Cambridge University Press & Assessment
978-0-521-87700-8 — Nanostructures and Nanotechnology
Douglas Natelson
Excerpt
More Information

1 Introduction and overview

Nanotechnology has seized both the public and scientific communities' imaginations, and it's not hard to see why. From K. Eric Drexler's and Ray Kurzweil's visions of self-reproducing engineered nanomachines building macroscale structures out of single crystal diamond and swimming through our capillaries repairing damaged cells, to talk of building an elevator directly to geosynchronous orbit, the promise of nanoscale science and engineering has been described to an enthusiastic public. Billions of dollars of research funding have been directed into this area, and further billions of dollars are already being spent on commercial products that are self-described as examples of nanotechnology. Estimates of the global economic impact of nanotechnology in the next ten years exceed $1 trillion, and are rising.

How seriously should we take these exciting visions? Is nanotechnology a "disruptive technology", a distinct advance that will reshape the world? What are the *real* potential impacts of the ability to engineer the structure and composition of materials on the nanometer scale? What are the limitations imposed by Nature (that is, physics and chemistry) on what is possible? What physical principles become relevant at small scales that on the one hand set limitations, but on the other provide new opportunities? What scientific questions remain to be answered at the nanometer scale, and why? How can people manipulate and characterize materials on these scales?

Hopefully this book will help you answer these questions, or at least give you a good idea of what to consider when trying. Even though the nanometer scale is very different from our everyday experience, it is possible to develop an intuition about how matter will behave at that extremely small level.

1.1 What is nanotechnology?

Nanotechnology is an extraordinarily broad term. According to the US government,[1] nanotechnology is "Research and technology development at the atomic, molecular or macromolecular levels, in the length scale of approximately 1 – 100 nanometer range, to provide a fundamental understanding of phenomena and materials at the nanoscale and to create and use structures, devices and systems that have novel properties and functions because of their small and/or intermediate size." That's quite a mouthful.

[1] See www.nano.gov.

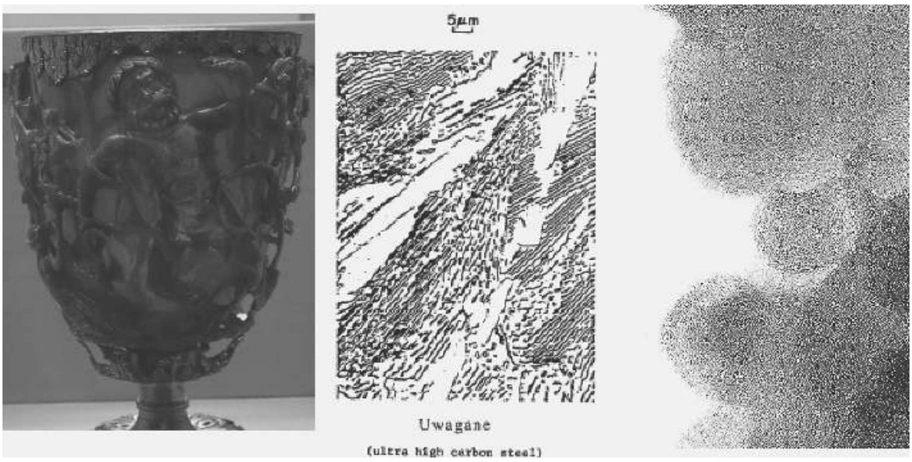

Fig. 1.1 Nanotechnology in history. Left: the Lycurgus cup, a Roman goblet made using glass containing colloidal gold nanoparticles (British Museum [1]). Center: an electron micrograph of micro- and nanostructured steel, *uwagane*, of the type used to produce Japanese *katana* swords. From [2]. Right: carbon black nanoparticles used to improve the stiffness and strength of rubber (Electrochemistry Encyclopedia [3]).

Colloquially, nanotechnology has come to encompass just about any technology that involves controllably engineering structures with at least some critical dimension less than 100 nm in extent. By that definition, we have been living in the Age of Nanotechnology for hundreds of years (see Fig. 1.1). Colored Roman glass and some medieval stained glasses owe their remarkable optical properties to embedded colloidal metal particles on that scale. Damascene steel and Japanese *katana* sword steel are made by elaborate processes that end up producing nanostructured materials with impressive mechanical properties. Metal films for coating other metals or glass (mirrors) have long been produced with thicknesses less than 100 nm. Carbon black, the sooty stuff used as an additive to enhance the mechanical properties of rubber, contains particles as small as a few nanometers.

