

Chapter 1

Introduction to Multilayered Nanostructures

On December 29, 1959, Richard Feynman addressed physicists at the banquet of the annual meeting of the American Physical Society. The title of his talk was “There’s plenty of room at the bottom” [Feynman (1961)]. There was much conjecture amongst the audience as to what a talk with that title would be about, but Feynman kept it secret. When he delivered his speech, Feynman described the new field of nanotechnology, although he did not coin that term. He described how one could write all of the information published in all the books in the world on the head of a pin using manipulation of atoms in three dimensions. At the time, the talk seemed to be more science fiction than fact (see Chapter 4 of [Regis (1995)] for a historical account), even though the scientific press published many articles about the presentation; the field of nanoscience has only blossomed since the early 1990s and now there are many devices that work with or manipulate the properties of individual atoms, molecules, or small groups of atoms or molecules.

The semiconductor industry has been reducing the size of structures in its microprocessors at a rapid rate; they now create line features and transistors that are smaller than 100 nm. Current research on quantum dots treat quantum-mechanical boxes that contain a few hundred to a few thousand electrons in a small spatial region. Fabrication techniques have become so sophisticated that novel devices can be made that involve the transport of current through single molecules trapped between metallic electrodes. The discovery of conducting carbon nanotubes has provided the nano world with a possible electrical wiring system. It is clear that the future will hold many surprises and technological advancements coming from nanotechnology.

As device features are made smaller and smaller, in particular, as they become on the order of a few atoms (or nanometers) in size, quantum-

mechanical effects begin to take over, and ultimately determine the device performance. It is the job of theorists to understand how to explain, model, and design devices when quantum-mechanical effects cannot be ignored. In this book we discuss one particular kind of nanotechnology—the field of multilayered nanostructures, which are composed of stacked atomic planes of different materials, with the thickness of some of the layers in the nanometer regime. Usually these devices are operated by attaching them to a voltage (or current) source, which transports electrical or heat current perpendicular to the stacked planes.

The approach and focus of this book are different from those of others. Most work on nanostructures focuses on devices that are small in all (or all but one) dimensions, so it is appropriate to start from an atomic or molecular picture and build up to the nanoscale devices (like quantum dots or wires). This class of nanoscale devices usually have strong surface effects, because the surface-to-volume ratio is usually large. Here we take an alternative “top-down” approach as opposed to the more traditional “bottom-up” approach, and consider systems in the thermodynamic limit that have only one dimension on the nanoscale (more precisely only one dimension has nanoscale inhomogeneity). This allows us to employ dynamical mean-field theory to solve the many-body problem because this technique is accurate when the number of nearest neighbors for each lattice site is large. In a multilayered nanostructure, there are no surfaces, so every lattice site maintains approximately the same number of neighbors as in the bulk. Furthermore, multilayered nanostructures are already being employed in technology, and are easier to manufacture and to use in devices than systems that are nanoscopic in all dimensions. Hence, it is likely that most applications that are commercially viable will involve multilayered nanostructures (at least for the not-too-distant future). Indeed, this is the motivation for producing this work.

1.1 Thin Film Growth and Multilayered Nanostructures

Multilayered nanostructures are the most common electronics devices that have at least one length scale in the nano realm. They have been in use for over five decades! The original devices are based mainly on semiconductors and the so-called pn junction. But research has been performed on superconducting variants for over four decades, and there are commercial devices in use for niche markets.

Electronics devices often rely on nonlinearities to function. Either it is the nonlinear current-voltage relation that determines the functionality of the device (like in a pn junction where current flows in essentially one direction), or it is the avalanche breakdown, or other nonlinear behavior, that ultimately determines when the device ceases to work. The classic multilayered nanostructure is a tunnel junction, consisting of a sandwich of two metallic electrodes separated by a thin layer of insulator. They can be easy to manufacture if the insulator is formed by exposing the metal surface to air (or other oxygen containing gas mixtures like oxygen and argon) where a native oxide layer will form. Since the two metallic regions are connected by a “weak link” due to the proximity or tunneling effect (described in Section I.3), the connection is inherently due to quantum-mechanical effects and the uncertainty principle: electrons in the metal cannot remain localized within the metal, but can leak through the barrier into the other metal. If the electrodes are superconducting and the barrier is thin enough, then the device is a Josephson junction.

A quantum-mechanical wavefunction is highly nonlinear. In classically allowed regions, it will oscillate and have nodes, while in classically forbidden regions, it will exponentially decay. Both behaviors are nonlinear, and ultimately lead to the nonlinear behavior of multilayered nanostructures. We will not discuss nonlinearities much in this work, but we mention this fact to remind the reader that whenever quantum-mechanical behavior governs the transport through a device, it is likely to have some underlying nonlinear features. Tuning and controlling these nonlinear features is often necessary to make the device useful. Examples of nonlinear current-voltage characteristics in Josephson junctions are shown in Fig. 1.1.

Another useful feature in devices is controllability. Many semiconductor devices have a voltage gate which can be varied to change the behavior of the device. Strongly correlated materials (described in Section I.2) often have properties that can be sharply tuned by external fields, pressure or chemical doping, and provide an interesting alternative of materials to use in devices from the conventional metals, semiconductors, and insulators currently in use. They are of particular interest when one considers controlling the transport of the spin of the electron (so-called spintronics devices), since magnetism is inherently quantum mechanical in nature, and many strongly correlated systems also display interesting magnetic properties. But, due to their quantum-mechanical behavior, involving correlated motion of electrons, they are less well understood than semiconductors, and fewer devices have been made from them. At the moment they hold




Fig. 1.1 Current-voltage curve for (a) a hysteretic Josephson junction and (b) a non-hysteretic Josephson junction. The bottom figure is a schematic of a Josephson junction which corresponds to a superconductor-barrier-superconductor sandwich; the superconductor “leaks” through the barrier from one superconductor to the other carrying current with a nonlinear current-voltage relation. The Josephson junction can carry current at zero voltage up to the critical current I_c , and then it moves into a resistive state. If the current-voltage curve is multivalued (left panel), then it is a hysteretic junction, while a single-valued curve (right panel) corresponds to a nonhysteretic junction. Both curves ultimately join up to the linear curve of Ohm’s law ($I = V/R_n$) at high voltage (R_n is the normal-state resistance). The characteristic voltage where the current-voltage curve starts to become linear is $V_c \approx I_c R_n$ which is typically no larger than a few meV.

great promise and interest. This work hopes to aid with the design of novel devices that use strongly correlated materials by enabling one to calculate properties based on the underlying features of the materials that comprise the device.

Modern science has made great strides in its ability to artificially grow multilayered nanostructures. There are a number of different growth techniques that are used, and they each have their set of advantages and disadvantages. All growth processes start with a substrate material that is chosen either for the lattice match with the candidate material to be grown (to serve as a template and to relieve strain), for the chemical inertness with respect to the growth material (to reduce interdiffusion and creation of unwanted chemical species at the interface), or for practicality in sub-




Fig. 1.2 Transmission electron micrograph of a sputtered device for use in spintronics. The TEM image allows us to see individual atomic planes, and is able to discern the chemical composition of each layer. *Figure reprinted with permission from [Wang, et al. (2005)]* (©2005 American Institute of Physics).

sequent device processing. The ultimate goal of material growth is to lay down atomically flat planes of each desired material, one plane at a time, and modify the constitution of the growth planes as desired to make the device of interest. In reality, this is never fully achieved with any technique, but in current state-of-the-art device growth, it is possible to achieve almost atomic flatness of the epitaxial growth planes, and in some cases the interface regions can be nearly atomically flat with limited interdiffusion or chemical reactions.

