This is not the only possible source of oscillating signals; check that the oscillations are not 60 Hz noise or 120 Hz (rectified 60 Hz) noise by counting the oscillations per scanline, and dividing by half the scan period. If it is, there is most likely an electrical fault. Report it to the tool administrator. On the scale of the session, there's not a lot that can be done; it will have to be removed in image processing.

A third consideration on the gains is that from time to time there is noise proportional to the integral gain. This may depend on sample and tip characteristics, or it could be a technical fault in the two systems it was observed in. See if bringing the integral gain down to 0.01 or some other low value reduces the noise of the image. These low values are usually consistent with an acceptable phase, as mentioned above.

If the image simply seems noisy, one possible explanation is that the sample is rough. Try shrinking the scan size by a substantial factor – 4 or 10, say – and see if what appeared to be noise was actually a sparsely sampled uneven surface. You will frequently not need to take a whole field of data to determine this. Don't forget that you can increase the scan rate for smaller scans without making the tip move above the speed limit of 10 microns per second; just remember to lower the scan rate before restoring the large field.

### A1.4.2 Contingencies that can arise during imaging

As taking a full frame is a lengthy process, things can go wrong in the middle. A prompt appropriate response can turn a frame-ruining incident into two bad lines. The usual thing that occurs is that something has become attached to the tip (or detached from the tip).

If the attachment occurred right at the end of the tip, a single scan-line will have a large jump in it. The ensuing lines will return to normal, until the junk on the tip comes off, at which point there will be another jump in the opposite direction. These tend to occur repeatedly, picking up and dropping parts, possibly the same one. Lowering the amplitude setpoint tends to reduce the severity of the problem. Possible reasons for this are that the firmer contact knocks the junk away, that jamming the tip down onto it firmly enough attaches it more permanently, and that pressing it more firmly causes the junk to attach further up the tip rather than at the bottom where it can interfere with imaging. Very low amplitude setpoints may be required to achieve this. The tip is already dirty enough not to be a top-quality tip, so if the sample can tolerate the additional forces, there

is little reason not to lower it this way. Once the debris is gone, the amplitude setpoint can be restored to its earlier value. On longer time scales than a single session, if this continues to be a problem, try cleaning the samples more thoroughly. Cleaning by the procedure given earlier in section A.1.1 should be adequate, though this is not compatible with most finished samples. Simply heating the sample can get rid of sticky volatile substances. These do not always work.

Another common problem is that the tip loses contact with the sample. Something about the tip has changed that makes the old amplitude setpoint too relaxed; lower it until the tip regains contact. Watch out, this can happen again. At the end of the frame, this should be dealt with. Withdraw and re-tune the tip.

A most pernicious problem can be easily identified if contrast appears inverted, where features extending up appear to be extending downward instead. The cause of this inversion is that an electrostatic or magnetostatic attraction between debris on the tip and the feature on the substrate is distorting the force felt by the tip. In this case, it is vital to knock the debris off the tip.

Sometimes, restoring a tip is possible. For sticky things, putting it in an ozone-rich environment will remove organic contamination. For the electrostatic or magnetic debris, heating can remove the trapped charges or demagnetize the fragment.

### A1.4.3 Image processing

For AFM image processing, Gwyddion[50] is an open-source SPM analysis program that functions on a wide range of platforms; it is better in most respects than the Nanoscope software.

The standard image processing goals are to level the data, remove artifacts, and extract profiles.

A comprehensive workflow would be too much to describe, not only in this context, but any context; the problem of cleaning up AFM data is highly open-ended. This will explain the processes typically applied. Before beginning processing, an important step is to select reasonable colors for the image. In Gwyddion's Edit menu, the Color Gradients item selects the color scale; the Default Mask Color sets the color for arbitrarily shaped selections. The Gwyddion.net red-yellow color scale is very conventional. A mask color should be selected that is not among the colors in the gradient (blue works nicely with this red-yellow scale).

First, remove polynomial background, levelling the whole image. Normally, try to clean up the data without resorting to high-order polynomials. It is vital that the points being fit are actually supposed to be coplanar.