What distinguishes current efforts from these past achievements are understanding and control. In the past few decades, there have been tremendous advances in the chemistry and physics of materials relevant at the nanometer scale. Simultaneously and symbiotically, new tools have been developed that allow people to “see” what is happening at that scale with unprecedented precision. Computational prowess has also grown astronomically over the same time period, enabling calculations that formerly would have been impractical. The net result of all this is that people now have the ability to design, engineer, and construct a vast array of structures on the nanometer scale to take advantage of the unique properties of matter when tailored on those dimensions.

One reason it is difficult to define nanotechnology is that much nano research takes place at the boundaries between traditional disciplines. For example, consider a single recent innovation: the use of semiconductor nanocrystals as markers for particular biological processes. The unique optical properties of the nanocrystals result directly from physics, specifically the impact of finite crystal size on the allowed electronic states in the material. The nanomaterial itself is made using a batch synthesis process in solution at moderate temperatures in a typical chemistry laboratory fume hood. Surface chemistry is exploited to coat the nanocrystals with a surfactant layer that prevents them from aggregating and falling out of solution. The nanocrystals are characterized using standard materials science tools such as transmission electron microscopy (TEM) and x-ray photoemission spectroscopy (XPS). Biochemistry expertise is required to get the nanocrystals to bind appropriately in cells, and biological microscopy techniques are needed to do the actual sensing of processes in living systems. In this respect nanotechnology is truly an interdisiplinary endeavor.

1.2 Sizes of things

It helps to have a "big picture" view of the very small so we can see what's likely to be relevant. Just how small *is* the nanometer scale? Figure 1.2 is an example of the typical answer to this question, trying to bridge the realm of everyday experience with the nanoscale. Rather than just looking at the relative sizes of things, let's consider the figure in terms of the physics relevant in the different size regimes.

At the comparatively macroscopic scale is a human hair, on the order of 100 μm in diameter. From the physics perspective, this object and its interactions with the world are

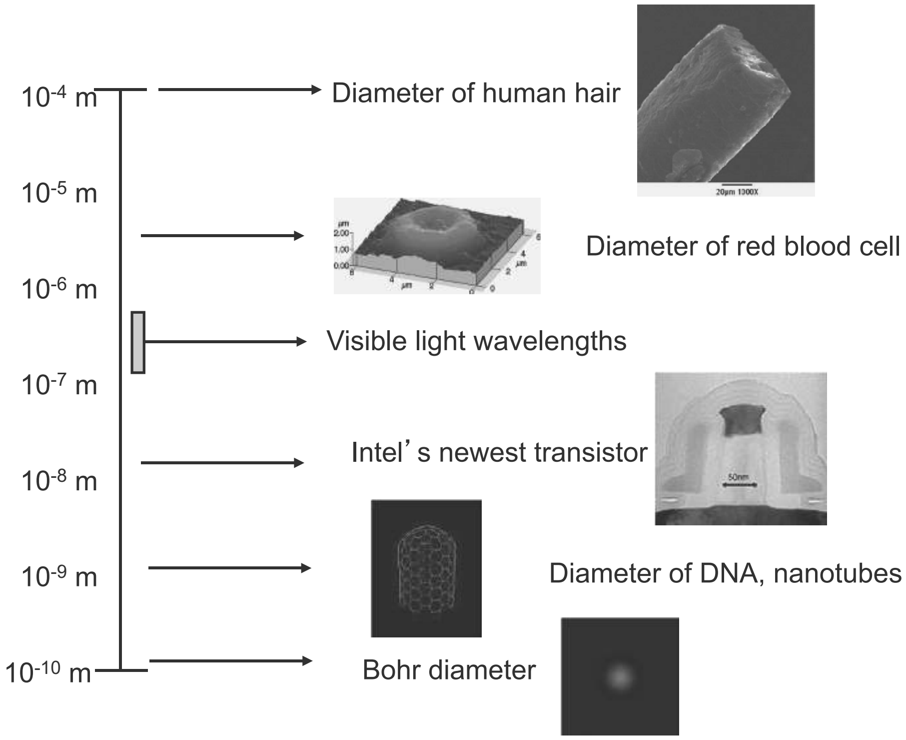

Sizes of things.

