

**Investigating Atomic Scale Disordered Stripes
in the Cuprate Superconductors
with Scanning Tunneling Microscopy**

A dissertation presented

by

Elizabeth Main

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

December 2011

©2011 - Elizabeth Main

All rights reserved.

Thesis advisor

Jennifer E. Hoffman

Author

Elizabeth Main

**Investigating Atomic Scale Disordered Stripes
in the Cuprate Superconductors
with Scanning Tunneling Microscopy**

Abstract

The high- T_c cuprate superconductors have been studied for 25 years in the search for the mechanism underlying their superconductivity. In the process, experiments learned that the correlated electrons in these materials organize themselves in a variety of patterns. One such pattern is a type of short-range charge modulations that exist both in and outside of the superconducting phase, which has been linked to the cuprate pseudogap phase. In optimal and slightly underdoped $\text{Bi}_2\text{Sr}_2\text{CuO}_{6+\delta}$ this charge order (labeled Q^*) has wavelength $\sim 5a_0$. A second and, I argue, related order (Q^{**}) has wavelength of $\sim 1.25a_0$. These modulations are highly disordered and for this reason their nature is not fully understood. In this thesis I use scanning tunneling microscopy (STM) to study the disorder of these charge modulations, as an avenue to understanding the nature of the charge order itself.

Locally, the charge modulations have a preferred orientation, with a wavevector pointing along one crystal axis or the other. But globally, there is no preferred direction. Our most striking finding is that the local orientation is the same for the

Q^* and Q^{**} orders, strong new evidence that these two types of charge modulations have the same physical cause. Next, we find that Q^* and Q^{**} are subject to two kinds of disorder. Disorder in the optimal local wavelength competes with defect pinning of crests and troughs to produce the disordered modulations that we see. To get our final result, I view the local orientation of the charge modulations as an Ising spin, and compare the resulting Ising maps to theoretical predictions for different classes of disorder. I find the disorder to be consistent with 3D Random Field disorder.

New analytical tools were necessary to carry out these measurements. I describe a new algorithm to map the local wavevector of a modulation. Then I present a second new algorithm to correct an STM image for the effects of a slightly anisotropic tip. This thesis also describes the construction of a new spin-polarized STM that might (but hasn't yet) shed further light on electronic order in the cuprates.

Contents

Title Page	i
Abstract	iii
Table of Contents	v
Dedication	vii
1 Introduction	1
1.1 Background and Motivation	1
1.2 Just Enough STM Background	3
1.3 Materials	6
1.3.1 Datasets	8
1.4 STM Observations of Charge Order in BSCCO	9
1.4.1 Note about Nomenclature	11
1.5 Outline	12
2 Design and Construction of a Potentially Spin-Polarized STM	13
2.1 System Overview	14
2.1.1 Requirements of a Spin-Polarized STM	14
2.1.2 Other Design Features	14
2.2 STM Scan Head	16
2.2.1 Rotor	17
2.2.2 Piezo Tube Assembly	20
2.2.3 Set Screws	21
2.3 In Situ Sample and Tip transfer	22
2.4 Atomic Resolution Tests	25
3 Spatial Lockin: A New Algorithm for Mapping Local Wavevector	26
3.1 Background	26
3.2 Mathematical Explanantion	27
3.2.1 Finding the local fourier component	27
3.2.2 Choosing the Best Local Wavevector	29
3.2.3 Choosing L	30

3.3	Additional Steps	33
3.3.1	Preprocessing	34
3.3.2	Postprocessing	36
3.4	Applying the Spatial Lockin	36
3.5	Variations on the Spatial Lockin	37
4	The Disordered Local Wavevector of the Q^* and Q^{**} Peaks	40
4.1	Materials	44
4.2	Local Wavevector	51
4.2.1	Correlation Coefficients	51
4.2.2	Standard Deviations	55
4.3	Conclusion	56
5	Correcting for Anisotropy of the STM Tip	59
5.1	Motivation	59
5.2	Quantifying Tip Anisotropy	60
5.3	Stoll Formula	63
5.4	Recipe for Tip Anisotropy Correction	65
5.4.1	Tip Angle	67
5.5	Limitations	69
5.6	Results	70
5.7	Conclusion	73
6	Local Anisotropy of the Q^{**} Peak	76
6.1	Background	76
6.2	Stripes or Checkerboard?	78
6.3	Mapping local anisotropy	80
6.4	Finding critical exponents	84
6.5	Results: Discussion of critical exponents	90
6.6	Conclusion	92
7	Conclusion	95
	Bibliography	98

Dedicated to my family on the West coast.

