Chapter 2

Materials and Techniques

This thesis documents the use of a scanning tunneling microscope (STM) to study the high temperature superconductor $Bi_2Sr_2CaCu_2O_{8+\delta}$ (BSCCO). This chapter introduces both the material BSCCO, and the scanning tunneling microscope which can be used to measure its density of states. I also explain the relationship between STM and several other experimental techniques, such as angle-resolved photoemission spectroscopy, neutron scattering, and nuclear magnetic resonance, which can be used to measure or infer complementary information about the density of states.

2.1 $Bi_2Sr_2CaCu_2O_{8+\delta}$

We choose to study BSCCO with our STM, because BSCCO has weak bonds between the two BiO layers, so it can be easily cleaved to achieve an atomically flat surface. The structure of BSCCO is shown in figure 2.1(a).

Because an STM rasters a sharp tip within a few Ångstroms of the sample surface, it is essential that the surface be atomically flat. One loose atom on the surface could hop to the tip and cause the tip to become energetically unstable. The atom will flop around in several closely spaced energy states so that no further useful information can be obtained from the sample surface because it is swamped by spurious effects from the tip. The following are two requirements for STM study of a surface: (1) The surface must cleave exactly between layers, leaving no residual chunks of a mostly missing layer. (2) The surface must be free from other contaminants, for example helium or water molecules which may land on it.

To achieve a satisfactory surface, we cleave the BSCCO sample while it is surrounded

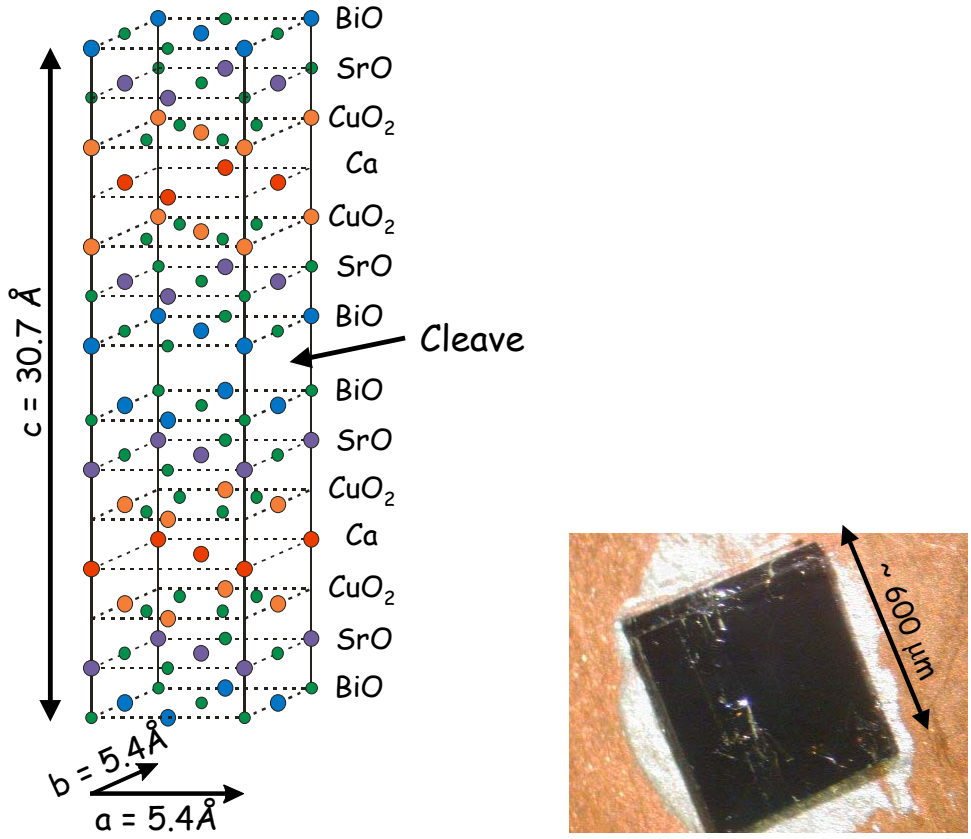

(a) (b)

Figure 2.1: (a) Chemical structure of $Bi_2Sr_2CaCu_2O_{8+\delta}$. The sample cleaves easily between the two BiO layers, and we image the exposed BiO layer. Note that two layers beneath every Bi atom (blue) lies a Cu atom (orange). (b) Photograph of a cleaved BSCCO sample, glued with conducting epoxy to a copper sample holder. The samples we use are typically ∼1 mm square and ∼10s of μm thick.

by a cryogenic UHV environment. Cleaving is mechanically simple. First we glue a small sample to the copper sample holder, then we glue a small rod to the other side of the sample. When the whole setup has been inserted into the correct cryogenic UHV environment, we knock the rod off, and it splits the BSCCO sample in two, exposing an uncontaminated surface. A photo of a cleaved surface (after the experiment, now sitting in air) is shown in figure 2.1(b). The rod is typically knocked off by swinging another rod perpendicular to it just before the sample is inserted into the STM.

In practice, there are many details, such as use of the correct glue to avoid cracking when exposed to cold temperatures, or shattering the BSCCO due to differential thermal contraction. The temperature may also vary: BSCCO may be cleaved in a cryogenic UHV

environment without itself actually being at a cryogenic temperature, if it is suspended from a hot rod to room temperature. Another variable is the speed at which the cleaving occurs. No controlled tests have been carried out, but the speed of cleaving must influence the local temperature at the cleaved surface, and could in the future be tuned to increase the success rate of cleaves.

One concern is whether the violent act of cleaving locally heats the surface and/or causes some kind of surface redistribution. It is apparent from topographic measurements that the bismuth atoms remain ordered as expected upon cleaving, but is possible that this violent action has some effect on the redistribution of the oxygen dopant atoms, which are so far invisible to STM.

The related technique of angle-resolved photoemission spectroscopy (ARPES) is not as surface sensitive as STM, but it is still somewhat surface sensitive. A surface reconstruction resulting in different electronic surface states would prevent ARPES from seeing the true bulk density of states. However, a few unwanted junk molecules on the surface, or a flaky or terraced cleave, would probably not preclude a good ARPES measurement, as they would preclude any kind of STM measurement.

Therefore, at the present time, BSCCO is the preferred material for study by STM and ARPES, although ARPES has successfully studied many other HTSC materials, and even STM is starting to branch out as more reproducible cleaving recipes are slowly figured out for other materials.

The BSCCO crystals studied in this thesis were grown by the floating zone method with superconducting transition temperatures ranging between underdoped (T_c = 65 K) and slightly overdoped (T_c = 85 K). The samples are cleaved at the BiO plane in cryogenic ultra high vacuum and immediately inserted into the STM head.