Fig. 1.2

"big". It contains a large number, $\sim 10^{21}$, of atoms, bound together at the molecular level into macromolecules called proteins, primarily of a single type, keratin.

To describe the mechanical properties of the hair, no one would ever consider solving the quantum mechanical or even classical equations of motion for all the consituent atoms; that would be computationally untenable, and manifestly unwise. Fortunately, it is also unnecessary from a practical standpoint. The atoms, locked in the protein molecules, interact with each other strongly enough and in the right ways to take on certain *emergent properties* that describe the mechanical response of the whole ensemble. It makes sense to talk about an average mass density, and an elastic modulus that relates the elongation of the hair to the average internal forces in the hair – related to the Hooke's Law spring constant you think of from first-year physics. If the hair is pulled tight and plucked, the intelligent way to think about the collective motion of all the atoms is in the form of sound waves with wavelengths spanning many thousands of atoms.

Similarly, no one would ever consider describing the optical properties of the hair starting from the optical transitions known from gas phase spectroscopy of the component atoms. Instead, the ensemble of molecules collectively has optical properties that can be succinctly understood as some (complex) index of refraction.

These collective properties and their emergence in "large" systems are the intellectual core of *condensed matter* physics. Even when the underlying rules governing the properties and interactions of individual objects are simple, the ensemble can show collective properties that are profound, surprising, and extremely rich. As Nobel laureate physicist Phil Anderson said in his now famous 1972 Science paper, "More is *different*".

Moving down the ladder, we come to a single red blood cell a few microns in diameter. It's not too inaccurate to describe this cell as a bag made from a lipid bilayer membrane, filled with a bunch of liquid. Even though the membrane is only a couple of molecules thick, it is quite extended and has some features of macroscopic objects, like an elastic modulus within the layer. The mostly-water inside the cell is also pretty fairly described by "continuum" averaged quantities, like a mass density, a viscosity that describes dissipation within the fluid, an index of refraction, ionic concentrations, etc. Except in some very reductionist sense, quantum mechanics is not necessary to think about in understanding these properties.

We now pass through the wavelengths of visible light. While electromagnetic radiation comes in quanta called photons, it is still often useful to consider the continuum (huge number of photons) limit of thinking about light as propagating waves of electric and magnetic fields. There is nothing particularly special about this length scale, except that it corresponds to photons with energies of about 1–2 eV that our eyes are photochemically adept at detecting. Generally it is challenging to resolve (via scattered light) objects separated by distances much smaller than the wavelength of light being used for illumination. In short, other tools are needed to "see" objects smaller than this.

The next item down is a cross-sectional image of a silicon metal-oxide-semiconductor field-effect transistor (MOSFET), the basis for much of modern electronics. We will discuss these in detail later. The example shown here is state-of-the-art as of 2006, with a critical spacing between incoming (source) and outgoing (drain) electrodes of 50 nm, and a crucial insulating layer only 1.2 nm thick. In 2006 there were $\sim 10^{8}$ of these in a typical high end microprocessor. Clearly the semiconductor industry has been working at the nanoscale for some time! Interestingly, even with dimensions such as these, the material properties are still well described as "bulk-like"; the operating characteristics of the MOSFET are not terribly different from those of appropriately scaled much larger devices.

This would no longer be the case, however, if the Si were structured on the 1–2 nm scale, the next stop on the road to the atomic regime. As indicated in the figure, this is approximately the diameter of an individual single-walled carbon nanotube, or the double helix of deoxyribonucleic acid (DNA). At this size scale it is often no longer appropriate to use continuum quantities to understand the mechanical, electronic, and optical properties of materials, and quantum mechanical details begin to matter. A significant portion of this book will examine why this is so, and the implications.

Finally, at the bottom end of the size scale for neutral, stable matter, we are left with individual atoms. For hydrogen, the Bohr diameter is approximately 0.1 nm. As should be familiar from chemistry and physics classes back to high school, the electronic and optical properties of single atoms are completely dominated by quantum mechanics.