Chapter 1

Introduction

1.1 Background and Motivation

It has now been 25 years since the discovery of high T_c superconductivity in the family of materials known as the “cuprates.” Yet the mechanism that allows electrons to pair—and therefore to superconduct—remains a mystery. Understanding this phenomenon might allow scientists to create materials with dramatically higher T_c a feat which would have a profound impact on technology.

The cuprate superconductors are formed by adding dopant atoms to a parent compound. In the absence of doping, the material neither conducts nor superconducts. It belongs instead to a class of materials known as Mott insulators, whose electrons should naively be free but are instead confined to their lattice site by the Coulomb repulsion of neighboring electrons. This correlation between electrons leads to the phase diagram shown in Fig. 1.1a. With few or no dopant atoms (at the left of the phase diagram), the cuprates are in the antiferromagnetic Mott phase. When the

dopant fraction reaches about $p = 0.08$, superconductivity arises at the lowest temperatures. In between lies the “pseudogap” phase, so called because the electronic spectrum is partially gapped around the Fermi surface but this gap doesn’t correspond to zero resistance. The doping-dependent temperature T_c denotes the upper bound of the superconducting phase, while T^* is the upper limit of the pseudogap phase. The pseudogap phase may encompass several regions in which electrons organize themselves into different patterns. Fig. 1.1b displays one of many possible phase diagrams showing these ordering tendencies, intended to overlap with the diagram in a. These phases likely involve breaking different symmetries of the host crystal. While studying Mott insulators may bring us closer to understanding high T_c superconductivity, the ordering susceptibilities of a Mott insulator are a rich condensed matter puzzle in their own right!




Figure 1.1: **(a)** Generic phase diagram of the high- T_c cuprate superconductors. **(b)** Speculative phase diagram that includes several different electronic orders. Reproduced from Kivelson *et al* [12]

Let us narrow our scope to a particular form of electronic stripe order, in which electrons work together to break both the translational and rotational symmetry of the host crystal. Below their phase transition we call them “static.” “Fluctuating” stripes may exist above but near the phase transition, but they show no broken symmetry at length scales greater than some coherence length ξ . Bulk probes (neutron [33], [8] and x-ray scattering [2]) have observed static long-range stripe order only in a particular family of cuprates (LSCO and LBCO). In these compounds charge density waves (CDW) exist near $1/8$ doping, have wavelength of $4a_0$ (where a_0 is the length of one unit cell), and are stabilized by structural distortion of the lattice [35]. The cartoon in Fig. 1.2 displays the particular susceptibility to stripes in compounds near $1/8$ doping, and illustrates why this charge order has four unit cell periodicity (and why the associated spin density waves have $8a_0$ wavelength). The wavevector is in the direction of the Copper-Oxygen bond. T_c is reduced markedly in these compounds near $1/8$ doping [19], [8], suggesting that stripe order competes with superconductivity. Weaker signatures of incommensurate stripes have been observed at nearby dopings¹. In the compound $\text{Bi}_2\text{Sr}_2\text{Ca}_{n-1}\text{Cu}_n\text{O}_{2n+4+\delta}$ (BSCCO), the focus of this dissertation, static long-range-ordered charge stripes have not been observed (and not for lack of trying! [5], [40].)

1.2 Just Enough STM Background

Before detailing STM observations of charge modulations in the cuprates, I will pause and briefly explain the experimental setup in an STM, and the particular

¹See references in the excellent reviews in Refs. [12] and [35]




Figure 1.2: Cartoon of charge and spin stripes in a doped Mott insulator, illustrating the susceptibility to stripe order at $p = 1/8$ doping. Each site represents the $3d_{x^2-y^2}$ orbital of a copper atom. Each arrow represents an electron with a particular spin direction. Empty circles are holes. Reproduced from Vojta [35]. This cartoon depicts site-centered stripes; Ref. [29] proposes alternative diagrams of bond-centered and bond-ordered charge modulations.

cuprate compound that we most like to study. Excellent introductions to tunneling microscopes exist in the literature; this does not attempt to be one of them (see, for example, Refs. [6], [32], [36]). But it would be remiss not to mention the very basics of STM.