If all the interesting interactions of the superconductivity are thought to occur in the CuO plane, one might ask why are we satisfied with studying the BiO plane? Two layers beneath the BiO plane is the CuO plane. So in order to access the CuO plane, we need to tunnel through the BiO and SrO planes. Luckily, it is thought that both of these planes are insulating (or at least semi-conducting with a large gap)[50] so they can be treated as an extension of the vacuum tunneling barrier. In fact, their presence on top of the CuO layer of interest may be a blessing in disguise, because they protect the oxygen dopant atoms on *both* sides of the CuO plane, and thus preserve the charge carrier concentration in the CuO plane. If the top BiO layer were removed in the act of cleaving, the CuO plane would likely lose half of its charge carriers.

2.2 Scanning Tunneling Microscopy

The scanning tunneling microscope was invented in 1982 by Binnig and Rohrer,[51] for which they shared the 1986 Nobel Prize in Physics. The instrument consists of a sharp conducting tip which is scanned with respect to a flat conducting sample. When a voltage V is applied between tip and sample, a current will flow, and this current can be measured as a function of (x, y) location and as a function of V. This is illustrated schematically in figure 2.2.

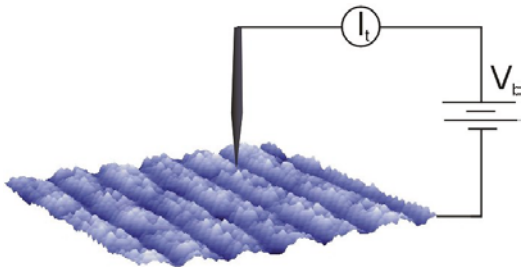

Figure 2.2: Schematic of STM tip and sample. A bias voltage V is applied to the sample, the tunneling current I is measured as V and/or the (x, y) position of the tip are varied.

2.2.1 Calculation of Tunneling Current

The current which flows between the tip and the sample can be calculated by time-dependent perturbation theory. If the sample is biased by a negative voltage $-V$ with respect to the tip, this effectively raises the Fermi level of the sample electrons with respect to the tip electrons. Electrons will tend to flow out of the filled states of the sample, into the empty states of the tip. This situation is illustrated in figure 2.3. The elastic tunneling current from the sample to the tip for states of energy ε (with respect to the Fermi level of the sample) is:

$$I_{\text{sample}\to\text{tip}} = -2e \cdot \frac{2\pi}{\hbar}|M|^2 \underbrace{(\rho_s(\varepsilon) \cdot f(\varepsilon))}_{\substack{\#\text{ filled sample states} \\ \text{for tunneling from}}} \underbrace{(\rho_t(\varepsilon + eV) \cdot [1 - f(\varepsilon + eV)])}_{\substack{\#\text{ empty tip states} \\ \text{for tunneling to}}} \quad (2.1)$$

where the factor of 2 out front is for spin, $-e$ is the electron charge (we are tunneling single electrons, not quasiparticles or Cooper pairs), $2\pi/\hbar$ comes from time-dependent perturba-

tion theory, $|M|^2$ is the matrix element, $\rho_{s(t)}(\varepsilon)$ is the density of states of the sample (tip), and $f(\varepsilon)$ is the Fermi distribution:

$$f(\varepsilon) = \frac{1}{1 + e^{\varepsilon/k_B T}} \tag{2.2}$$

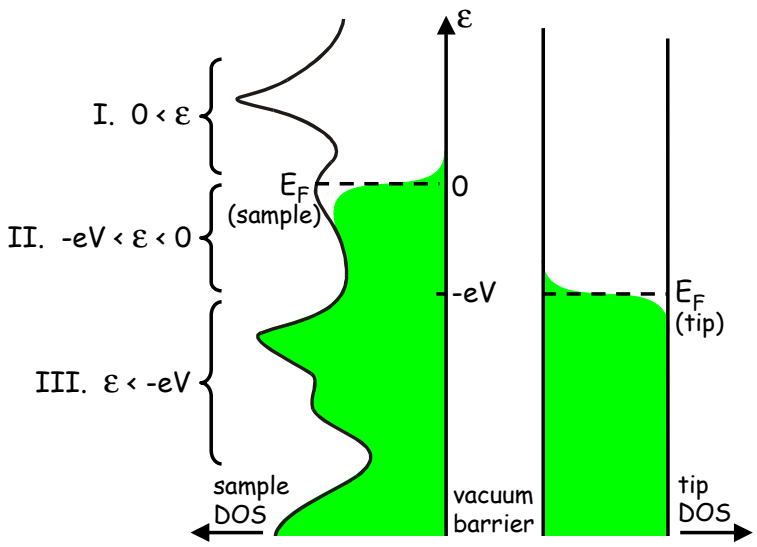

Figure 2.3: Schematic of tip-sample tunneling. Energy is along the vertical axis, and density of states of the sample and tip are shown along the horizontal axes. Filled states are shown in green. In this case, a negative bias voltage $-V$ has been applied to the sample, which effectively raises its Fermi level by eV with respect to the Fermi level of the tip. This allows for filled states on the left (sample) to tunnel into empty states on the right (tip). The tunneling current is measured by an external circuit.

Though the dominant tunneling current for negative sample voltage $-V$ will be from sample to tip, there will also be a smaller tunneling current of electrons from tip to sample:

$$I_{\text{tip}\to\text{sample}} = -2e \cdot \frac{2\pi}{\hbar}|M|^2 \underbrace{(\rho_t(\varepsilon + eV) \cdot f(\varepsilon + eV))}_{\substack{\#\text{ filled tip states} \\ \text{for tunneling from}}} \underbrace{(\rho_s(\varepsilon) \cdot [1 - f(\varepsilon)])}_{\substack{\#\text{ empty sample states} \\ \text{for tunneling to}}} \tag{2.3}$$

When we sum these, and integrate over all energies ε, we arrive at the total tunneling current from sample to tip:

$$I = -\frac{4\pi e}{\hbar}\int_{-\varepsilon_F(\text{tip})}^{\infty} |M|^2 \rho_s(\varepsilon)\rho_t(\varepsilon+eV)\left\{f(\varepsilon)\left[1-f(\varepsilon+eV)\right]-\left[1-f(\varepsilon)\right]f(\varepsilon+eV)\right\}d\varepsilon \tag{2.4}$$

We can simplify this expression in several ways. First, all measurements reported in this thesis were taken at $T = 4.2$ K. At this low temperature, the Fermi function cuts off very sharply at the Fermi surface, with a cutoff width of $k_BT = 0.36$ meV. (For reference, the width of the band in BSCCO is several eV, and many of the features studied in this thesis have widths on order several meV.)