If we consider the surface of the hair on this kind of distance scale, it is extremely lumpy and inhomogeneous. If the hair is exposed to ambient atmosphere, it is undoubtedly coated with a physisorbed layer of water, as well as some amount of small molecular weight hydrocarbon compounds. Further, its surface is in constant motion, with adsorbed contaminants reshuffling themselves, bonds vibrating at terahertz frequencies, ions and solvated electrons swapping charge over a broad distribution of time scales, gas molecules from the air impinging at a rate high enough that each atom on the surface is hit on average once a nanosecond. At the single nanometer level, apparently clean surfaces are often dirty, and apparently quiescent equilibrium is alive with activity.

Somewhere between the macroscale and the atomic scale, we have passed from the classical, bulk world of Newton's laws, continuum elasticity theory, and fluid mechanics. Instead we have entered the realm of quantum mechanics, interfacial effects, local fluctuations and deviations from equilibrium, and the breakdown of continuum approximations. This transition happens because as a system's size, L, is reduced, length scales relevant for particular physical processes shift from being negligibly small compared to L to being comparable to or much larger than L. In the next section, we will see just how this works in a bit more detail, and introduce several of these physically motivated lengths, using a particular example familiar, at least crudely, from everyday experience.

1.3 Important length scales: breaking a wire

What happens when a metal wire breaks? As a concrete example, let's consider a gold wire with a diameter $d = 25$ μm. The choice of gold already simplifies our thought experiment: gold is ductile, so we won't have to worry about crack nucleation and propagation, and bulk gold is chemically inert under ambient conditions, so we needn't be concerned about oxidation. To the naked eye, nothing particularly interesting or surprising occurs. When a tensile force is applied to the ends of a wire of length L, the wire deforms, first elastically (so that when the force is removed the wire springs back to its original length), then plastically (irreversibly). Under continued pulling, a "neck" forms somewhere along the wire, where the plastic deformation reduces the local diameter. The deformation speeds up, the neck shrinks further, and eventually the wire breaks completely, popping into two separate pieces. If one were to measure the electrical conductance through the wire every 0.1 s by applying a known voltage, V, across it and measuring the resulting current, I, one would see that the conductance, $G \equiv \mathrm{d}I/\mathrm{d}V$ remains fairly constant, perhaps decreasing slightly before dropping to zero after the wire breaks.

1.3.1 Length scales in play

Examining the wire and the breaking process on smaller length scales and more finely in time is very revealing, and serves as a guide to some of the physics that we will see later. First, consider the microstructure of the wire before the breaking process. The wire is polycrystalline, composed of many individual Au grains, typically 20–40 nm in size. Each grain consists of Au atoms stacked in a close-packed arrangement called face-centered cubic (FCC), with atom centers separated by a *lattice parameter*, $a \approx 0.4$ nm. Generally the grains are randomly oriented, so that the surface of the wire consists of a collection of different crystallographic faces.

The detailed crystal structure is determined in part by the electronic interactions between the atoms. It turns out not to be too bad an approximation to think of the metal as consisting of a lattice of FCC-stacked ion cores, with each ion having a positive charge, surrounded by an electronic "fluid" with a density of around 5.9×10^{22} electrons per cm^3. Even though the electrons relevant to electronic conduction in the wire are negatively charged and therefore interact repulsively, we can often get away with ignoring their interactions. One reason for this is that, by rearranging themselves, the electrons can *screen* excess charge on a distance scale of roughly the *Thomas–Fermi screening length*, $r_{\mathrm{TF}} \approx 0.4$ nm in this system.

One can think of these electrons "semiclassically",[2] and ask how far a typical electron travels before scattering – that is, before participating in some process that reorients its direction of propagation. Processes that do not change the energy of the electron are called elastic, while those in which the electron gains or loses energy to some other degree of freedom in the material are inelastic. The typical distance traveled by an electron before such a scattering process is called the *mean free path*, ℓ. As we shall see later, any break in the periodicity of the stacking of the atoms (vacancies, grain boundaries, surfaces) can cause elastic scattering, In the absence of inelastic scattering, the grain size mentioned above is a reasonable estimate of ℓ_{e}, the elastic mean free path.