The mask is a set of selected points. These points can be excluded from the fit, leaving only coplanar points. Masks can be can be marked by threshold, and/or edited manually; other methods exist, but these two should suffice for selecting coplanar regions.

Even after an initial fit, the individual lines are frequently not calibrated to each other. To resolve this, the line correction tools adjust individual lines, usually only shifting them up and down. Of these, the two most common are the Median Line Correction and the Median Difference Line Correction. The former shifts each line so they all have the same median value; the latter shifts each line so that the median of the differences from the previous line is zero. The difference correction is better if large features comprise a

majority of some lines but not others, but is susceptible to causing drift if the data is noisy.

If neither of these is satisfactory, there is a tool, 'Level rows using intersections with given lines', which can produce decent results – you select lines, each of which is at a known constant height, and the line heights are adjusted to be best consistent with this assertion. However, it is more subject to noise than the others, as it only samples a small portion of each row. Like the others, this tool cannot change the slopes of scanlines, only their offsets.

Using just these tools is sufficient – though sometimes iterating them is required - to clean up many images. Sometimes, additional measures must be taken.

After the above steps is the best time to use the 'Correct horizontal scars (strokes)' tool, as replacing the bad portions of these lines with a Laplace interpolation before the neighboring points have settled down is pointless.

The 'Level data by fitting a plane through three points' tool is often the most reliable way of flattening, once there is no remaining curvature in the sample.

60 (or 120) Hz noise can be greatly attenuated with the "Data Process > Correct Data > 1D FFT Filtering…" tool. Stretch the tool window out as long as possible so the spectrum can be seen in detail, and look for a sharp peak at mid or high wavevector.

The Extract Profiles tool makes profiles; you must press 'Apply' to create the profiles as independent objects that can be measured, rescaled, and saved as images or data tables. Be sure to rename the profile to something indicating where the profile came from, as this information is not otherwise stored. The profile can then be found in the Data Browser, in the Graphs tab. Saving graphs as text or image is done via the Graph top-level menu, not File.

## Appendix 2: Calculation of Dipole Energy

This appendix details the method of evaluating an integral arising in the theory of nanoparticle shape selection:

$$U_d = \frac{\Delta\Phi^2}{32\pi^2 e^2} \iint d^2\vec{r'} d^2\vec{r''} f\left(\vec{r'}\right) f\left(\vec{r''}\right) g\left(\left|\vec{r'} - \vec{r''}\right|\right)$$

*f* is the support function for the nanoparticle: 1 where the particle is, 0 elsewhere.

Now, so we can avoid having to use piecewise functions, we need to pick a near-field form for the dipole interaction. A convenient form is

$$g(r) = \left(r^2 + d^2\right)^{-\frac{3}{2}}$$

with d the dipole length. This is chosen as it gives the correct far-field form and a non-divergent near-field form, and is simple to work with.

By taking the Fourier transform of this equation we can reduce the double integral to a single integral over wavevector q. Letting the Fourier transform of g(r) be v(q) and the Fourier transform of f(r) be S(q), the new equation is

$$U_d = \frac{\Delta\Phi^2}{32\pi^2 e^2} \int \frac{d^2\vec{q}}{2\pi} v\left(\vec{q}\right) \left|S\left(\vec{q}\right)\right|^2$$

While the Fourier transform of the interaction form is straightforward,

$$v(q) = \frac{1}{d} e^{-qd}$$

the Fourier transform of the support function contains all of the fresh complexity in the problem. In full,

$$S(\vec{q}) = \int_0^{2\pi} \int_0^{r(\phi)} r\, dr\, d\phi\, e^{-i\vec{q}\cdot\vec{r}} f(\vec{r})$$

We split this S into the components from the various Fourier components of r(ϕ). The S from the isotropic mode is

$$S_0(\vec{q}) = \int_0^{2\pi} \int_0^{r(\phi)} r\, dr\, d\phi\, e^{-i\vec{q}\cdot\vec{r}}$$