In the approximation of a perfectly abrupt cutoff, the integral breaks into 3 parts (where ε is still referenced to the Fermi surface of the sample):

I. $\underline{0 < \varepsilon}$

$$\begin{aligned} &f(\varepsilon) \approx 0; \quad f(\varepsilon+eV) \approx 0 \\ &\Rightarrow f(\varepsilon)\left[1-f(\varepsilon+eV)\right]-\left[1-f(\varepsilon)\right]f(\varepsilon+eV) \approx 0\cdot 1 - 1\cdot 0 = 0 \end{aligned} \tag{2.5}$$

II. $\underline{-eV < \varepsilon < 0}$

$$\begin{aligned} &f(\varepsilon) \approx 1; \quad f(\varepsilon+eV) \approx 0 \\ &\Rightarrow f(\varepsilon)\left[1-f(\varepsilon+eV)\right]-\left[1-f(\varepsilon)\right]f(\varepsilon+eV) \approx 1\cdot 1 - 0\cdot 0 = 1 \end{aligned} \tag{2.6}$$

III. $\underline{\varepsilon < -eV}$

$$\begin{aligned} &f(\varepsilon) \approx 1; \quad f(\varepsilon+eV) \approx 1 \\ &\Rightarrow f(\varepsilon)\left[1-f(\varepsilon+eV)\right]-\left[1-f(\varepsilon)\right]f(\varepsilon+eV) \approx 1\cdot 0 - 0\cdot 1 = 0 \end{aligned} \tag{2.7}$$

Therefore, the relevant range of ε over which we must integrate to find the tunneling current is reduced to $-eV < \varepsilon < 0$. (Likewise, if we had applied a positive bias voltage V to the sample, our range of integration would be $0 < \varepsilon < eV$.) So we are left with approximately:

$$I \approx -\frac{4\pi e}{\hbar}\int_{-eV}^{0} |M|^2 \rho_s(\varepsilon)\rho_t(\varepsilon+eV)d\varepsilon \tag{2.8}$$

In reality equation 2.8 will be modified by an apparent smearing of energy features with $\approx k_BT$ width.

Second, we pick a tip material which has a flat density of states within the energy range of the Fermi surface that we wish to study. For example, if we wish to study the sample density of states within 200 meV of the Fermi surface, then the measured tunneling current will be a convolution of the density of states of the tip and sample in this energy range. So we pick a tip material which has a flat density of states in this range, so that $\rho_t(\varepsilon + eV)$ can be treated as a constant and taken outside the integral.

$$I \approx \frac{4\pi e}{\hbar}\rho_t(0)\int_{-eV}^{0}|M|^2\rho_s(\varepsilon)d\varepsilon \tag{2.9}$$

In practice, most of the measurements reported in this thesis were taken with a tungsten tip, sharpened in situ by field emission onto a gold surface. Tungsten is thought to have a flat density of states near its Fermi surface. As an extra check of this assumption, one can verify that density of states spectra taken with a tungsten tip are flat from -200meV to +200meV at a number of locations on the gold surface (because gold definitely has a flat density of states in this energy range). Other common tip materials include Pt and PtIr.

Bardeen first laid out a basic theory for vacuum tunneling in 1961.[52] Perhaps most importantly, he showed that under the realistic assumptions that (1) the tip and the sample each have their own independent density of states, (2) each of their wavefunctions fall exponentially to zero in the tunneling barrier, and (3) the overlap is small enough (i.e. tip-sample separation is large enough) that each side is insignificantly influenced by the tail of the wavefunction from the other side, then the matrix element for tunneling will be virtually independent of the energy difference between the two sides of the barrier. The matrix will remain unchanged even if one side transitions from the normal state to the superconducting state. In other words, to a reasonable approximation, we can take the matrix element outside the integral and treat it as a constant. Cohen *et al.* in 1962 reformulated this tunneling theory in a slightly easier to read fashion, in terms of a tunneling Hamiltonian.[53]

$$I \approx \frac{4\pi e}{\hbar}|M|^2\rho_t(0)\int_{-eV}^{0}\rho_s(\varepsilon)d\varepsilon \tag{2.10}$$

But what is $|M|^2$? The matrix element comes from the assumption that both tip and sample wavefunctions fall off exponentially into the vacuum gap. Basically, we assume that the vacuum barrier is a square barrier, and we do a WKB approximation. In reality, there will be some tilt to the top of the barrier, but the tilt will be the applied voltage (on order 100 meV) while the height of the barrier will be on order the energy required to remove an electron from a metal, i.e. the work function, of several eV. Therefore, the tilt of the barrier

will be much smaller than the height of the barrier, and can be ignored.

WKB says that the tunneling probability through a square barrier will be $|M|^2 = e^{-2\gamma}$ with γ given by:

$$\begin{aligned} \gamma &= \int_0^s \sqrt{\frac{2m\varphi}{\hbar^2}} dx \\ &= \frac{s}{\hbar}\sqrt{2m\varphi} \end{aligned} \tag{2.11}$$

where m is the mass of the electron, s is the width of the barrier (tip-sample separation), and φ is the height of the barrier, which is actually some mixture of the work functions of the tip and sample.

We can measure the work function by recording the tunneling current as a function of tip-sample separation.

$$I \propto e^{-2s\sqrt{2m\varphi}/\hbar} \tag{2.12}$$

Therefore, we can measure φ from the slope of $\ln I$ vs. s. Typically, we find $\varphi \approx$ 3-4 eV. The higher φ, the more the tunneling current will vary for a given change in s, therefore a higher φ gives a better resolution tip.

Lang made specific calculations in 1987 with a Na atom tip,[54] to show the crossover regime between point contact (where we expect Bardeen's assumptions to break down and universal conductance phenomena to dominate) and vacuum tunneling, where we expect wavefunctions on either side to fall off exponentially. He found for the simplest-to-calculate case of a sodium atom on the end of the tip, this crossover regime occurs around 5-8 Å (nuclear distance between closest tip and sample atoms).

Empirically, our tunneling experiments are occurring in the exponential regime, so we may guess that nucleus-to-nucleus tip-sample separation s exceeds 8 Å. However, due to the exponential fall-off, we have no way of measuring the absolute value of s. This at times causes us much grief, because there is no way to know for sure if we are making measurements at a constant tip-sample separation. So if we see variation from one point on the sample surface to another, we can't be sure whether the variation is due to intrinsic inhomogeneities in the sample at the specific energy of measurement, or due to a varying tip-sample separation.