The quantum mechanical nature of the conduction electrons means that they have wave-like properties. The effective wavelength of the electrons relevant for conduction, the *Fermi wavelength*, is around $\lambda_{\mathrm{F}} \approx 0.5$ nm in bulk Au.[3] Like all wave phenomena, the electronic waves have some phase. Waves that are in phase at a given position interfere constructively, while those that are out of phase by π radians interfere destructively. These interference

[2] We'll learn about the details of this in Chapter 2.
[3] As we shall see, the similar magnitudes of a, r_{TF}, and λ_{F} are not coincidental.

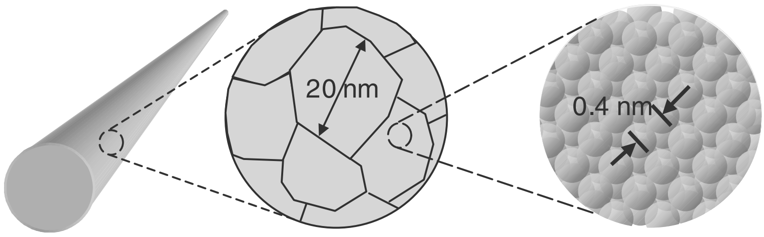

Fig. 1.3 A gold wire that seems roughly homogeneous is really polycrystalline with a grain size of 20 nm, and each grain is a face-centered cubic crystal with atomic separation of 0.4 nm.

effects, which have no classical analog, are generally undetectable on macroscopic length scales because of *decoherence*. Inelastic interactions with environmental degrees of freedom effectively randomize the phases of the electronic waves on a distance scale, L_ϕ, the *coherence length*. At room temperature in Au, $L_\phi \sim 1-2$ nm.

Some of these same inelastic interactions, including electron–electron and electron–vibrational scattering, are what establish thermal equilibrium for the electrons. There is therefore a related distance scale, L_{i}, the inelastic length, over which a nonthermal distribution of electrons equilibrates to the temperature of the system as a whole. At room temperature it's not too bad to assume $L_{\mathrm{i}} \sim L_\phi$, though this can change at low temperatures.

Finally, for completeness recall that electrons possess intrinsic angular momentum called "spin". In the absence of magnetic correlations like ferromagnetism, the spins point in no particular direction. One can ask, if a particular electron is prepared with its spin aligned along some known direction, how far does the electron propagate before some process effectively randomizes the direction of the electron's spin. In gold, there is strong spin–orbit scattering; that is, from the point of view of the moving electron, the electric field from the gold ion cores looks like an effective magnetic field that causes the spin to precess. The *spin–orbit scattering length* in this material is around $L_{\mathrm{SO}} \approx 30$ nm. On distances short compared to L_{SO}, it's not bad to think about the spin of an electron as a well defined quantity, while on longer distance scales, the electron trades angular momentum back and forth with the lattice.

So, the stage is set. Before the breaking begins, the hierarchy of lengths is:

$$L \gg d \gg \ell_e \sim L_{\mathrm{SO}} > \ell_i \sim L_\phi \sim L_{\mathrm{i}} > \lambda_{\mathrm{F}} \sim r_{\mathrm{TF}} \sim a. \tag{1.1}$$

There are actually two more length scales relevant during the breaking process, as shown in Fig. 1.4: δ, defined here as the rough diameter of the most constricted part of the "neck", and λ, the rough length of that region. Before the breaking process, it doesn't make sense to think about δ and λ, since there is no constriction.