The angle integral produces a Bessel function. Integrating by parts, we get the (angle-independent)

$$S_0(q) = \frac{2\pi r_0}{q} J_1(q r_0)$$

which gives an isotropic component of the energy,

$$U_0 = \left(\frac{\Delta\Phi^2}{8e^2 d}\right) \int_0^{\infty} d\xi v'(\xi) \frac{J_1^2(\xi r_0 d)}{\xi} r_0^2 = g_0 r_0^2$$

The anisotropic component requires approximation. If we consider that $|r_m| << r_0$, we can turn the integration over r into a multiplication, thus yielding

$$S_m(\vec{q}) = 2\pi i^m J_m(q r_0)\, r_0 r_m e^{im\phi_q}$$

Bringing these back out to the equation for the energy, we get

$$U_{shape} = \frac{\Delta\Phi^2}{8e^2 d} \left(\frac{r_0}{d}\right)^2 \sum_m \left[ |r_m|^2 \int \zeta d\zeta v(\zeta) J_m^2 \left(\frac{\zeta r_0}{d}\right)\right]$$

$$= \sum_m g_m(r_0) |r_m|^2$$

with m > 1.

This yields the expresion used in the text.

## Bibliography


1 M. D. Abramoff, P. J. Magelhaes, and S. J. Ram, "Image Processing with ImageJ," Biophotonics International **11** (7), 36-42 (2004).

2 Jens Als-Nielsen, Didier Jacquemain, Kristian Kjaer, Franck Leveiller, Meir Lahav, and Leslie Leiserowitz, "Principles and applications of grazing incidence X-ray and neutron scattering from ordered molecular monolayers at the air-water interface," Physics Reports **246** (5), 251-313 (1994).

3 Adrian Bachtold, Peter Hadley, Takeshi Nakanishi, and Cees Dekker, "Logic Circuits with Carbon Nanotube Transistors," Science **294** (5545), 1317-1320 (2001).

4 Sukang Bae, Hyeongkeun Kim, Youngbin Lee, Xiangfan Xu, Jae-Sung Park, Yi Zheng, Jayakumar Balakrishnan, Tian Lei, Hye Ri Kim, Young Il Song, Young-Jin Kim, Kwang S. Kim, Barbaros Ozyilmaz, Jong-Hyun Ahn, Byung Hee Hong, and Sumio Iijima, "Roll-to-roll production of 30-inch graphene films for transparent electrodes," Nat Nano **5** (8), 574-578 (2010).

5 Jingwei Bai, Xiangfeng Duan, and Yu Huang, "Rational Fabrication of Graphene Nanoribbons Using a Nanowire Etch Mask," Nano Letters **9** (5), 2083-2087 (2009).

6 Jingwei Bai, Xing Zhong, Shan Jiang, Yu Huang, and Xiangfeng Duan, "Graphene Nanomesh," Nature Nanotechnology **5** (3) (2010).

7 N. C. Bartelt, W. Theis, and R. M. Tromp, "Ostwald ripening of two-dimensional islands on Si(001)," Physical Review B **54** (16), 11741 (1996).

8 Claire Berger, Zhimin Song, Tianbo Li, Xuebin Li, Asmerom Y. Ogbazghi, Rui Feng, Zhenting Dai, Alexei N. Marchenkov, Edward H. Conrad, Phillip N. First, and Walt A. de Heer, "Ultrathin Epitaxial Graphite: 2D Electron Gas Properties and a Route toward Graphene-based Nanoelectronics," The Journal of Physical Chemistry B **108** (52), 19912-19916 (2004).

9 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** (6430), 605-607 (1993).

10 A. Bianco and M. Prato, "Can Carbon Nanotubes be Considered Useful Tools for Biological Applications?," Advanced Materials **15** (20), 1765-1768 (2003).

11 G. Binnig, C. F. Quate, and Ch Gerber, "Atomic Force Microscope," Physical Review Letters **56** (9), 930 (1986).

12 G. Binnig, H. Rohrer, Ch Gerber, and E. Weibel, "Surface Studies by Scanning Tunneling Microscopy," Physical Review Letters **49** (1), 57 (1982).