In summary, the tunneling current is fairly well approximated by:

$$I \approx \frac{4\pi e}{\hbar} e^{-s\sqrt{\frac{8m\varphi}{\hbar^2}}} \rho_t(0) \int_{-eV}^{0} \rho_s(\varepsilon) d\varepsilon \tag{2.13}$$

2.2.2 Measurement Types

Our STM holds the tip at virtual ground, and applies a bias voltage to the sample. The tunneling current is measured from the tip. The tip sits at the end of a piezo tube scanner (with 4 quadrants $\pm x$ and $\pm y$ on the outside, and a single electrode for z on the inside), which can control its motion with sub-Ångstrom precision in 3 directions. The voltage output to the scanner is ± 220 V, with a 16-bit DAC, which is sent through a simple RC filter with rolloff around 3 kHz, to remove high frequency noise leaking out from the electronic control unit. This gives a bit size of $440V/2^{16} \sim 6.7$ mV and a spatial control resolution of approximately 0.03 Å at 4.2K (with a 200 nm total scan range).

We can measure tunneling current as a function of 4 variables: x, y, z, and V. In practice, we attempt to keep z constant by employing a feedback loop to keep the tunneling current constant at a fixed bias voltage. Assuming that z is constant, we then take measurements as a function of x, y, and V.

Topography

The most common mode of STM measurement employed by STM groups around the world is a topography. In this mode, we raster the tip across the surface at a fixed sample bias voltage $-V_{\text{set}}$, and employ a feedback loop which controls the voltage on the z piezo to keep the tunneling current constant at I_{set}. By recording the voltage to the z piezo, we can effectively map the height of the surface.

It's actually not clear what we mean by the "height" of the surface. One obvious suggestion would be some contour of constant charge density. However, as we can see from equation 2.13, the tunneling current does not depend on the total charge density, but only the charge density within eV below the Fermi surface, where $-V$ is the applied bias voltage.

One might argue then that we should apply an arbitrarily large voltage so that we can capture more of the charge density, but we then run into two problems: (1) BSCCO is a fragile compound with weak bonds, and if we apply a large voltage locally, pieces of the surface will literally rip off. (2) If V is too high (on order φ), our tunneling approximation breaks down.

So we are stuck with this somewhat arbitrary definition of the "height of the surface" as the tip-sample separation for which tunneling current is fixed at a particular constant value I_{set} for a particular applied bias voltage V_{set}. In practice we usually choose to fix the current at -100 pA, for a bias voltage of -100 mV. This is arbitrary, but gains some validity from the fact that we do see atoms and other structural features as expected, even over a wider range of choices of I_{set} and V_{set}. The most widely varying density of states features of our BSCCO samples seem to be within 75 meV of the Fermi level.

Density of states

From equation 2.13 we see that if we hold the tip-sample separation constant, at a given (x, y) location, and put a negative bias voltage $-V$ on the sample, we have:

$$I = I_0 \int_{-eV}^{0} \rho_s(\varepsilon) d\varepsilon \tag{2.14}$$

In other words, we can measure the integral of the density of states, down to any energy $-eV$, by varying $-V$. Note that for a negative bias voltage on the sample, we are tunneling electrons from sample to tip, and we are measuring the integrated density of full states below the Fermi level in the sample. For a positive bias voltage on the sample, we are tunneling electrons from tip to sample, and we are measuring the integrated density of empty states above the Fermi level in the sample.

OK, that's nice, we have the integrated density of states (IDOS). But it would be much nicer to just have the density of states (DOS). After acquiring an IDOS vs. V curve, we could take a numerical derivative of our data to get the DOS. But taking a derivative numerically is a horribly noisy thing to do. It is much better to measure the derivative directly.

So we employ a lock-in amplifier to modulate the bias voltage by dV (typically a few mV) around a voltage V of interest. As a result of the voltage modulation dV we can measure a current modulation dI at the same frequency. We call this dI/dV the conductance $g(V)$. Now, we can write:

$$g(V) \equiv \frac{dI}{dV} \propto \text{DOS}(eV) \tag{2.15}$$

Therefore, by using a lock-in and varying V, we can map out an entire density of states curve.

The energy resolution is limited by the amplitude of the wiggle (until the modulation becomes less than $\approx k_B T = 0.36$ meV). So ideally, we could make the voltage modulation smaller than 0.36 mV. But in practice, we can't get enough signal-to-noise at this low amplitude without prohibitively long averaging times. Most of the data reported in this thesis was measured with a 2 mV RMS modulation, therefore blurring our energy resolution by approximately 5.6 meV.

Linecut

In the previous section, we discussed a single DOS curve at a single location. Since we have (x, y) control over the location of our tip using the piezo tube scanner, we can measure DOS curves anywhere we want. Some samples, like good metals (without impurities), should have a completely homogenous DOS everywhere.

But some more interesting samples, for example BSCCO, are inhomogeneous. For example, we can measure a full DOS curve at every point along a straight line, spaced a few Å apart, and we see a "linecut".

DOS map

What we have discussed so far is a 3-dimensional data set: 2 spatial dimensions x and y (by varying the position of the tip) and 1 energy dimension (by varying V). We can view this 3-dimensional data set as a series of DOS-vs-energy curves at every location (x, y), or we could view it as a series of 2-dim DOS-maps at each energy eV. Mapping the DOS at a specific energy is a good visual way to see the inhomogeneities in the density of states. The various ways of viewing the STM measurements are shown in figure 2.4.

Most of the action in the density of states due to the superconducting gap happens within the energy range -100 meV to +100 meV. Most energy features, for example due to impurity resonances, have widths on order a few meV. So in order to get a good picture of the density of states, a minimal spectra has 101 points, which gives 2 meV energy resolution from -100 meV to +100 meV.

The high temperature superconductors have a short coherence length, $\xi \sim 10 - 20$ Å. Therefore, the density of states may be changing significantly on this short length scale. So in order to capture spatial variations in the density of states, we require spatial resolution at least one pixel per 5 Å.

Therefore, for a minimal survey of the density of states in a ~ 500 Å square area, we

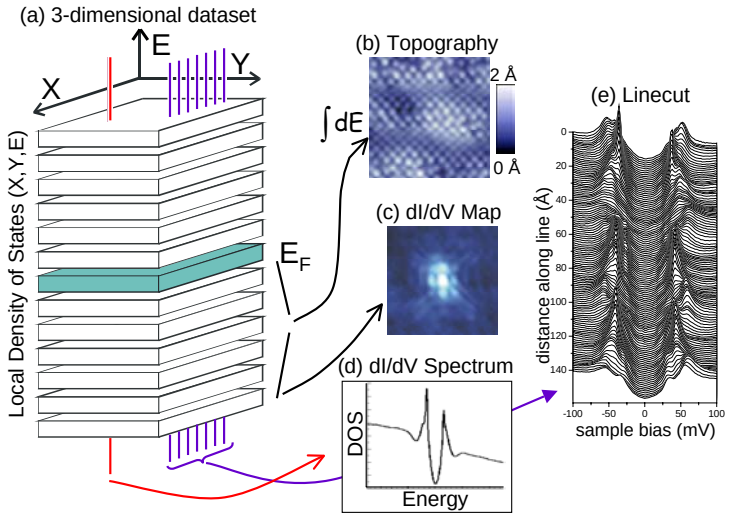

Figure 2.4: (a) An STM has access to an essentially 3-dimensional dataset: two spatial dimensions (x, y) and energy E. This 3-dimensional data space can be explored in 4 different ways (b)-(e). (b) Topography: the tip height z is adjusted through a feedback loop to maintain a constant current I_{set} at a constant bias voltage V_{set}; recording the tip height effectively maps out the height of the surface. (c) dI/dV map: the density of states at a fixed energy E is mapped as a function of position (x, y) on the sample surface. (d) dI/dV spectrum: the density of states as a function of energy is measured at a single point on the sample surface. (e) Linecut: the density of states as a function of energy is mapped at several points along a line on the sample surface.