1.3.2 The breaking process up close

Start pulling on the ends of the wire. The tensile force is distributed across the wire cross-section. Pull hard enough, and somewhere within the wire, grains start to change

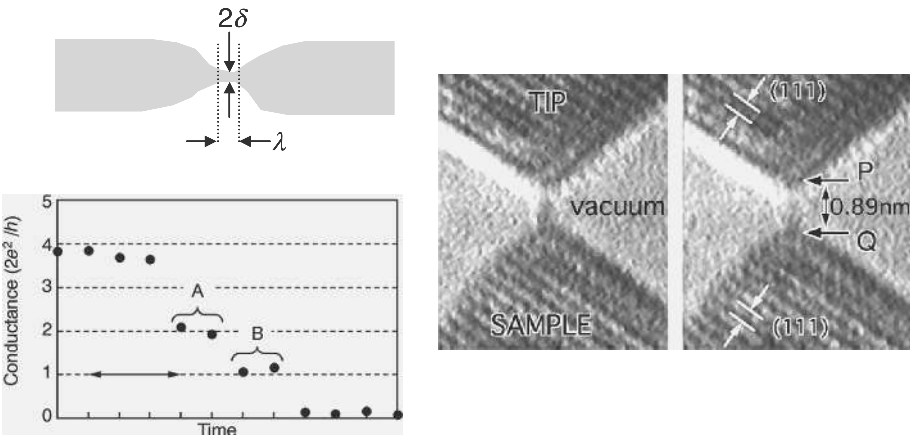

Fig. 1.4 Breaking a gold wire. Top left: definitions of λ and δ, the characteristic length scales of the necking constriction. Bottom left: conductance as a function of time as a gold wire is pulled apart. Right: *in-situ* transmission electron micrographs of the end of the breaking process, showing the formation of a one-dimensional chain of Au atoms just before complete breaking. Adapted from [4].

irreversibly. Atoms at the boundaries between grains are less constrained than those within grains, so they are more free to move in response to local forces. Defects within the grains (vacancies, interruptions in the fcc stacking called dislocations) can also move, propagating to grain boundaries under the applied forces. Once the cross-section somewhere along the wire is reduced below $\pi d^2/4$, the local tensile *stress* increases, since the same applied force can only be spread over a smaller area.

Interesting things start to happen as δ and λ shrink, and trade places in the hierarchy of Eq. (1.1) with the intrinsic scales of different physical processes. The first of these crossovers occurs when λ (on the order of δ for such a ductile material) and δ become smaller than a typical grain size, and therefore smaller than ℓ_e. If a slight voltage bias is applied between the two ends of the wire, one finds that the conductance of the whole system is now limited by the constriction. Electrons passing through the narrow neck are more likely to scatter elastically from the boundaries of the neck than within the neck itself. Similarly, since $\lambda < L_{SO}$, the spin of the electrons remains approximately conserved when traversing the neck.

When $\delta, \lambda \to \ell_i$, another crossover occurs. Now the constriction is so short and narrow that electrons that traverse it do so without any significant scattering taking place in the constriction itself. This is the *ballistic* or *point contact* regime, and raises some interesting questions. The neck now strongly limits the total conduction through the wire, but does so even though electrons pass through that region ballistically. What is the origin of the constriction's resistance? As we shall see in Chapter 5, a ballistic constriction's resistance in this situation is an example of a *contact resistance*, rooted in the difficulty of getting electrons to enter the ballistic region in the first place. If there is no scattering within the constriction, yet the constriction has some resistance R_c, where is the power dissipated from the simple I^2R_c Joule heating? We shall find that the dissipation takes place where there is inelastic scattering, typically a distance L_{i} away from the constriction itself in the bulk wire material. This implies that when a voltage is applied to the wire, the electrons in the constriction are not in thermal equilibrium; this complicates the detailed description of such junctions. Indeed, once $\lambda < L_{\phi}$, properly treating conduction through such a constriction requires consideration of both nonequilibrium effects and quantum interference corrections to the conduction.

The real excitement comes when $\delta \rightarrow \lambda_{\mathrm{F}}, a$, when the diameter of the constriction at its narrowest is only a couple of lattice spacings and (effective) electron wavelengths. The wave character of the electrons becomes of paramount importance. The number of electronic "modes" that span from one side of the constriction to the other is a handful, and depends on the detailed atomic configuration at the junction between the two wire pieces. Measurements of the conductance while passing through this regime reveal discrete steps in G as the pulling continues, corresponding to particular arrangements of the junction atoms. The exact G values vary from junction to junction, but a histogram reveals that the conductance in Au often changes by integer multiples of $G_0 \equiv 2e^2/h$, where e is the charge of the electron (-1.602×10^{-19} C), and h is Planck's constant (6.63×10^{-34} Js). This is *conductance quantization*, and is an example of electronic conduction dominated entirely by quantum mechanics that can be seen at room temperature on a benchtop!