13 J. Kent Blasie, Songyan Zheng, and Joseph Strzalka, "Solution to the phase problem for specular x-ray or neutron reflectivity from thin films on liquid surfaces," Physical Review B **67** (22), 224201 (2003).
14 D. K. G. de Boer, A. J. G. Leenaers, and W. W. van den Hoogenhof, "Influence of roughness profile on reflectivity and angle-dependent X-ray fluorescence," J. Phys. III France **4** (9), 1559-1564 (1994).
15 K. I. Bolotin, K. J. Sikes, Z. Jiang, M. Klima, G. Fudenberg, J. Hone, P. Kim, and H. L. Stormer, "Ultrahigh electron mobility in suspended graphene," Solid State Communications **146** (9-10), 351-355 (2008).
16 A. N. Broers, A. C. F. Hoole, and J. M. Ryan, "Electron beam lithography--Resolution limits," Microelectronic Engineering **32** (1-4), 131-142 (1996).
17 Helen Cathcart, Valeria Nicolosi, J. Marguerite Hughes, Werner J. Blau, John M. Kelly, Susan J. Quinn, and Jonathan N. Coleman, "Ordered DNA Wrapping Switches on Luminescence in Single-Walled Nanotube Dispersions," Journal of the American Chemical Society **130** (38), 12734-12744 (2008).
18 Jian-Hao Chen, Chaun Jang, Shudong Xiao, Masa Ishigami, and Michael S. Fuhrer, "Intrinsic and extrinsic performance limits of graphene devices on SiO2," Nat Nano **3** (4), 206-209 (2008).
19 C. E. Cross, J. C. Hemminger, and R. M. Penner, "Physical Vapor Deposition of One-Dimensional Nanoparticle Arrays on Graphite: Seeding the Electrodeposition of Gold Nanowires," Langmuir **23** (20), 10372-10379 (2007).
20 Robert F. Curl and Richard E. Smalley, "Probing C60," Science **242** (4881), 1017-1022 (1988).
21 Yaping Dan, Ye Lu, Nicholas J. Kybert, Zhengtang Luo, and A. T. Charlie Johnson, "Intrinsic Response of Graphene Vapor Sensors," Nano Letters **9** (4), 1472-1475 (2009).
22 Marie-Christine Daniel and Didier Astruc, "Gold Nanoparticles: Assembly, Supramolecular Chemistry, Quantum-Size-Related Properties, and Applications toward Biology, Catalysis, and Nanotechnology," Chemical Reviews **104** (1), 293-346 (2003).
23 S. S. Datta, D. R. Strachan, S. M. Khamis, and A. T. C. Johnson, "Crystallographic etching of few-layer graphene," Nano Letters **8** (7), 1912-1915 (2008).
24 Sujit S. Datta, Douglas R. Strachan, E. J. Mele, and A. T. Charlie Johnson, "Surface Potentials and Layer Charge Distributions in Few-Layer Graphene Films," Nano Letters **9** (1), 7-11 (2009).
25 C. R. Dean, A. F. Young, I Meric, C Lee, L Wang, S Sorgenfrei, K Watanabe, T Taniguchi, P Kim, K. L. Shepard, and J Hone, "Boron nitride substrates for high-quality graphene electronics," Nat Nano **5** (10), 722-726 (2010).