require 101 (energy) $\times$ 128 $\times$ 128 (spatial), or greater than 10^6 measurements! Even at a meager 100 ms per measurement, such a DOS map requires almost 48 hours. Even making a single measurement in 100 ms is no small feat, because each measurement requires changing the DC bias to a new value V, modulating a small AC bias dV on top, and locking in to the new dI, which necessarily takes several frequency cycles to stabilize after the change in DC bias V. The AC modulation frequency for maps is typically $\sim$ 1.1 kHz, and other carefully optimized parameters such as ramp time, wait time, and measurement averaging time can be found in Eric Hudson's PhD thesis.[55]

Of course, a more detailed DOS map with higher spatial resolution for identification of smaller features (such as short wavelength DOS standing waves, or impurity resonances) may require 101 (energy) $\times$ 512 $\times$ 512 (spatial) measurements, which takes 16 times longer. Therefore the spatial registry must remain locked in place for up to a month (spanning multiple tip retractions for liquid helium transfers)!

2.2.3 Noise Considerations

Vibration Noise

Vibration noise can affect the STM measurement two ways: (1) change in the tip-sample separation, which is amplified exponentially in the tunneling current, and (2) motion of the wire carrying the tip current, which capacitively couples to its environment and therefore causes current spikes when moved. The latter issue can be addressed by carefully clamping all of the wires in place, so that they cannot move with respect to each other due to helium boiling vibrations or external building vibrations. Measurement of external building vibration noise is discussed in more detail in Appendix B.

Electronic Noise

In practice, after fastidious attention to vibration isolation, the largest source of noise is electronic noise. The current preamplifier has a gain of 10^9 which means that the op-amp has a feedback resistor of 1 GΩ. The Johnson noise at room temperature is $0.13\sqrt{R(\Omega)}\text{nV}/\sqrt{\text{Hz}} = 4.11\mu\text{V}/\sqrt{\text{Hz}}$.

In topographic mode, we are essentially integrating over a 3 kHz bandwidth, which amounts to 0.23 mV of rms noise. This is only 0.23 pA of effective current noise, which is generally not a problem in topographic mode where our setpoint current is typically $I_{\text{set}} =$ 100 pA.

In dI/dV mode, the effect of the noise is limited to a band around the frequency of bias voltage modulation. Typically the effective width of this band with feasible averaging times is $\sim$ 100 Hz. The rms voltage noise is therefore 41 μV which means an effective minimum current noise of only 0.04 pA. In fact, due to other imperfections in the electronics, the current noise can be somewhat higher than this. When attempting to measure small features in the density of states, this can be a significant problem.

This current noise could be improved by putting the current preamplifier down on the fridge at 4.2 K. Then the Johnson noise would be reduced by a factor $\sqrt{300\text{K}/4\text{K}} = 8.66$.

Capacitive Coupling

Our STM holds the tip at virtual ground, and applies a bias voltage to the sample. The tunneling current is measured from the tip. Many successful STMs instead bias the tip and measure the tunneling current from the sample. However, the current measurement is

the most sensitive part of the whole operation, and most subject to noise from capacitive pickup. It is necessary to shield the current-measurement side with a ground plane from as many angles as possible, for example to prevent capacitive pickup from the high and rapidly sweeping voltages on the scan piezo. Since the tip is more geometrically compact than the sample (& attached sample holder), and thus easier to shield, it seems logistically easier to bias the sample and measure the tip current.

2.2.4 Inhomogeneity in $\vec{R}$-space vs. Anisotropy in $\vec{k}$-space

An STM is an excellent tool for measuring the density of states as a function of position with sub-Å resolution. Real-space resolution is important for the study of BSCCO in the following experiments: quasiparticle scattering imaging (discussed in chapter 3), magnetic vortex imaging (discussed in chapter 4), and impurity imaging (discussed in chapter 5).

But what about momentum space?

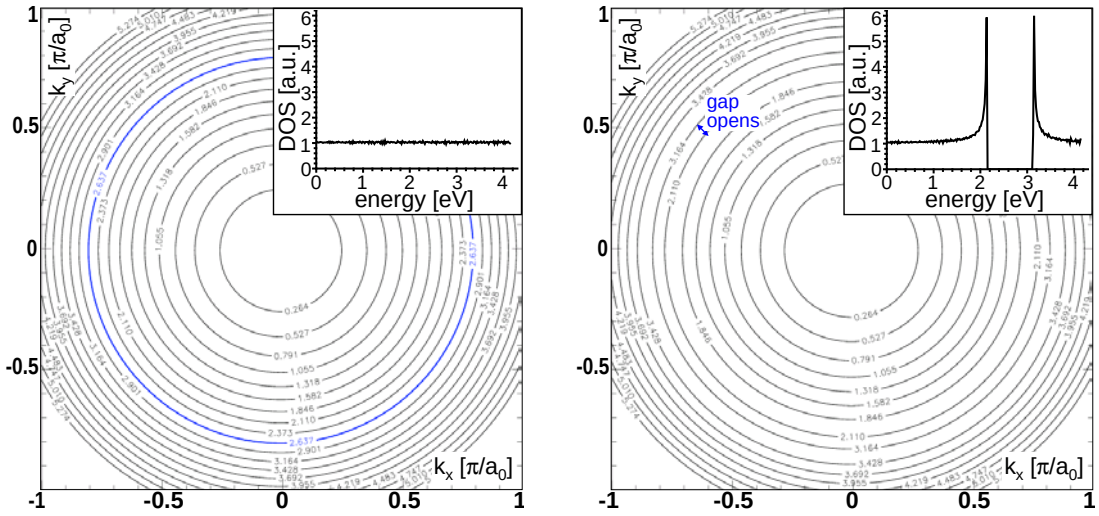

(a) Normal state. (b) Superconducting state (Δ = 500 meV).