The simplest way to grow materials is via sputtering, which involves bombarding a target with inert ions, forcing the target atoms to be expelled and shower onto the substrate where the thin film will be grown (the word sputtering comes from the Greek verb *sputare* which means to spit). Sputtering is a simple growth process because one need not worry about the relative vapor pressures of the constituents, since the material grows in a nonequilibrium fashion. It also grows with the same stoichiometry as that of the target (essentially because the atoms that are emitted all come from the surface of the target). Sputtering is generally not believed to be able to grow atomically sharp interfaces, and it can be difficult to guarantee uniform coverage during the growth process; its main advantages are

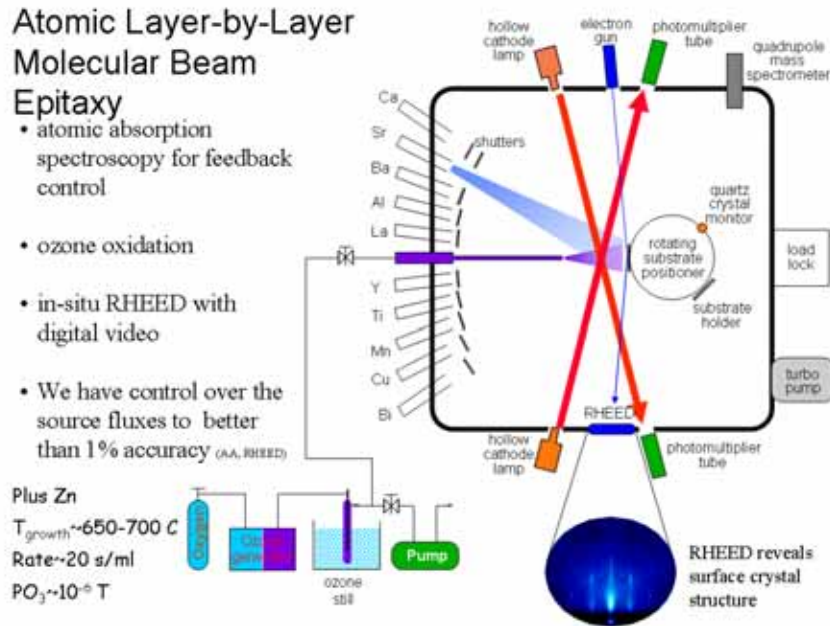


Fig. 1.3 Schematic of a molecular beam epitaxy growth chamber. The MBE growth takes place in ultra high vacuum. Different sources are introduced by opening shutters that allow the heated material to evaporate into the chamber. Many different means to characterize the sample during growth are possible. For example, RHEED oscillations show when a monolayer of growth is completed. *Figure adapted with permission from [Eckstein and Bozović (1995)].*

that it grows stoichiometrically and it is fast, so impurities may not have a chance to enter the device in high concentrations. It can achieve high quality growth, as illustrated in a spintronics device grown via sputtering that has nearly atomically flat interfaces for a variety of magnetic and nonmagnetic multilayers [Wang, *et al.* (2005)].

Molecular beam epitaxy (MBE) is arguably the most precise of the growing techniques. An MBE machine has growth conditions controlled to high precision. The growth chamber is inside an ultra high vacuum (UHV) chamber that has a sample holder and a series of growth materials inside separate furnaces; shutters in front of the furnaces open to allow the evaporated vapors of the different materials into the chamber, which will hit the sample and stick. A schematic of such a device is given in Fig. 1.3.

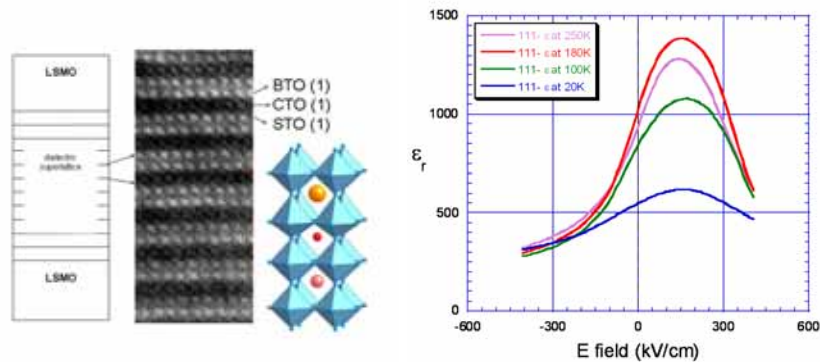
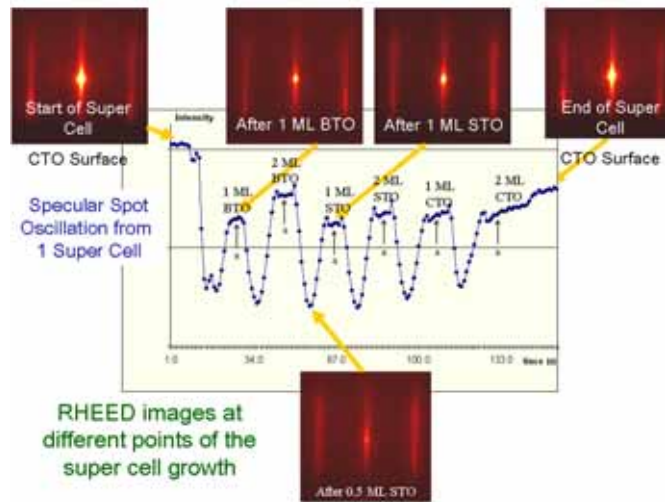


Fig. 1.4 (Top) RHEED oscillations during growth showing the completion of each monolayer. (Bottom) Left panel: schematic of the complex dielectric oxide formed from CaTiO_3 , BaTiO_3 , and SrTiO_3 along with a TEM image; right panel: dielectric response of different devices. *Figure adapted with permission from [Warusawithana, *et al.* (2003)] (original figure © 2003 the American Physical Society) and [Warusawithana, Chen, O'Keefe, Zuo, Weissman and Eckstein (unpublished)].*

The growth process can be monitored by RHEED oscillations which repeat as each atomic monolayer is placed down. The growth is usually slow, with perhaps a few seconds for each atomic layer. An example of the growth of an artificially engineered dielectric is given in Fig. 1.4. The top panel




Fig. 1.5 Schematic of a PLD system for growing MgB_2 . The magnesium and boron targets (heated up by the UV laser pulse) are supplemented by a so-called Knudsen (or effusion) cell which is an evaporator of a beam of magnesium to maintain high enough Mg pressure for stoichiometric growth. A residual gas analyzer monitors the gases in the chamber, where the growth takes place in vacuum. *Figure reprinted with permission from [Kim and Newman (unpublished)].*

shows the RHEED oscillations, while the bottom left panel is a TEM of the different layers (with a schematic of the device) and the bottom right is an example of the dielectric response as a function of the applied field.

Pulsed laser deposition (PLD) is another high precision growth technique. It involves ablating materials targets with a high power UV laser pulse, which creates a plume that is directed at the sample. The growth proceeds in spurts, in this fashion, and can achieve nearly atomic flatness, but it is not as common to monitor the layer-by-layer growth as in MBE. It is, however, typically much faster than growth in an MBE system, and has emerged as a popular choice for thin-film device growth in research laboratories because of its speed combined with its innate ability to preserve the target's stoichiometry. An example of a PLD system is shown in the schematic picture of Fig. 1.5. A trilayered $\text{TiNbN-Ta}_x\text{N-TiNbN}$ sample grown with PLD is imaged with a TEM in Fig. 1.6.