26 B. G. Demczyk, Y. M. Wang, J. Cumings, M. Hetman, W. Han, A. Zettl, and R. O. Ritchie, "Direct mechanical measurement of the tensile strength and elastic modulus of multiwalled carbon nanotubes," Materials Science and Engineering A **334** (1-2), 173-178 (2002).
27 T. Dürkop, S. A. Getty, Enrique Cobas, and M. S. Fuhrer, "Extraordinary Mobility in Semiconducting Carbon Nanotubes," Nano Letters **4** (1), 35-39 (2003).
28 E. L. Evans, O. P. Bahl, and J. M. Thomas, "The decoration of, and epitaxial growth of gold on, graphite surfaces," Carbon **5** (6), 587-588, IN523-IN524, 589 (1967).
29 Tore Frängsmyr and Gösta Ekspång (eds), *Nobel Lectures, Physics 1981-1990*. (World Scientific Publishing Co., Singapore, 1993).
30 Marcus Freitag, A. T. Johnson, Sergei V. Kalinin, and Dawn A. Bonnell, "Role of Single Defects in Electronic Transport through Carbon Nanotube Field-Effect Transistors," Physical Review Letters **89** (21), 216801 (2002).
31 K. Fukushima and et al., "Fresnel fringes in electron microscope images," Journal of Physics D: Applied Physics **7** (2), 257 (1974).
32 A. K. Geim and K. S. Novoselov, "The rise of graphene," Nat Mater **6** (3), 183-191 (2007).
33 M. Glicksman and P. Voorhees, "Ostwald Ripening and Relaxation in Dendritic Structures," Metallurgical and Materials Transactions A **15** (6), 995-1001-1001 (1984).
34 Brett Goldsmith, Jr. Joseph J. Mitala, Jesusa Josue, Ana Castro, Mitchell B. Lerner, Timothy H. Bayburt, Samuel M. Khamis, Ryan A. Jones, Joseph G. Brand, Stephen G. Sligar, Charles W. Luetje, Alan Gelperin, Paul A. Rhodes, Bohdana Discher, and A.T. Charlie Johnson, "Biomimetic chemical sensors based on carbon nanotube transistors functionalized with olfactory receptor proteins," ACS Nano (submitted) (2011).
35 T. Guo, P. Nikolaev, A. Thess, D. T. Colbert, and R. E. Smalley, "Catalytic growth of single-walled manotubes by laser vaporization," Chemical Physics Letters **243** (1-2), 49-54 (1995).
36 Jason H. Hafner, Chin-Li Cheung, Tjerk H. Oosterkamp, and Charles M. Lieber, "High-Yield Assembly of Individual Single-Walled Carbon Nanotube Tips for Scanning Probe Microscopies," The Journal of Physical Chemistry B **105** (4), 743-746 (2001).
37 J. Hass, R. Feng, Mill, aacute, J. E. n-Otoya, X. Li, M. Sprinkle, P. N. First, W. A. de Heer, E. H. Conrad, and C. Berger, "Structural properties of the multilayer graphene/ 4H-SiC(0001[over -bar]) system as determined by surface x-ray diffraction," Physical Review B **75** (21), 214109 (2007).
38 X. Hong, A. Posadas, K. Zou, C. H. Ahn, and J. Zhu, "High-Mobility Few-Layer Graphene Field Effect Transistors Fabricated on Epitaxial Ferroelectric Gate Oxides," Physical Review Letters **102** (13), 136808 (2009).