Figure 2.5: Isotropic band structure and density of states for a hypothetical perfect s-wave superconductor. (a) The blue line shows the Fermi surface in the normal state. (b) The Fermi surface is completely gapped in the superconducting state. The magnitude of the superconducting gap, $\Delta = 500$ meV, has been exaggerated for visual clarity. The density of states curves shown in the insets are calculated by numerical integration from the band structure contour maps.

An STM tip collects (emits) electrons from (to) all directions in the sample indiscriminately. Therefore, the DOS curves measured by an STM are an average over $\vec{k}$-space.

Why do we care? For many materials, we don't. For example, a conventional low-temperature superconductor has an isotropic s-wave order parameter. This means that the electrons bound into Cooper pairs have an isotropic wave function. They don't care which direction they're traveling; they still see the same binding energy. The energy landscape in the first Brillouin zone for an ideal isotropic s-wave superconductor in the normal and superconducting states (with an exaggerated large gap for clarity) is shown in figure 2.5. The corresponding density of states (found by numerical integration over the contour map) is shown in the insets. When we drop into the superconducting state, we merely push the contours of constant energy (circles) closer together at the Fermi level. The contour corresponding to the Fermi level itself disappears; there are no more states at that energy due to the opening of the superconducting gap. The density of states is the same in every direction; i.e. no matter which direction $\hat{k}$ we take a cut out from zero energy in the center, we see the same thing.

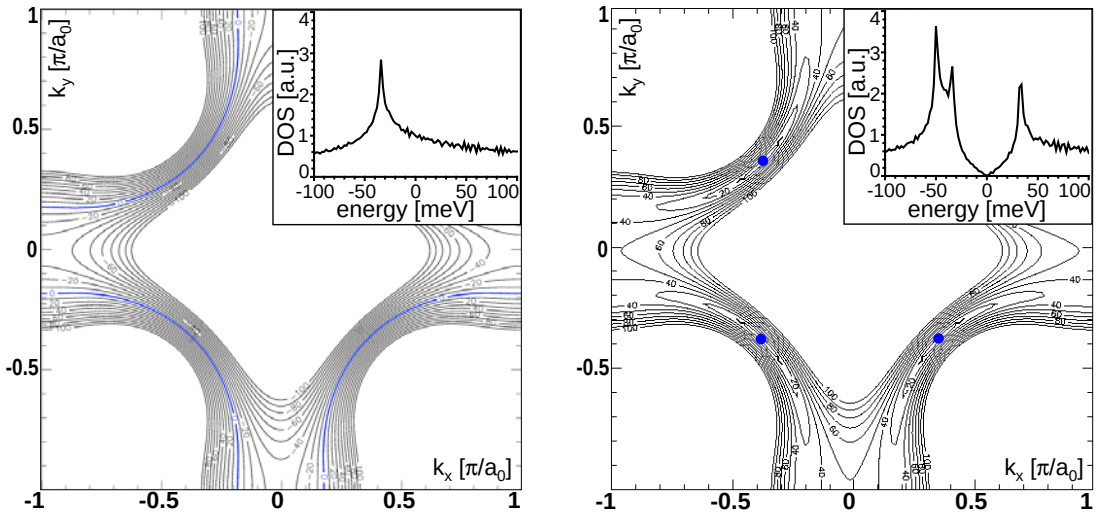

(a) Normal state.

(b) Superconducting state ($\Delta = 36$ meV).

Figure 2.6: Anisotropic band structure and density of states for BSCCO, a $d_{x^2-y^2}$-wave superconductor. The blue line shows the Fermi surface in the normal state (a); the Fermi surface is gapped in the superconducting state (b) except for four nodal points, shown as blue dots. The band structure and magnitude of the superconducting gap $\Delta = 36$ meV are realistic for optimally doped BSCCO. The density of states curves shown in the insets are calculated by numerical integration from the band structure contour maps.

However, BSCCO is a d-wave superconductor; the bound electrons in the Cooper pairs have d-wave, four-fold symmetry. This means that the energy landscape is different in different directions. Figure 2.6 shows the energy landscape of BSCCO in the normal state and in the superconducting state.

In BSCCO, an electron entering the tip from one direction may see a 10 meV square gap, while an electron entering from another direction may see a 30 meV square gap. What we measure with an STM is a sum of all these processes from every direction. The average of all the square gaps gives us a V-shaped gap. Therefore, each individual STM spectrum forfeits the $\vec{k}$-space information. The BSCCO cross-shaped energy landscape in $\vec{k}$-space shown in figure 2.6 was actually deduced from angle-resolved photoemission (ARPES), not from STM.

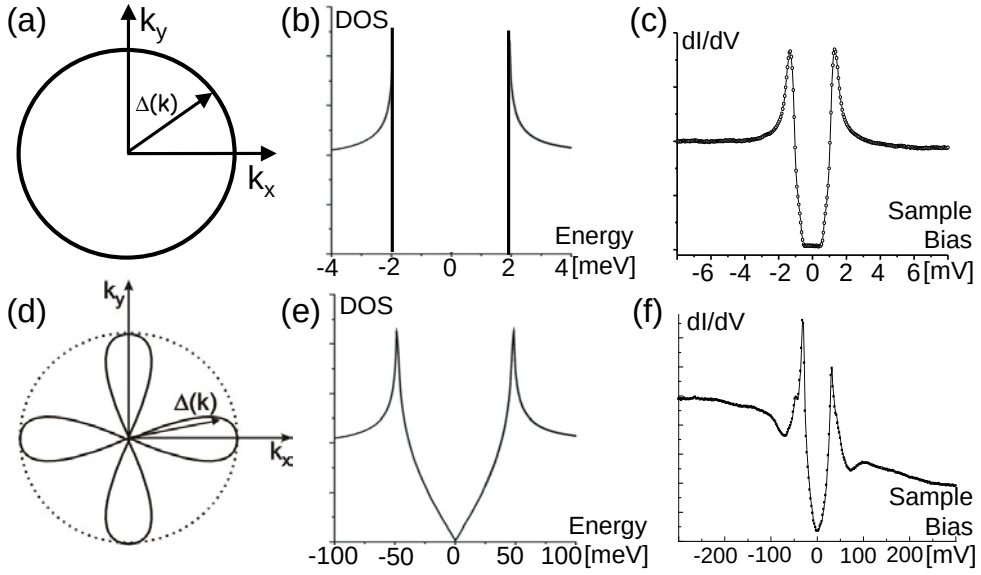

Figure 2.7: Demonstration of the density of states, as seen by an STM, averaged over $\vec{k}$ for s-wave and d-wave superconductors. (a) Gap $\Delta(\vec{k})$ is constant as a function of angle for an s-wave superconductor. (b) Therefore when the STM averages over angle, the resultant density of states still shows a square gap. (c) Real data:[49] density-of-states spectrum on s-wave superconductor $NbSe_2$. The imperfection of the square gap is due in part to thermal broadening, and in part to the very slight anisotropy in the $NbSe_2$ s-wave gap. (d) Gap $\Delta(\vec{k})$ is angle-dependent for a d-wave superconductor. (e) Therefore, when the STM averages over angle, it combines square-gapped spectra with all different values of Δ, and the resultant average shows a V-shaped gap. (f) Real data: typical density-of-states spectrum on d-wave superconductor BSCCO.