Fig. 1.6 TEM images of a trilayered TiNbN-Ta_xN-TiNbN sample suitable for processing into a Josephson junction. The sample was made with the PLD process. The left panel has the widest field of view, which is blown up in the upper right and then lower right images. Note that although the interfaces meander across the sample, the barrier width is quite uniform throughout the growth process. *Figure reprinted with permission from [Yu, et al. (2006)].*

Chemical vapor deposition (CVD) is a technique often used in industrial manufacturing. A series of different gaseous phases of materials are directed toward the sample, where a chemical reaction takes place at the surface, facilitating the growth. CVD is complicated by the need to find the right precursor chemical gases for a given growth process. It can be combined with other techniques, such as in the growth of MgB₂ a recently discovered 40 K conventional electron-phonon superconductor, which uses a gaseous phase for the boron, but thermal evaporation of solid metal for the magnesium.

A schematic of this hybrid physical chemical vapor deposition (HPCVD) procedure is illustrated in the left panel of Fig. 1.7 and is the process used in making high quality MgB₂ films [Zeng, et al. (2002)]. It shows the sample substrate region in black, atop the red sample holder. The boron gaseous precursor flows continuously past the sample, and Mg vapor is generated around the sample by the heating of solid Mg. The quality of the films can

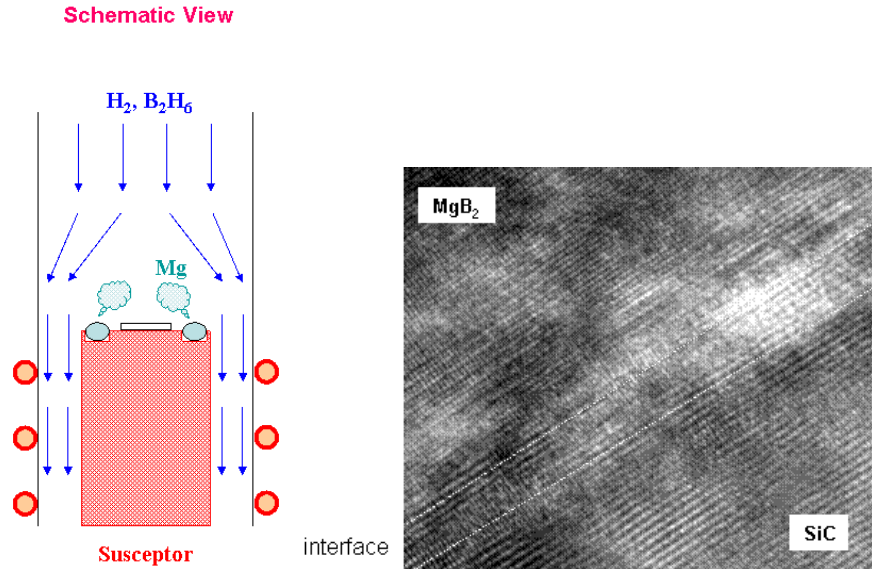


Fig. 1.7 Left panel: schematic diagram of the hybrid physical chemical vapor deposition process used to make ultra high quality MgB_2 films. Right panel: cross-sectional TEM image of the films showing a narrow interface region, where the sample quality is degraded (diagonal region about five atomic planes thick near center of figure). *Right panel reprinted with permission from [Xi (unpublished)]. Left panel reprinted with permission from [Progrebnyakov *et al.* (2004)]* (© 2004 the American Physical Society).

be seen in the cross-sectional TEM image in the right panel, which shows the substrate (SiC), the high quality atomically flat layers of MgB_2 , and a thin interface region (about five atomic planes thick) where substrate steps and dislocation defects are located and degrade the sample quality. These films are such high quality because the degraded region is so thin.

There are many ways to characterize the quality of the final device that has been grown. We have already shown a number of TEM images, which can determine where the atoms sit, and thereby provides information on the flatness of the interfaces, and of interdiffusion or chemical reactions at the interfaces. But a TEM image is a destructive process, because one needs to slice, polish, and thin the sample until it can be imaged. Furthermore, we are often interested in understanding properties of the transport in a device, and such information cannot be revealed by TEM measurements. Another technique that is quite useful is called ballistic electron emission microscopy or BEEM for short. This measurement is shown schematically in Fig. 1.8. A

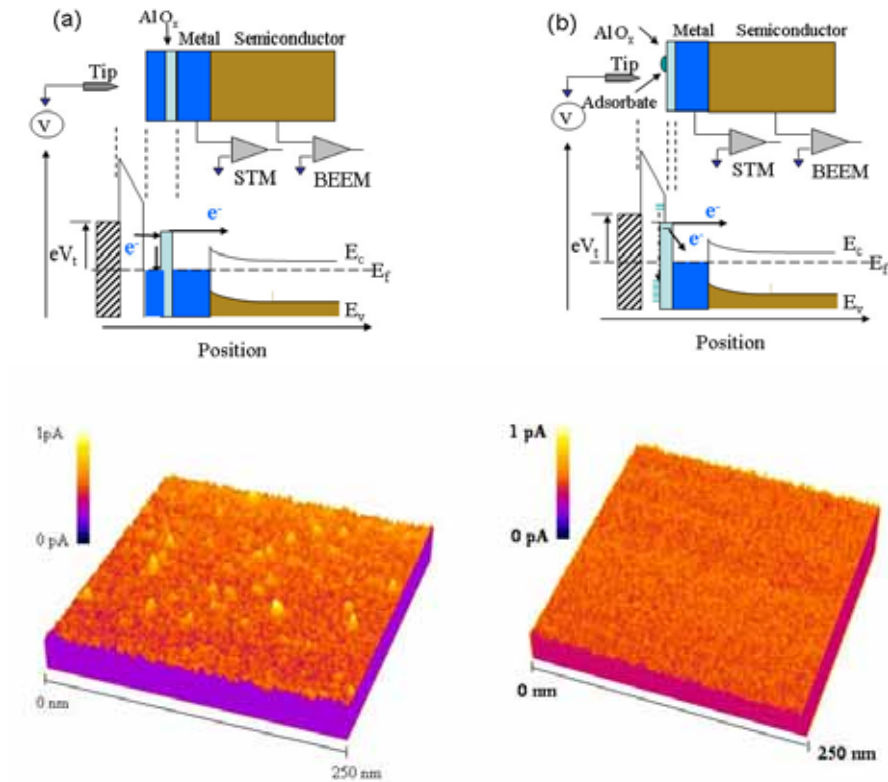


Fig. 1.8 Top panel: BEEM schematic for two different types of samples. The STM is always scanned over the surface with a bias voltage applied to it, and there always is a Schottky barrier formed by the electronic charge reconstruction at a metal-semiconductor interface to provide a barrier to electrons moving through the device. Energetic electrons will pass over the barrier and be collected. In this fashion, one can determine the local contributions to current flow through the device. Such a map is presented in the bottom two panels for a thin (left) and thick (right) disordered AlO_x barrier. The thin barrier has pinholes, while the thick barrier is pinhole free, and has nearly uniform current flow. *Reprinted with permission from [Buhrman (unpublished)].*

scanning tunneling microscope (STM) tip is scanned over the surface of the sample with a voltage difference applied so that it can eject electrons into the sample. Since the sample sits on top of a metal-semiconductor interface, the electron needs to have enough energy to get over the Schottky barrier that forms due to an electronic charge reconstruction at the interface, in order to be collected. By monitoring this collection current versus the position of the STM tip, one can directly measure the uniformity of the

sample for perpendicular transport. In other words, one can actually image the so-called pinholes, which are “hot spots” in the device that allow current to flow more easily and provide an inhomogeneous current flow through the device; usually one does not want to have pinholes, because the random nature for how they form can significantly effect the uniformity of device parameters across a chip. Two BEEM images of a disordered aluminum oxide barrier are shown in the bottom panels of Fig. 1.8 [Rippard, *et al.* (2002); Perrella, *et al.* (2002)]. The left panel has a very thin layer, and the right panel has a thicker layer. One can clearly see the pinholes on the left (bright yellow regions), which then become much more uniform on the right. In both cases, however, the barrier is still quite disordered, because the aluminum oxide is not stoichiometric. This can be inferred, in part, from the fact that the barrier height to tunneling, which can also be measured in the BEEM experiment, is far below half of the band gap of Al_2O_3 . What is interesting from a device standpoint is that the disordered aluminum oxide barrier creates a uniform tunnel barrier for transport, even if it is nonstoichiometric, as long as it is thick enough [Rippard, *et al.* (2002); Perrella, *et al.* (2002)]. This is one reason why it is so useful in so many different types of multilayered nanostructures.