39 Sumio Iijima, "Helical microtubules of graphitic carbon," Nature **354** (6348), 56-58 (1991).
40 Sumio Iijima and Toshinari Ichihashi, "Single-shell carbon nanotubes of 1-nm diameter," Nature **363** (6430), 603-605 (1993).
41 Liying Jiao, Xinran Wang, Georgi Diankov, Hailiang Wang, and Hongjie Dai, "Facile synthesis of high-quality graphene nanoribbons," Nat Nano **5** (5), 321-325 (2010).
42 A. Johnson, S. M. Khamis, G. Preti, J. Kwak, and A. Gelperin, "DNA-Coated Nanosensors for Breath Analysis," Sensors Journal, IEEE **10** (1), 159-166 (2010).
43 Robert R. Johnson, A. T. Charlie Johnson, and Michael L. Klein, "The Nature of DNA-Base–Carbon-Nanotube Interactions," Small **6** (1), 31-34 (2010).
44 Robert R. Johnson, A. T. Charlie Johnson, and Michael L. Klein, "Probing the Structure of DNA-Carbon Nanotube Hybrids with Molecular Dynamics," Nano Letters **8** (1), 69-75 (2007).
45 Robert R. Johnson, Axel Kohlmeyer, A. T. Charlie Johnson, and Michael L. Klein, "Free Energy Landscape of a DNA-Carbon Nanotube Hybrid Using Replica Exchange Molecular Dynamics," Nano Letters **9** (2), 537-541 (2009).
46 Eugenii Katz and Itamar Willner, "Biomolecule-Functionalized Carbon Nanotubes: Applications in Nanobioelectronics," ChemPhysChem **5** (8), 1084-1104 (2004).
47 S.M. Khamis, R. A. Jones, A. T. Charlie Johnson, G. Preti, J. Kwak, and A. Gelperin, "DNA-decorated carbon nanotube-based FETs as ultrasensitive chemical sensors: Discrimination of homologs, structural and optical isomers," (2011).
48 Samuel M. Khamis, Robert R. Johnson, Zhengtang Luo, and A. T. Charlie Johnson, "Homo-DNA functionalized carbon nanotube chemical sensors," Journal of Physics and Chemistry of Solids **71** (4), 476-479 (2010).
49 Keun Soo Kim, Yue Zhao, Houk Jang, Sang Yoon Lee, Jong Min Kim, Kwang S. Kim, Jong-Hyun Ahn, Philip Kim, Jae-Young Choi, and Byung Hee Hong, "Large-scale pattern growth of graphene films for stretchable transparent electrodes," Nature **457** (7230), 706-710 (2009).
50 "Gwyddion". Petr Klapatek and David Nečas. gwyddion.net
51 Jing Kong, Hyongsok T. Soh, Alan M. Cassell, Calvin F. Quate, and Hongjie Dai, "Synthesis of individual single-walled carbon nanotubes on patterned silicon wafers," Nature **395** (6705), 878-881 (1998).
52 H. W. Kroto, J. R. Heath, S. C. O'Brien, R. F. Curl, and R. E. Smalley, "C60: Buckminsterfullerene," Nature **318** (6042), 162-163 (1985).
53 L. D. Landau, "To a Theory of Phase Transitions, part II " Phys. Z. Sowjetunion **11**, 26-35 (1937).

54 Changgu Lee, Xiaoding Wei, Jeffrey W. Kysar, and James Hone, "Measurement of the Elastic Properties and Intrinsic Strength of Monolayer Graphene," Science **321** (5887), 385-388 (2008).
55 Xiaolin Li, Xinran Wang, Li Zhang, Sangwon Lee, and Hongjie Dai, "Chemically Derived, Ultrasmooth Graphene Nanoribbon Semiconductors," Science **319** (5867), 1229-1232 (2008).
56 Xuesong Li, Weiwei Cai, Jinho An, Seyoung Kim, Junghyo Nah, Dongxing Yang, Richard Piner, Aruna Velamakanni, Inhwa Jung, Emanuel Tutuc, Sanjay K. Banerjee, Luigi Colombo, and Rodney S. Ruoff, "Large-Area Synthesis of High-Quality and Uniform Graphene Films on Copper Foils," Science **324** (5932), 1312-1314 (2009).
57 David R. Lide (ed), *CRC Handbook of Chemistry and Physics*. (CRC Press, 2001), 82nd ed.
58 I. M. Lifshitz and V. V. Slyozov, "The kinetics of precipitation from supersaturated solid solutions," Journal of Physics and Chemistry of Solids **19** (1-2), 35-50 (1961).
59 J. M. B. Lopes dos Santos, N. M. R. Peres, and A. H. Castro Neto, "Graphene Bilayer with a Twist: Electronic Structure," Physical Review Letters **99** (25), 256802 (2007).
60 Ye Lu, B. R. Goldsmith, N. J. Kybert, and A. T. C. Johnson, "DNA-decorated graphene chemical sensors," Applied Physics Letters **97** (8), 083107-083103 (2010).
61 Ye Lu, Brett Goldsmith, Douglas R. Strachan, Jong Hsien Lim, Zhengtang Luo, and A. T. Charlie Johnson, "High-On/Off-Ratio Graphene Nanoconstriction Field-Effect Transistor," Small **6** (23), 2748-2754 (2010).
62 Zhengtang Luo, Ye Lu, D.W. Singer, M. E. Berck, L. A. Somers, B.R. Goldsmith, and A. T. Charlie Johnson, "Growth of Uniform Wafer-size Graphene on Electropolished Copper," Journal of the American Chemical Society (submitted) (2011).
63 M. Makkee, E. J. A. X. van de Sandt, A. Wiersma, and J. A. Moulijn, "Development of a satisfactory palladium on activated carbon catalyst for the selective hydrogenolysis of CCl2F2 (CFC-12) into CH2F2 (HFC-32)," Journal of Molecular Catalysis A: Chemical **134** (1-3), 191-200 (1998).
64 Sharali Malik, Stephanie Vogel, Harald Rösner, Katharina Arnold, Frank Hennrich, Anne-Kathrin Köhler, Clemens Richert, and Manfred M. Kappes, "Physical chemical characterization of DNA-SWNT suspensions and associated composites," Composites Science and Technology **67** (5), 916-921 (2007).
65 E. J. Mele, "Commensuration and interlayer coherence in twisted bilayer graphene," Physical Review B **81** (16), 161405 (2010).
66 Hongki Min and A. H. MacDonald, "Chiral decomposition in the electronic structure of graphene multilayers," Physical Review B **77** (15), 155416 (2008).