2.3 Other Experimental Techniques

There are many other experimental techniques which provide complementary information about the cuprates. In order to better understand the relevant results of these experiments, I will summarize the basic techniques here.

2.3.1 Angle-Resolved Photoemission Spectroscopy

ARPES is a technique for measuring the DOS of a sample with momentum resolution instead of spatial resolution. ARPES takes advantage of the photoelectric effect: shoot in photons with a well known energy, then measure the direction and energy of the emitted electrons. A schematic of the experiment geometry (courtesy of Z.-X. Shen) is shown in figure 2.8. Conservation of energy and momentum allows reconstruction of the energy-momentum relation (dispersion relation) of the electrons within the crystal.

$$E_{\text{kin}} = h\nu - \phi - |E_B| \tag{2.16}$$

$$\vec{p}_{\parallel} = \hbar\vec{k}_{\parallel} = \sqrt{2m_e E_{\text{kin}}}\sin\vartheta \tag{2.17}$$

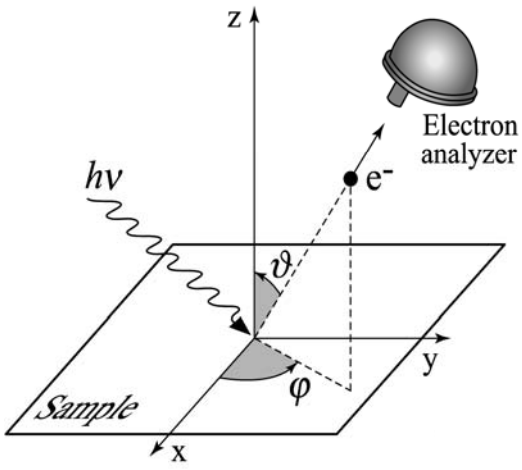

Figure 2.8: Angle-resolved photoemission spectroscopy (ARPES) experimental geometry. Figure borrowed from Damascelli *et al.*[56] The incoming photon has significant energy $h\nu$ but negligible momentum. The energy and momentum of the outgoing electron are measured with a high-resolution Scienta detector (the hemisphere shown).

A few caveats:

(1) ARPES works well only on 2-dimensional materials. If the electrons in the material are moving around with an unknown z component to their momentum as well as the x and y components, then it is impossible to reconstruct the dispersion relation: there are too few measurables, and too many variables.

(2) ARPES spatially averages. Therefore, all of the compelling reasons for studying the spatial dependence of BSCCO DOS (such as electronic inhomogeneity, vortices, and impurities) are unaddressed. Spatial averaging can also create apparent spurious effects in the data. For example, suppose we had some short-length-scale segregation between two different phases. Suppose that one phase had a narrow gap width with sharp coherence peaks, while the other had a wide gap with rounded gap-edge peaks. If we spatially average, we will come up with an intermediate spectrum that looks nothing like any true spectrum in the crystal.

Spatially averaging can be deceiving in more subtle ways. Suppose we have a continuous distribution of gap widths Δ, each with sharp coherence peaks. But if we spatially average, we will arrive at a spectrum with broader coherence peaks. The width of the coherence peak is inversely proportional to the quasiparticle lifetime. So an ARPES measurement on an inhomogeneous surface could give a false impression of a very short quasiparticle lifetime. But ARPES experiments on BSCCO do see long quasiparticle lifetimes in some parts of $\vec{k}$-space and short quasiparticle lifetimes in other parts.

(3) ARPES measures only filled states (i.e. states with $\varepsilon < \varepsilon_F$), not empty states above the Fermi level. This is because ARPES measures only photo-ejected electrons; an electron can only be ejected at a certain energy if there was a filled state in the crystal at the corresponding energy to begin with. Unlike STM, ARPES cannot measure empty states by reversing the process and putting electrons back into the sample.

(4) ARPES cannot work in a magnetic field, because it depends on the ejected electrons traveling in a straight line from the sample to the detector at the angle of interest.

Fermi surface

Despite these caveats, ARPES has been tremendously successful at mapping out the Fermi surface of many 2-dimensional materials, such as high-T_c superconductors, by measuring number of ejected electrons vs. energy at many different detector angles.

What ARPES actually measures is the spectral function $A(\vec{k}, \omega)$ times the Fermi func-

tion $f(\omega)$ times some matrix element $I_0(\vec{k}, \nu, A)$ which may have some $\vec{k}$ and A dependence and also some dependence on ν, the energy of the incoming photon. A more detailed discussion of the ARPES matrix element may be found in a comprehensive review of ARPES on high-T_c superconductors by Damascelli *et al.*[56] In practice, ARPES experiments typically make measurements at several different incoming photon energies and polarizations, to verify that any effects they may be seeing are not dependent on a specific incoming photon energy $h\nu$. As for matrix-element dependence on $\vec{k}$ and A, several approximations show these should not be so large as to overwhelm the useful $A(\vec{k}, \omega)$ information.

For our purposes, we note that both STM and ARPES have matrix elements which may influence experimental results. But the important point is that the matrix elements have totally different physical origins and totally different forms and dependencies for the two experiments. So if the same result can be found from both STM and ARPES, then we can probably conclude that the matrix elements are not affecting our ability to extract the density-of-states information from either experiment.

So let's assume ARPES is just measuring the spectral function times the Fermi function. By measuring this spectral function along a line in $\vec{k}$-space which crosses the Fermi surface, they will see a peak in emitted electron intensity which disperses towards the Fermi level, then disappears as $\vec{k}$ (i.e. the angle of the detector) passes through the Fermi surface. The intensity disappears above the Fermi surface, because the Fermi surface is *defined* as the location above which there are no more filled electron states. And if there are no filled electron states, then no electrons can be ejected. So by mapping the locations in $\vec{k}$-space at which the intensity peaks disappear, ARPES can map the location of the Fermi surface. ARPES also maps the dispersion E vs. k for the electrons below the Fermi surface.