There is a simple model that explains why the oxygen defects form in aluminum oxide [Mather, *et al.* (2005)], and we describe this model in Fig. 1.9. The common way to form an aluminum oxide layer is to first put down a layer of aluminum, and then to introduce oxygen gas for a certain period of time at a certain pressure to allow the aluminum to oxidize. In some devices, like Josephson junctions, there is no device degradation if some unoxidized aluminum remains, because it will be made superconducting by the proximity effect, while in other cases, like in magnetic tunnel junctions for spintronics, one wants all of the aluminum to oxidize, because metallic aluminum will degrade the tunnel magnetoresistance. The model for the oxidation process is that the oxygen first sits on the surface of the aluminum before it is driven into the sample. After some oxygen has moved in, the oxygen vacancies reach a steady state with the chemisorbed oxygen surface layer, and no more oxygen will flow through to oxidize the aluminum further. When the device is then processed to add additional layers, the oxygen surface layer will either be driven in (due to the processing conditions) or will react with the new layers being added on top, which can potentially degrade the top interface. Heating the sample prior to additional growth of multilayers can drive the chemisorbed oxygen into the aluminum and reduce the number of defects. Indeed, if the device is

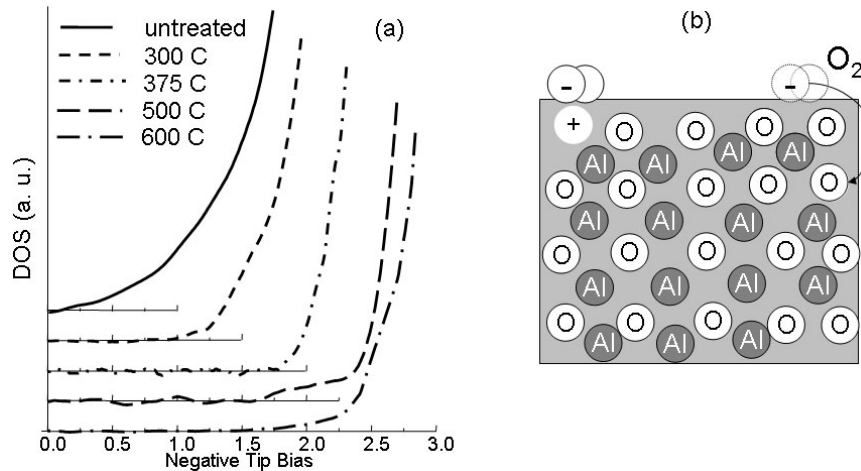


Fig. 1.9 Model for aluminum-oxide growth by exposing a thin film of aluminum to oxygen. On the right, one can see how a chemisorbed oxygen layer can form on the surface, by binding electronically to the defect sites; this chemisorbed layer does not allow further oxygen to flow into the barrier. By heating the sample, one can thermally activate the oxygen to move over the barrier and be driven into the aluminum layer. This is confirmed in the left panel, which shows how the tunnel current turns on at a higher and higher voltage as the sample is annealed at higher temperatures, and eventually a barrier height equal to half the Al_2O_3 band gap develops. *Left panel reprinted with permission from [Mather, et al. (2005)] (©2005 American Institute of Physics) and right panel reprinted with permission from [Buhrman (unpublished)].*

annealed at higher and higher temperatures, one sees the expected barrier height for Al_2O_3 begin to develop (see the left panel of Fig. 1.9).

In this section, we have described a number of different growth processes and characterization tools for multilayered nanostructures. The growth process is often quite complex, and significant care must be taken to achieve high quality results, but the state-of-the-art does allow quite good devices to be grown in research laboratories. Characterization tools used both during growth and after growth allow the device properties to be determined and understood, helping to find new ways to grow even better devices in the future. We will be concentrating on describing the theoretical and numerical formalisms for how to determine the transport through such devices throughout this book.

1.2 Strongly Correlated Materials

The first successful semiclassical attempt to describe the conduction of electrons in metals was given by Paul Drude in 1900 [Drude (1900a); Drude (1900b)]. This model assumes that electrons move independently through the crystal without feeling the effects of the other electrons but they do scatter off of defects, impurities, lattice vibrations, *etc.*, with a constant scattering time called the relaxation time. From this simple assumption, one can produce a constant electrical current from an applied electrical field (as described in virtually every solid state physics text). This theory was modified by Arnold Sommerfeld in 1927 to include the quantum-mechanical effects of the Fermi-Dirac distribution of electrons and the Pauli principle [Sommerfeld (1927)]. In spite of its incredible simplicity, the Drude-Sommerfeld model works remarkably well in describing the behavior of a wide variety of metals. The theoretical basis for understanding why such a simple model works so well was established by Lev Landau with the introduction of Fermi-liquid theory [Landau (1956)]. Fermi-liquid theory maps the elementary excitations of the interacting electronic system onto the excitations of a noninteracting system, and describes the residual weak interactions with a small set of phenomenological parameters. Nearly all metals can be described by Fermi-liquid theory (or “dirty” Fermi liquid theory, which corresponds to Fermi liquids with some additional static disorder that creates a finite relaxation time at the Fermi energy when $T = 0$). The basic result of Landau’s Fermi-liquid theory is that some fraction of the electrons, corresponding to the electrons with the lowest available energies, behave like noninteracting electrons with an infinite relaxation time at the Fermi energy when $T = 0$. Hence they can be described well by semiclassical approaches at finite temperature even though the electrons do feel an electron-electron repulsion from the other electrons in the material.

Strongly correlated electrons are, in general, different from these “garden-variety” electrons found in most metals. In strongly correlated electron materials, the electrons feel strong effects of the other electrons, and hence their motion is constrained by the positions of the neighboring electrons, which can lead to interesting phenomena, most notably a metal-insulator transition, as was first described by Nevill Mott [Mott (1949)].

The Mott metal-insulator transition is easiest to describe with an artificial material of atomic hydrogen placed on a crystal lattice with a continuously varying lattice parameter. If we assume that the electrons do not congregate between the hydrogen nuclei, and hence rule out the forma-

tion of molecular hydrogen, then the system can be described by electrons that hop on a lattice constructed by the periodic arrangement of the hydrogen nuclei. If the lattice parameter is very large, then each electron is tightly bound to a nucleus, and we have a collection of isolated hydrogen atoms, which will not conduct electricity because the electrons are localized, and cannot be unbound by applying a small electric field. This state is an insulator. If we now shrink the lattice spacing, bringing the atoms closer together, then the wavefunctions of the electrons will begin to overlap. When this occurs, the electrons can hop from one hydrogen atom site to a neighboring hydrogen atom site if the electrons have opposite spins. Once such a process is allowed, the electrons become delocalized, and then they can screen out the bare Coulomb attraction with the nuclei, which will tend to make them even more delocalized, and eventually they will become metallic, easily conducting electricity when a small electric field is applied. The change in character from a metal to an insulator as the lattice spacing increases is the classic example of the Mott metal-insulator transition.