67 Kohei Mizuno, Juntaro Ishii, Hideo Kishida, Yuhei Hayamizu, Satoshi Yasuda, Don N. Futaba, Motoo Yumura, and Kenji Hata, "A black body absorber from vertically aligned single-walled carbon nanotubes," Proceedings of the National Academy of Sciences **106** (15), 6044-6047 (2009).
68 Alpha T N'Diaye, Tim Gerber, Carsten Busse, Josef Mysliveček, Johann Coraux, and Thomas Michely, "A versatile fabrication method for cluster superlattices," New Journal of Physics **11** (10), 103045 (2009).
69 M. V. Nikolic, S. M. Radic, V. Minic, and M. M. Ristic, "The dependence of the work function of rare earth metals on their electron structure," Microelectronics Journal **27** (1), 93-96 (1996).
70 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** (7065), 197-200 (2005).
71 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** (5696), 666-669 (2004).
72 K. S. Novoselov, D. Jiang, F. Schedin, T. J. Booth, V. V. Khotkevich, S. V. Morozov, and A. K. Geim, "Two-dimensional atomic crystals," Proceedings of the National Academy of Sciences of the United States of America **102** (30), 10451-10453 (2005).
73 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** (5817), 1379-1379 (2007).
74 A. Oberlin, M. Endo, and T. Koyama, "Filamentous growth of carbon through benzene decomposition," Journal of Crystal Growth **32** (3), 335-349 (1976).
75 L. G. Parratt, "Surface Studies of Solids by Total Reflection of X-Rays," Physical Review **95** (2), 359 (1954).
76 Reginald M. Penner, "Mesoscopic Metal Particles and Wires by Electrodeposition," The Journal of Physical Chemistry B **106** (13), 3339-3353 (2002).
77 L. Pietronero, S. Strässler, H. R. Zeller, and M. J. Rice, "Charge Distribution in c Direction in Lamellar Graphite Acceptor Intercalation Compounds," Physical Review Letters **41** (11), 763 (1978).
78 L. Pietronero, S. Strässler, H. R. Zeller, and M. J. Rice, "Electrical conductivity of a graphite layer," Physical Review B **22** (2), 904 (1980).
79 B. Radisavljevic, A. Radenovic, K Brivio, V Giacometti, and A Kis, "Single-layer MoS2 transistors," Nat Nano **6** (3), 147-150 (2011).
80 Lorenz Ratke and Peter W. Voorhees, *Growth and Coarsening: Ostwald Ripening in Material Processing*. (Springer, 2002).