Below T_c

Below T_c in a d-wave superconductor, things get more interesting because an anisotropic gap opens up in the density of states at the Fermi level. ARPES can measure the energy of the gap as a function of position in $\vec{k}$-space by using gold (in contact with the BSCCO) as a reference for the real Fermi level, and noting the shift in the energy where the BSCCO peak disappears vs. the gold peak. What they find is that below T_c, the gap is zero only at 4 points, along the a and b crystal axes (the "nodal" points), and largest along the x and y axes (the "anti-nodal" direction). The angle dependence of the gap has been mapped out by Ding, *et al.*[57] and other groups.[58, 59, 60, 61, 62, 63]

In summary, we have from STM the energy landscape of BSCCO in $\vec{R}$-space, and from

ARPES the energy landscape of BSCCO in $\vec{k}$-space.

2.3.2 Neutron Scattering

Neutron scattering is the best way to obtain $\vec{k}$-space information about the magnetic properties of a material. A good introduction to the experimental techniques of neutron scattering can be found in a conference presentation by Boothroyd.[64]

Neutron scattering is a bulk probe with no spatial resolution (but with the benefit that it is definitely not sensitive to surface effects or to the specifics of cleaving procedure). Neutrons are so weakly interacting that a sample of size ~ 1 cm^3 is required for adequate signal strength. Therefore most neutron scattering experiments are performed on LSCO, the best material for growth of large crystals.

Neutrons can probe two different aspects of a crystal: the locations of the atomic nuclei, or the magnetic structure of the electrons. Elastic neutron interactions with the nuclei give us Bragg diffraction peaks, which enable determination of the crystal structure. Inelastic neutron interactions with the nuclei give the phonon dispersion relations of the crystal.

For comparison with STM searches for alternative phases in HTSCs, the neutron interactions with the electron spins are more relevant. These elastic and inelastic interactions measure the static and dynamic electron spin ordering in the crystal, respectively. Specifically, neutron scattering off electron spins measures the magnetic susceptibility as a function of $\vec{k}$-space.

The neutron scattering setup is similar to ARPES: columnated, monochromatic neutrons are incident on the cuprate crystal parallel to the CuO_2 plane. The energy-resolved neutron detector is then moved through a range of solid angles above the crystal, and the number of incident neutrons are counted for each solid angle interval $d\Omega$ and (in the case of inelastic neutron scattering) each energy interval dE.

Inelastic Neutron Scattering

An inelastic neutron scattering event involves both a momentum transfer and an energy transfer:

$$\hbar\mathbf{Q} = \hbar(k_i - k_f) \tag{2.18}$$

$$\hbar\omega = E_i - E_f = \hbar^2(k_i^2 - k_f^2)/2m \tag{2.19}$$

The scattering event is therefore characterized by $(\mathbf{Q}, \omega)$. The scattering probability is represented by the differential cross-section:

$$\frac{d^2\sigma}{d\Omega dE_f} = \frac{k_f}{k_i} S(\mathbf{Q}, \omega) \tag{2.20}$$

The k_f/k_i factor is sometimes important, for example if the neutron loses a lot of energy ($k_f \ll k_i$) then the intensity is much reduced. But $S(\mathbf{Q}, \omega)$ contains all the physics of the system, and is given by the fluctuation-dissipation theorem:

$$S(\mathbf{Q}, \omega) \propto \underbrace{\frac{1}{1 - e^{-\hbar\omega/k_B T}}}_{\text{Bose factor}} \underbrace{\text{Im}\{\chi(\mathbf{Q}, \omega)\}}_{\chi''(\mathbf{Q}, \omega)} \tag{2.21}$$

where $\chi(\mathbf{Q}, \omega)$ is magnetic susceptibility, i.e.:

$$M(\mathbf{Q}, \omega) = \chi(\mathbf{Q}, \omega) H(\mathbf{Q}, \omega) \tag{2.22}$$

Inelastic neutron scattering has been used to look for dynamic spin fluctuations in LSCO[65] and YBCO[66] and recently even BSCCO[67] (albeit with much larger linewidths than LSCO or YBCO).

Elastic Neutron Scattering

Elastic neutron scattering is really just a special case of inelastic neutron scattering, with $\omega = 0$. Therefore, elastic neutron scattering measures the static spin ordering. The technique is analogous to mapping the positions of nuclei in crystals using X-ray or neutron diffraction, except it is the static spins of the electrons which are being mapped.

Studies of static spin ordering have been carried out in LSCO[46] and LCO.[68]

2.3.3 Nuclear Magnetic Resonance

Nuclear magnetic resonance (NMR) can explore the spatial distribution of magnetic field within a sample. Reviews of NMR on high-T_c superconductors have been written by Rigamonti *et al.*[69] and Berthier *et al.*[70]

For a material such as BSCCO with a complicated unit cell, NMR can individually access the magnetic environment at each atomic species, because each species will have a

different Larmor precession frequency. The linewidth around the central frequency for each species will give information about the degree of disorder and local field distribution in the sample. In fact, large NMR linewidths in BSCCO were some of the first evidence for gross electronic inhomogeneity.

NMR also measures the inverse spin-lattice relaxation time $1/T_1$, which is a measure of spin fluctuations. NMR studies showing additional spin susceptibility at low temperature in BSCCO indicated strong electronic disorder.[71] Since YBCO is a more ordered crystal than BSCCO (the oxygen dopants fit into known crystal lattice sites in the chain plain, instead of squishing in interstitially like BSCCO), the narrower linewidths in YBCO make it easier to extract quantitative NMR information. Therefore, most NMR studies on the cuprates have been carried out on YBCO.

NMR is very useful for studies of vortices in superconductors. The Larmor frequency of the probe nucleus is a measure of the local field. Therefore, the full Larmor lineshape can be seen as a sort of "histogram" of atomic locations relative to the vortex center.[72] This can be used to determine the spatial symmetry of the vortex lattice (triangular, square, etc.), or the properties of quasiparticles existing in different local fields.

2.4 Summary

In summary, this thesis documents the use of a scanning tunneling microscope (STM) to study the real-space dependence of the density-of-states in $Bi_2Sr_2CaCu_2O_{8+\delta}$ (BSCCO), with sub-Å resolution. Angle-resolved photoemission spectroscopy (ARPES) is a complementary technique which accesses the $\vec{k}$-space dependence of the density-of-states in BSCCO. Both STM and ARPES are surface sensitive techniques which require clean, flat surfaces for effective study. Of the 3 families of hole-doped high-T_c superconductors, BSCCO is easiest to cleave and produce such a clean, flat surface.

Neutron scattering and nuclear magnetic resonance (NMR) are two more techniques which access the magnetic information in a crystal. Both are bulk probes, insensitive to surface effects. Neutron scattering measures the dynamic and static spin ordering in the crystal in $\vec{k}$-space, and is typically limited in application to large single crystals ($\sim$ 1 cm) such as the lanthanum family of cuprates. NMR measures the distribution of magnetic fields in a sample, and is a useful tool for studying vortices. NMR is typically most effective on well-ordered crystals such as $YBa_2Cu_3O_{7-\delta}$ (YBCO).