Strongly correlated electrons are a little bit different from the hydrogen example above, because it is the repulsion of the electrons with each other that determines their behavior, rather than the attraction with the ion cores (which in most crystals determines the band structure). Hubbard devised the simplest model for this behavior [Hubbard (1963)]. In his model, which is described in detail in Chapter 2, we have electrons that move in a single band on a lattice. They can hop to their nearest neighbors with a hopping integral t . When two electrons sit on the same lattice site, there is a screened Coulomb repulsion U . All other long-range Coulomb interactions are neglected. If we have on average one electron per site, then if $U \ll t$, the electrons are delocalized in a band and their motion is only slightly modified by the electron-electron interaction. If, on the other hand, we have $U \gg t$, then the Coulomb repulsion is so strong we cannot have two electrons (of opposite spin) occupy the same lattice site. Hence we have exactly one electron per site, and this configuration is frozen with respect to charge excitations, so the system is an insulator. This implies that there is a Mott-Hubbard metal-insulator transition as a function of U . The transition occurs at $U \rightarrow 0^+$ in one dimension [Lieb and Wu (1968)], but at finite U values for higher dimensions.

Predicting when a real material will display Mott-Hubbard insulating behavior is quite difficult. One simple rule is that if a density functional theory calculation predicts the system is metallic, but experiment shows it to be insulating, then it is a strongly correlated insulator. But such a




Fig. 1.10 Pressure-temperature phase diagram of the κ -Cl material. Transport measurements on this system identified four regions: (1) a Mott insulator; (2) a semiconductor; (3) a bad or anomalous metal; and (4) a Fermi-liquid metal. These four regions, along with the antiferromagnetic phase are shown in the phase diagram (the superconducting phase, which is also present, has not been depicted). The general character of this phase diagram, in particular, the first-order phase transition between the metal and insulator at intermediate temperatures, can be explained by numerical solutions of the Hubbard model using dynamical mean-field theory. *Figure reprinted with permission from [Limelette et al. (2003)]* (© 2003 the American Physical Society).

definition is neither rigorous, nor does it allow for much predictive power in finding new Mott insulators. Gebhard goes to great lengths to carefully describe conditions under which one has a Mott insulator [Gebhard (1997)], and the interested reader is referred there. More recently, a combination of density functional theory plus dynamical mean field theory shows promise in being able to provide a numerical framework for predicting Mott insulators and determining their properties, but the current techniques require huge investments in computer time, so it is not yet a practical tool for numerically exploring new materials (see Sec. 1.7).

We end this section by giving a recent explicit example of experimental work and calculations that illustrate the Mott insulating behavior of a strongly correlated material. This new material is of high interest, because the transition to different regions of the phase diagram can be reached by relatively small changes in either pressure or temperature. Experiments

on the organic material κ -(BEDT-TTF)₂Cu[N(CN)₂]Cl (called κ -Cl) show that it can be tuned through the Mott transition by varying the pressure over a range of about 1 kbar and temperatures up to 80 K [Limelette, *et al.* (2003)]. Results for the phase diagram are shown in Fig. 1.10. As the pressure increases, the ratio U/t decreases, so we see a Fermi-liquid metal on the lower right and a Mott insulator (plus an antiferromagnetic ordered phase) on the lower left. When the system is heated up, the insulating phase becomes more semiconducting, and the Fermi-liquid behavior disappears above the renormalized Fermi temperature; as the system goes into this incoherent phase it is metallic, but with anomalous properties, and typically poor conductivity. If the temperature is raised even further, the first-order transition between the metallic and insulating phases disappears at a classical critical point, above which, the system can undergo a smooth crossover from a metal to an insulator as a function of pressure. The general behavior of this metal-insulator transition is similar to that of the liquid-gas phase diagram of many liquids.

1.3 The Proximity Effect

In quantum mechanics, a “box” determined by a finite potential barrier is a leaky box, because the wavefunction of the electron always extends out of the box boundaries with an exponentially decaying wavefunction. This is shown schematically for a one-dimensional box in Fig. 1.11, where the wavefunctions of the two lowest bound states are plotted (centered on their respective eigenenergies). Note how there is always a finite probability to find the electron lying outside the box due to the uncertainty principle.

In many-body physics, a similar phenomenon occurs whenever two different materials are joined together at an interface; the wavefunctions of the right material leak into the left and *vice versa*. This mild sounding observation leads to some amazing quantum-mechanical effects; indeed the rest of this book focuses on investigating such effects. This “leakage of electrons” across a barrier is called tunneling. It is in many respects a mature subject. Esaki [Esaki (1958)] described a tunnel diode made out of semiconductors in the late 1950s, which was shortly followed by the superconducting version studied by Giaever [Giaever and Megerle (1960)]. Josephson [Josephson (1962)] showed that one gets surprising effects in a superconducting tunnel junction when the barrier is made thin enough. All three shared the Nobel prize in 1973 for their work on tunneling.




Fig. 1.11 Lowest two wavefunctions for a particle in a one-dimensional box depicted by the thick solid lines (these are the only bound states for a box of this depth). The dashed lines are the values of the respective energy levels. Note how the wavefunction for each case leaks out of the “boundary” of the box.

The best known proximity effect occurs in a Josephson junction [Josephson (1962)] [Anderson and Rowell (1963)], which is a sandwich structure composed of a superconductor-barrier-superconductor. A superconductor is a metal that has a net electron-electron attraction mediated by a phonon (in conventional low-temperature superconductors), which causes electrons with opposite momentum and spin to pair together (due to the so-called superconducting correlations). The physical picture is similar to two marbles on a rubber sheet—each feels the depression of the other marble, and they roll toward each other. In real superconductors, the electrons also repel each other because they have the same electronic charge; the superconductivity occurs because there is a time delay for the interaction with the phonons, which allows them to pair electrons together that are not located at the same position at the same time. The pairing leads to an energy gap, so the superconductor has no low-energy excitations below the energy of the superconducting energy gap (typically on the order of 1 meV). In the Josephson junction, the pairing correlations of the superconductor on the left leak into the nonsuperconducting barrier region in the middle (be it a metal or an insulator), and join up with the superconducting correlations in the superconductor on the right. This weak

link between the two superconductors can carry current across it if the macroscopic quantum-mechanical phase changes across the barrier region. This leads to the Josephson supercurrent—a finite current carried by superconducting pairs with zero voltage across the barrier. There is also a corresponding inverse proximity effect, where the pairing correlations in the superconductor are weakened by the closeness to the interface with the barrier.

The physical picture for the proximity-effect coupling of Josephson junctions is different for insulating and metallic barriers. In metallic barriers, the barrier has low-energy states, but the superconductor has none. As a superconducting pair approaches the interface with the barrier, it meets a hole in the metal, which annihilates one of the electrons, while the other electron moves through the barrier to the next interface. There, the electron is retro-reflected as a hole (the hole has the opposite momentum and energy of the electron), leaving behind a superconducting pair to travel through the superconducting lead on the right. This process is called Andreev reflection [Andreev (1964)] (see Fig. 1.12); it takes place over a time scale on the order of \hbar/Δ independent of the barrier thickness L . In insulating barriers, the barrier has no low-energy states, so the electron pairs must tunnel through the barrier, which occurs due to the quantum-mechanical “leakage” through the barrier. Obviously the supercurrent decreases faster with the thickness of the barrier when it is an insulator than when it is a metal (although both decay exponentially with the thickness).