Unsurprisingly, in this limit of a few-atom constriction the details of the particular chemistry of the wire material become very important, and it is no longer sensible to describe the mechanical properties of the neck in terms of the elastic characteristics of bulk Au. Experiments have found that the final stages of neck deformation in breaking Au wires lead to remarkable changes in neck structure. The Au atoms take on certain particular configurations very different from the bulk fcc arrangement, including helical stacking and extended chains only one atom in diameter! In configurations like these, the chemical reactivity, thermodynamic properties, and even the detailed electronic structure of the metal can change; every atom is at a surface, with different chemical bond arrangements than in the bulk.

Finally, one critical Au–Au bond in the constriction breaks, and the wire becomes two separate pieces. The last steps of the breaking process take tens of microseconds, and are followed by complicated reconstructions of the wire surfaces as the applied forces on the freshly exposed wire ends drop precipitously to zero. As this simple example shows, even something as conceptually simple as pulling a wire into two pieces reveals a wealth of physical processes, including rich quantum mechanical effects, once critical system scales pass into the nanometer regime.

1.4 The structure of this book

This text is the outgrowth of teaching a popular two-semester graduate course sequence at Rice University. Students in these classes tended to be first-year graduate students or senior undergraduates, with backgrounds including physics, physical chemistry, electrical engineering, and materials science. This text is intended for a similar audience, though it is written more from a physics point of view largely because of the background of the author. A familiarity with quantum mechanics and statistical physics ideas is very helpful. Many students in the courses had never taken a solid state course of any kind, but they had seen the Schrödinger equation before.

There are two main challenges in approaching such a broad audience in a rapidly moving field: establishing a common foundation of knowledge from which to proceed, and formulating (tractable!) exercises and problems that convey fundamental concepts and ideas. I have tried to do both, at sufficient depth for students to develop an appreciation for and an intuition about nanoscale science, though by necessity individual subtopics (*e.g.* photonics, microfluidics) cannot be covered at a level similar to that of a focused graduate course. In each chapter I try to provide references to other more specialized works – review articles, textbooks, web resources – to enable readers to pursue their particular interests in more detail.

Exercises are a mixture of calculations, derivations, and short response writing. As anyone who has studied quantum mechanics, statistical mechanics, or solid state physics can attest, the number of exactly solvable homework problems is limited. As a result many of the exercises here are comparatively simple, and are meant to stimulate critical thought about major concepts rather than turn students into theorists capable of calculational *tours de force*. The short answer questions, while challenging to grade, have proven very worthwhile. Students need to be able to read critically and write concisely, and the more practice at this the better.

The structure of this book is straightforward. I assume reasonable familiarity with undergraduate level quantum mechanics from either the physics or physical chemistry perspectives. A few particularly important ideas (time independent and time dependent perturbation theory; Fermi's Golden Rule) are reviewed in an appendix. Some background in solid state physics is not essential, but wouldn't hurt.

Without an understanding of why macroscopic materials have the properties they do, it's very hard to appreciate what is different about nanostructured materials. Therefore, Chapter 2 is a lengthy review of solid state physics, with particular emphasis on properties and processes that are likely to be affected by nanostructuring. This chapter is not meant to be an exhaustive treatment or a substitute for a two-semester solid state course, but should give readers of various backgrounds an overview of the concepts and the beginnings of an intuition about solids at the nanoscale. This will include a comparison and contrasting of, for lack of a better description, the physicist's and chemist's approaches to the solid state.

Following this discussion, Chapter 3 briefly examines several of the most common material systems relevant for nanotechnology. In addition to common bulk inorganic materials such as silicon and gallium arsenide, specific attention is given to nanostructured semiconductors, complex and strongly correlated oxides, magnetic semiconductors, and fullerene nanomaterials. A more complete discussion of magnetic materials is left for Chapter 7.

Chapter 4 provides a very abbreviated overview of various techniques commonly used for the fabrication and characterization of materials down to the nanometer scale. Fabrication methods include top–down lithographic patterning, material deposition and removal approaches, and more bottom–up, self-assembly techniques. Characterization