81 S. Reich, J. Maultzsch, C. Thomsen, Ordej, oacute, and P. n, "Tight-binding description of graphene," Physical Review B **66** (3), 035412 (2002).
82 Ludwig Reimer, "Scanning Electron Microscopy - Physics of Image Formation and Microanalysis", in *Scanning Electron Microscopy - Physics of Image Formation and Microanalysis* (Springer-Verlag, Berlin, 1998), pp. 145.
83 Alfonso Reina, Xiaoting Jia, John Ho, Daniel Nezich, Hyungbin Son, Vladimir Bulovic, Mildred S. Dresselhaus, and Jing Kong, "Large Area, Few-Layer Graphene Films on Arbitrary Substrates by Chemical Vapor Deposition," Nano Letters **9** (1), 30-35 (2008).
84 S. A. Safran and D. R. Hamann, "Electrostatic interactions and staging in graphite intercalation compounds," Physical Review B **22** (2), 606 (1980).
85 R Saito, G Dresselhaus, and Mildred S. Dresselhaus, *Physical Properties of Carbon Nanotubes*. (Imperial, London, 1998).
86 Bin Shan and Kyeongjae Cho, "First Principles Study of Work Functions of Single Wall Carbon Nanotubes," Physical Review Letters **94** (23), 236602 (2005).
87 Brian W. Smith and David E. Luzzi, "Electron irradiation effects in single wall carbon nanotubes," Journal of Applied Physics **90** (7), 3509-3515 (2001).
88 Brian W. Smith, Marc Monthioux, and David E. Luzzi, "Encapsulated C60 in carbon nanotubes," Nature **396** (6709), 323-324 (1998).
89 Cristian Staii, Alan T. Johnson, Michelle Chen, and Alan Gelperin, "DNA-Decorated Carbon Nanotubes for Chemical Sensing," Nano Letters **5** (9), 1774-1778 (2005).
90 Y. W. Tan, Y. Zhang, K. Bolotin, Y. Zhao, S. Adam, E. H. Hwang, S. Das Sarma, H. L. Stormer, and P. Kim, "Measurement of Scattering Rate and Minimum Conductivity in Graphene," Physical Review Letters **99** (24), 246803 (2007).
91 H. Trinkaus and B. N. Singh, "Helium accumulation in metals during irradiation - where do we stand?," Journal of Nuclear Materials **323** (2-3), 229-242 (2003).
92 Peter Trucano and Ruey Chen, "Structure of graphite by neutron diffraction," Nature **258** (5531), 136-137 (1975).
93 P. W. Voorhees, "The theory of Ostwald ripening," Journal of Statistical Physics **38** (1), 231-252 (1985).
94 P. R. Wallace, "The Band Theory of Graphite," Physical Review **71** (9), 622 (1947).
95 Mark Winter, "WebElements Periodic Table of the Elements | Ytterbium | biological information", (WebElements Ltd., 2010), Vol. 2010.
96 Stanislaus S. Wong, Ernesto Joselevich, Adam T. Woolley, Chin Li Cheung, and Charles M. Lieber, "Covalently functionalized nanotubes as

nanometre- sized probes in chemistry and biology," Nature **394** (6688), 52-55 (1998).
97 Bin Xue, Ping Chen, Qi Hong, Jianyi Lin, and Kuang Lee Tan, "Growth of Pd, Pt, Ag and Au nanoparticles on carbon nanotubes," Journal of Materials Chemistry **11** (9), 2378-2381 (2001).
98 Andrea F. Young and Philip Kim, "Quantum interference and Klein tunnelling in graphene heterojunctions," Nat Phys **5** (3), 222-226 (2009).
99 Min-Feng Yu, Oleg Lourie, Mark J. Dyer, Katerina Moloni, Thomas F. Kelly, and Rodney S. Ruoff, "Strength and Breaking Mechanism of Multiwalled Carbon Nanotubes Under Tensile Load," Science **287** (5453), 637-640 (2000).
100 B. L. Zhang, C. Z. Wang, K. M. Ho, C. H. Xu, and C. T. Chan, "The geometry of large fullerene cages: C[sub 72] to C[sub 102]," The Journal of Chemical Physics **98** (4), 3095-3102 (1993).
101 Yian-Biao Zhang, Mandakini Kanungo, Alexander J. Ho, Paul Freimuth, Daniel van der Lelie, Michelle Chen, Samuel M. Khamis, Sujit S. Datta, A. T. Charlie Johnson, James A. Misewich, and Stanislaus S. Wong, "Functionalized Carbon Nanotubes for Detecting Viral Proteins," Nano Letters **7** (10), 3086-3091 (2007).