In a normal-metal–barrier–normal-metal nanostructure, there is also a proximity effect, and it is similar to the problem of a quantum-mechanical particle in a box (Fig. 1.11) when the barrier is an insulator, because the metallic wavefunctions see a potential barrier at the interface, since there are no low-energy states in the insulator. Hence the wavefunctions decay exponentially until they reach the center of the barrier, and then they grow until they reach the metallic interface on the other side. Since the wavefunction connects the two metallic leads, the electrons can directly tunnel from the right to the left (or *vice versa*). In the metallic case, the proximity effect is more subtle, dominated by generating oscillations (with the Fermi wavelength) in the metallic leads due to the mismatch of the wavefunctions between the two metals. Similar effects can also occur in the barrier.

The study of multilayered nanostructures relies heavily on understanding proximity effects between dissimilar materials brought close together in a heterostructure. This forms a significant part of the final five chapters.




Fig. 1.12 Schematic plot of the Andreev reflection process. The low-energy electrons in the metal are confined to the barrier due to the energy gap Δ in the superconductors (the energy of the superconducting ground state is chosen as the zero in this diagram), so they form an electron-hole bound state, which allows a superconducting pair to travel from the left to the right through the Josephson junction. A similar process allows for current to travel from right to left. The symbol p_s denotes the momentum of the superconducting pair. *Figure adapted with permission from [Shafraniuk (unpublished)].*

1.4 Electronic Charge Reconstruction at an Interface

In surface physics, the process of a surface reconstruction, where the atoms on the surface rearrange themselves in response to the dangling bonds resulting from the interface with the vacuum, is well-known. The surface reconstruction of silicon was one of the first systems to be imaged with the scanning tunneling microscope [Binnig, *et al.* (1983)]. Much of the study of surfaces and how they interact with material deposited on the surfaces relies on understanding how the surface reconstructs itself.

In multilayered nanostructures, there are no open surfaces, and there is limited freedom for ions to rearrange their spatial locations in response to the interface with a different material (small relaxations of atoms near the interface certainly occur). But there is no reason why the chemical potential of the leads of the device needs to match the chemical potential of the barrier. This puts the barrier in an unstable situation, where some of the electrons are forced to either leave or enter the barrier from the leads (depending on the relation of the chemical potentials). Because the Coulomb interaction is long-ranged, the charge redistribution will be confined to the interface regions, with a healing length on the order of the Thomas-Fermi screening length [Thomas (1927);

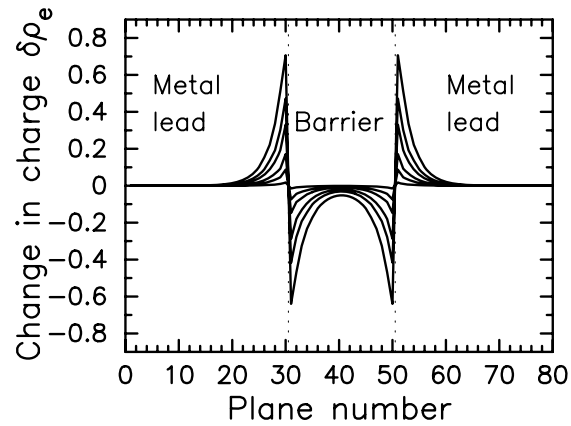


Fig. 1.13 Schematic of the electronic charge reconstruction at the interfaces of a multilayered nanostructure with a barrier that is 20 planes thick. The horizontal axis is the plane number (the barrier lies on planes numbered from 31 to 50—the interface is indicated with the dashed lines) and the vertical axis is the self-consistently calculated change in the charge density on each plane. The chemical potential of the barrier differs from that of the bulk metallic leads for each of the different curves. The screening length is chosen to be approximately 2.2 lattice spacings in both the metallic lead and the barrier. *Figure adapted with permission from [Nikolić, Freericks and Miller (2002a)] (original figure © 2002 the American Physical Society).*

Fermi (1928)] (usually less than an Angstrom in metals). The result is a screened-dipole layer at the interface, which creates an electric potential that causes scattering to electrons moving through the device and is plotted in Fig. 1.13 [Freericks, Nikolić and Miller (2002)]. One can see how charge spills from the barrier into the lead as the mismatch of the chemical potentials is increased. This effect is well known in the semiconductor community when a metal is placed in contact with a semiconductor creating a Schottky barrier [Schottky (1940)]. It is used to create a number of the different semiconductor-based devices.

The electric fields created by these screened dipole layers can be quite large. They do not cause current to flow, however, because they are exactly compensated by an opposite force due to the diffusion current arising from the change in the electron concentration. This is because the system has reached a static, equilibrium, rearrangement of the electronic charge.

One of the most interesting applications of interface charge reconstruction is the case of a metal-oxide-semiconductor field-effect transistor (MOS-

FET). In this device, one brings together a semiconductor and an insulator forming a sharp interface. The electronic charge reconstruction creates a thin layer of electrons that are trapped to lie in close proximity to the interface. If engineered properly, the dopant ions, which created the electron carriers in the first place, lie in the semiconductor, while the electrons lie in the insulator. Then the electrons are far away from scattering sites, and they can become incredibly mobile. It is within these systems that the quantum Hall effect and the fractional quantum Hall effect were both discovered. The creation of this “nearly free” two-dimensional electronic gas follows from the physics behind charge reconstruction at an interface.

Interface charge reconstruction will naturally occur in strongly correlated nanostructures as well, leading to even more interesting behavior when one of the materials is a strongly correlated insulator, since the charge depletion (or enhancement) can “dope” the insulator into a strongly correlated metal phase (or *vice versa* if the material is already a strongly correlated metal). These effects have been imaged in grain boundaries of high temperature superconductors, where the grain boundaries are known to be electrically active [Mannhart and Hilgenkamp]. A grain boundary occurs in the growth of a material where islands of different grains meet, and the temperature is too low for the system to anneal the crystallite boundaries out of the system. A TEM image of just such a grain boundary can be seen in the left panel of Fig. 1.14 [Browning, *et al.* (1993)]. This grain boundary has a large angle orientational mismatch, as is easily seen. Unfortunately, these grain boundaries have a significant deleterious effect on superconducting wires, as they create Josephson junction weak links between the grains, and the critical current of the weak link is much smaller than the maximal critical current of a bulk single crystal. This has proved to be the single largest hurdle to get over in making high temperature superconducting wires (of course, the presence of the grain boundaries can be employed to manufacture Josephson junctions, if desired).

The right panel of Fig. 1.14 [Browning, *et al.* (1993)] depicts the valence of the Copper atom as a function of the distance away from the grain boundary. Clearly the grain boundary is electrically active, and has a charge reconstruction. What is amazing is how far away from the grain boundary this charge rearrangement extends, which is likely due to the fact that the strongly correlated metal does not screen charge as efficiently as a more conventional metal. The charge distortion is reduced as the misorientation angle of the grain boundary is reduced; this is the underlying phenomenon that governs the reduction of critical current at a grain boundary.




Fig. 1.14 Left panel: high-angle grain boundary in a high temperature superconductor. Right panel: charge profile around the grain boundary (Copper valence) as measured with electron energy loss spectroscopy (see Fig. 1.16 below). The probe position (horizontal axis) is relative to the center of the grain boundary. The vertical axis is proportional to the valence on the copper atom, which changes from a maximum of 2.6 at the top to a minimum of 1.0 at the bottom, as the probe is moved across the grain boundary. Reprinted with permission from [Browning *et al.* (1993)].

Since diffusion of chemical species is easier along grain boundaries than within the grains themselves, it was discovered that the critical current across a grain boundary could be enhanced by diffusing Calcium ions to the grain boundary location [Hammerl, *et al.* (2000)]. The Ca ions must be modifying the local charge reconstruction at the grain boundary to do this. An interesting way to improve the critical current density of a high-temperature superconducting tape is to grow multilayers of pure Yttrium-Barium-Copper-Oxide, and of Calcium-doped YBCO. Between the grain boundaries the current will be carried predominately in the pure YBCO, but at the grain boundaries, because the presence of Calcium reduces the charge reconstruction, the critical current density is not reduced as much as in the pure YBCO. A schematic of this multilayered device is shown in left panel of Fig. 1.15, and the improvement in the critical current is shown in the right panel. At this point, it is not clear whether this process can be used to make high temperature superconducting wires into a viable technology.

Another example is the artificially engineered band-insulator/strongly correlated insulator heterostructure made from SrTiO_3 (a band insula-




Fig. 1.15 Left panel: schematic of the growth of pure YBCO (yellow) and Ca doped YBCO (red) for increasing the critical current at the grain boundary. Note how the Calcium dopes preferentially into the grain boundary region (the grain boundary is the black line), presumably changing the electronic charge reconstruction. Right panel: enhancement of the critical current density due to Ca doping (increase from the red to the orange curve after doping). Reprinted with permission from [Mannhart (2005)].

tor that is nearly ferroelectric) and LaTiO_3 (a strongly correlated insulator) [Ohtomo, *et al.* (2002)]. The heterostructures of these materials are made using PLD, and varying the Sr or La content within the titanate background. The heterostructures are grown with nearly atomically flat precision and excellent control over the thicknesses of the different layers. A detailed analysis of the structure shows little interdiffusion of the species across the interface. What is surprising, is that the system has metallic conducting channels in the transverse direction (along the planes rather than perpendicular to the planes), which vary with the thickness of and the spacing of the LaTiO_3 layers within the SrTiO_3 matrix. Sophisticated experimental equipment is needed to image the charge redistribution in multilayered nanostructures, because one needs to have both sensitivity to the local charge, and an ability to achieve atomic resolution. One way that this is accomplished is by combining electron microscopy observations with electron energy loss spectra (EELS) as shown in Fig. 1.16. This is done with a dedicated scanning transmission electron microscope (STEM) that is equipped with an annular detector and an electron spectrometer. In the STEM, the optics are devoted to focusing the electron beam to a very fine probe (0.13 nm diameter), which is raster scanned over the sample. The transmitted electrons scattered at high angles are collected into an annular dark field detector which is used for the imaging. Since these electrons are primarily Rutherford scattered by the ion cores, the image intensity will be roughly proportional to the square of the atomic number. This is why this technique is called Z-contrast microscopy (for a review of the instru-

ment see [Pennycook (2002)]). It is capable of producing incoherent images with atomic resolution and atomic specificity. Electrons traveling parallel to the optical axis (*i.e.* through the hole in the annulus) are collected into the EELS, so simultaneous EELS measurements can be obtained. These spectra can be employed to determine the local electronic charge, or the energies for the thresholds of different excitations, or the local chemical environment of a particular ion.

This imaging technique was used to measure the charge profile near the grain boundary, shown in Fig. 1.14 [Browning, *et al.* (1993)], and was used in the SrTiO₃/LaTiO₃ heterostructures [Ohtomo, *et al.* (2002)]. This imaging technique has also been applied to YBa₂Cu₃O_{7- δ} /La_{0.67}Ca_{0.33}MnO₃ heterostructures [Varela, *et al.* (2003); Varela, *et al.* (2005)]. They find that the interfaces are nearly atomically flat, with essentially no interdiffusion of chemical species across the interface (determined by examining the EELS results). They also can use the STEM-EELS apparatus to map out the local charge density, which is plotted in Fig. 1.16. One can see how the charge screening length is much shorter in the LCMO material than in the YBCO, but the heterostructure is not thick enough for the LCMO material to heal its charge to the bulk value.

The phenomena described above has been termed electronic charge reconstruction [Okamoto and Millis (2004a); Okamoto and Millis (2004b)], due to its similarity with the well-known surface reconstruction. Okamoto and Millis analyzed the SrTiO₃/LaTiO₃ system [Ohtomo, *et al.* (2002)] using a hybrid density functional theory/many-body theory approach. The low-energy bands are modeled with a tight-binding scheme, and Coulomb interactions are introduced to describe the electron correlations. The many-body theory was analyzed in a static mean-field theory approach [Okamoto and Millis (2004a)] and in another approximate many-body physics method that can produce the MIT [Okamoto and Millis (2004b)]; both produced much insight into the physics behind this behavior. In particular, since the different systems are at different chemical potentials in the bulk, there is a localized charge transfer at the interfaces, which artificially dopes each of the insulators. This leads to metallic regions near the interfaces that can conduct electricity in the transverse (planar) directions. The results of their calculations are summarized in Fig. 1.17.

Electronic charge reconstruction is a phenomenon that naturally occurs at the interface of any two materials unless they happen to have exactly the same chemical potential (which is unlikely to occur in any real system at all

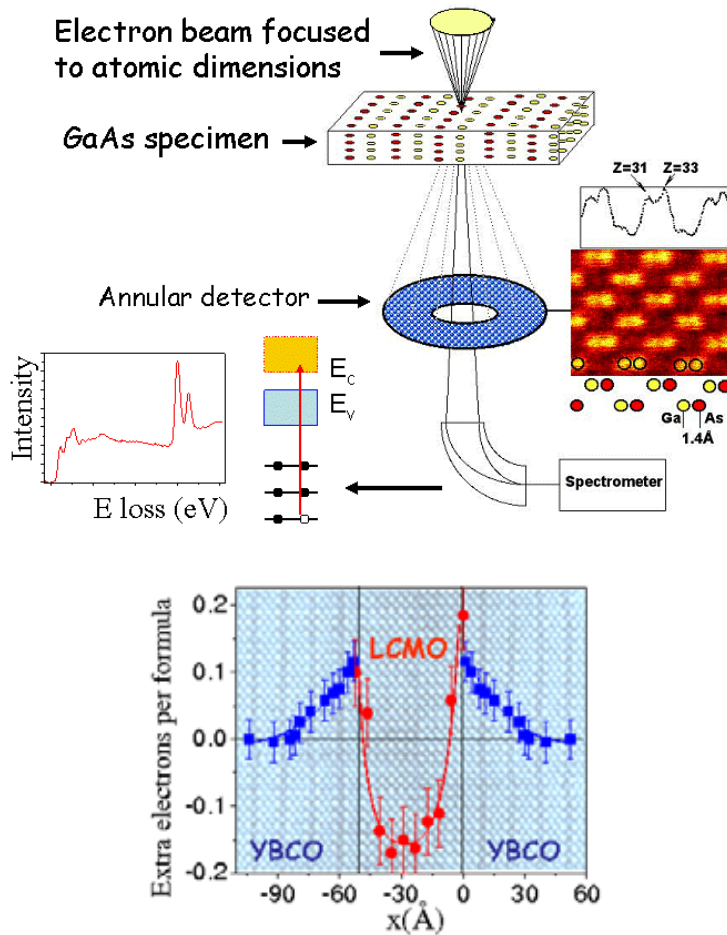


Fig. 1.16 Top panel: Experimental setup for a STEM-EELS measurement. This example shows a measurement on GaAs. A simultaneous measurement of the positions and types of the atoms and of the charge profile can be achieved. Bottom panel: Charge profile overlaid over the atomic positions of a $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}/\text{La}_{0.67}\text{Ca}_{0.33}\text{MnO}_3$ heterostructure. Top panel adapted with permission from [Pennycook (2002)]. Bottom panel reprinted with permission from [Varela (2005)].

temperatures). We describe how to perform self-consistent calculations of electronic charge reconstruction in Chapter 3 and give additional numerical examples in Chapter 6.




Fig. 1.17 Calculated local charge density near the interface of a SrTiO₃/LaTiO₃ heterostructure with six planes of LaTiO₃. The top curve shows the charge density as a function of distance in the inhomogeneous (longitudinal) direction. The bottom curve shows the charge density in the Fermi-liquid-like coherence peak near the electrochemical potential. It is clear from the curves, that the interface regions are conducting, with a thickness on the order of three atomic planes. *Reprinted with permission from* [Okamoto and Millis (2004b)] (© 2004 the American Physical Society).

1.5 Roadmap to Real-Materials Calculations

In this book, we concentrate on calculations for model Hamiltonians which usually include one itinerant electron band only. Model systems have been used for decades in many-body physics because they capture the important quantum-mechanical aspects of the problem, but are simpler than materials-specific calculations. Much can be learned about the many-body problem, and about strongly correlated nanostructures by examining these model systems.

But eventually we want to be able to handle real-materials problems in a “first principles” fashion. First principles calculations usually start from the density functional theory with the local density approximation or a generalized gradient expansion. Density functional theory is exact for the ground state energies of real materials if the exact exchange-correlation functional is known for the material [Hohenberg and Kohn (1964)]. The exchange-correlation functional is complicated, and not known in the general case.

The approximation, called the local density approximation, assumes that the functional is the same as the functional of a uniform electron gas with the same density at a given position in space. This approximation should be accurate for metals where the electron density does not change sharply through the material. It has been generalized to include gradient terms for the change of the density as well.

The common way that density functional theory is solved is to map the interacting problem onto a noninteracting problem with the same electron density (but in a complex potential) [Kohn and Sham (1965)]. Solving this problem numerically yields a band structure, which is believed to be similar to the true band structure of the material (assuming such a concept exists), but density functional theory provides no proof of this result. If we want to go beyond the mean-field-like treatment of the local density approximation, we first parameterize the density functional theory bands by a tight-binding model (where choosing the appropriate basis can be critical), and then add electron-electron interaction terms. These interactions can be solved with the techniques of dynamical mean-field theory under the assumption that the electronic self-energy is local (that is, independent of momentum). There has been much progress in solving for properties of strongly correlated materials in the bulk with this procedure, but it is a computationally intense project.

Plutonium is one of the most interesting materials from a solid-state physics context. It possesses numerous phases as functions of temperature and pressure, and it is generally believed that a number of these phase transitions arise from strong electron correlations. The $\alpha - \delta$ phase transition of Pu is interesting because it is accompanied by a 25% volume change, which is believed to be governed primarily by a change in character of the electrons from a band-like metal to a localized insulator. Since this transition can affect the stability of Pu when it is stored for long periods of time (Pu will be self-heated due to the nuclear radioactivity), it is of significant importance to understand its properties. This system was chosen as one of the first systems to apply the DFT+DMFT approach to [Savrasov, Kotliar and Abrahams (2001)]. The first calculations focused on the electronic structure, as modified by the strong electron interactions, and were then followed up by work on the phonons [Dai, *et al.* (2003)]. The phonon dispersions were later verified by inelastic X-ray scattering [Wong, *et al.* (2003); Wong, *et al.* (2005)] and summarized in a short review [Kotliar and Vollhardt (2004)] (see Fig. 1.18); this agreement between experiment and theory shows the power of these methods.

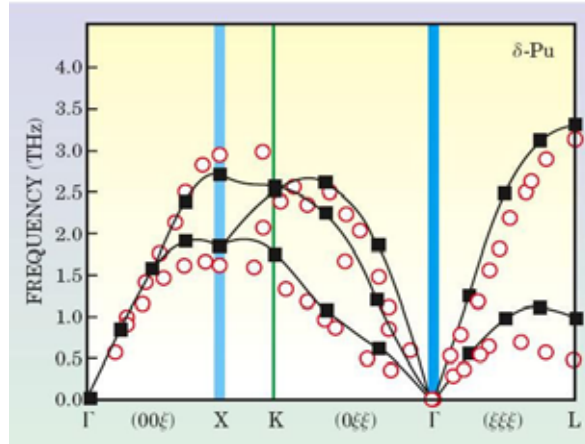


Fig. 1.18 Theoretically calculated phonon dispersion of δ -phase Pu (squares connected by full lines) which were later confirmed by inelastic X-ray scattering measurements (open red circles). *Figure reprinted with permission from [Kotliar and Vollhardt (2004)] (©2004 American Institute of Physics) (adapted with permission from [Dai *et al.* (2003)], [Wong *et al.* (2003)] and [Wong *et al.* (2005)]).*

The work of Vollhardt and collaborators has concentrated on examining transition-metal oxides. Initially they focused on V_2O_3 [Held, *et al.* (2001); Mo, *et al.* (2003); Keller, *et al.* (2004)], which is close to the Mott metal-insulator transition; more recent work has examined strongly correlated metals like $CaVO_3$ and $SrVO_3$ [Sekiyama, *et al.* (2004); Nekrasov, *et al.* (2005)]. In the left panel of Fig. 1.19, we show the spectral function (below the Fermi energy) obtained from a high energy (bulk) photoemission study. Photoemission is an experiment where high energy light is shone onto a clean surface of the material and it expels an electron (Einstein won his Nobel prize for this photoelectric effect). By varying the angle of incidence of the photon, one can map out the excitation spectra (multiplied by a Fermi factor) as a function of momentum. Here we show the local spectra, summed over all momenta. The agreement between experiment and theory is excellent, and the theory shows small differences between the two systems. The right panel is a comparison with X-ray absorption spectroscopy, since inverse photoemission data is not yet available (inverse photoemission corresponds to electrons shone onto a surface and light emitted; X-ray absorption spectroscopy measures the ease with which X-rays can be absorbed by forcing an electron to have a transition from a K-edge oxygen core state




Fig. 1.19 Comparison of the parameter free LDA+DMFT(QMC) spectra of SrVO₃ (solid red line) and CaVO₃ (dashed blue line) with experimental data (symbols) below and above the Fermi energy. The left panel is high resolution photoemission spectroscopy for SrVO₃ (red circles) and CaVO₃ (blue squares) [Sekiyama *et al.* (2004)]. The right panel is 1s X-ray absorption spectroscopy for SrVO₃ (red diamonds) and CaVO₃ (blue triangles) [Inoue *et al.* (1994)]. The horizontal line is the experimental subtraction of the uniform background intensity. Reprinted with permission from [Nekrasov *et al.* (2005)] (© 2005 the American Physical Society).

to the Fermi energy of the correlated bands—hence they measure the DOS for the unoccupied states modified by atomic matrix elements, which are not expected to change the signal in this case).

The same strategy can be used for multilayered nanostructures. First a density functional theory is used to determine the band structure and from that an effective (inhomogeneous) tight-binding model is created. The interactions are introduced to represent the strong electron correlation effects. These are treated with inhomogeneous dynamical mean-field theory in order to fully solve the problem. So far, no one has attempted such a calculation. It requires a generalization of the techniques developed in this book along the lines of the progress made with first-principles calculations of bulk strongly correlated materials.