Figure 1.3 shows a schematic of the setup of a scanning tunneling microscope. A voltage V is applied between a conductive sample and a tip. The resulting tunneling current I falls off exponentially with increasing tip-sample separation leading to atomic resolution. The tip is scanned along the surface with sub-Angstrom precision using a piezoelectric scan tube. The design and construction of a particular STM will be detailed in Chapter 2.

There are several different measurements that an STM can perform. In a topography, the tip scans the sample while a feedback loop varies the height of the tip in order to keep the current constant. Alternately, we can take a spectrum by turning




Figure 1.3: **(a)** Schematic of a scanning tunneling microscope (STM). A voltage V is applied between a conductive sample and an atomically sharp tip. The two are brought less than a nanometer apart. The resulting tunneling current I is measured. The current falls off exponentially with increasing tip-sample separation. **(b)** Anatomy of a Density of states map (DOS map).

off the feedback loop (holding height constant), sweeping the voltage and measuring the resulting current. To a good approximation the current is proportional to the integrated density of states IDOS:

$$I \propto e^{-\kappa z} \int_{-eV}^0 \text{DOS}(\epsilon) d\epsilon \equiv e^{-\kappa z} \text{IDOS}(V). \quad (1.1)$$

Typically we use a lockin method to measure the conductance directly:

$$dI/dV \propto e^{-\kappa z} (\text{DOS}(V)) \quad (1.2)$$

where $\text{DOS}(V)$ is shorthand for the density of states at energy $-eV$. In these equations, I have tacitly taken the low-temperature limit. All of the data in this disser-

tation was taken at low temperature ($\approx 6\text{K}$) leading to thermal broadening of less than 1 meV.

The third and most important type of image, known as a DOS map, is a combination of these two types. The voltage is set at V_{set} . For each pixel in the map, the tip goes into feedback, forcing the current to be I_{set} . Then the feedback loop is turned off and a dI/dV spectrum is taken. Afterward, the feedback loop is turned on while the tip travels to the position of the next pixel. The resulting three-dimensional set of data is diagrammed in Fig. 1.3. Each layer maps the density of states at a particular voltage. Such maps often have 200-500 pixels in each direction, and each spectrum takes around 1 second. So these maps can take up to several days. STM is truly a static probe!

It can be shown that at each pixel,

$$dI/dV \propto \frac{\text{DOS}(V)}{\text{IDOS}(V_{\text{set}})} \quad (1.3)$$

a phenomenon known as the “setup condition.” The integrated density of states $\text{IDOS}(V_{\text{set}})$ generally varies from point to point on the sample. If there are non-dispersive modulations at any subset of energies between $V = 0$ and V_{set} then IDOS will have a component with the periodicity of the modulations. For this reason, it is hard to trust the energy dependence of stripe phenomena as measured by STM.

1.3 Materials

The cuprates are layered ceramic materials related by the existence of one or more copper oxide (CuO_2) planes per unit cell. It is in the essentially two-dimensional CuO_2




Figure 1.4: **(a)** Crystal structure of $\text{Bi}_2\text{Sr}_2\text{CuO}_{6+\delta}$ (Bi-2201). **(b)** Crystal structure of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ (Bi-2212). The entire orthorhombic unit cell is shown in both (a) and (b).

layer that superconductivity emerges. There are many members of the $\text{Bi}_2\text{Sr}_2\text{Ca}_{n-1}\text{Cu}_n\text{O}_{2n+4+\delta}$ (BSCCO) family, two of which are explored in this dissertation. The first, $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ (Bi-2212), has two CuO_2 planes per unit cell, separated by a SrO layer. The second, $\text{Bi}_2\text{Sr}_2\text{CuO}_{6+\delta}$ (Bi-2201) has only a single CuO_2 plane in each unit cell. The crystal structures are shown in Fig. 1.4, and topographic images of the surface are shown in Fig. 1.5. Additionally, all of the Bi-2201 samples described in this dissertation have been doped with lead, which completely

suppresses the supermodulation (the structural modulations with wavelength $\tilde{6}.5a_0$ along the crystalline b-axis, that can be clearly seen in Fig. 1.5b.) The crystals cleave between the two BiO planes; each white dot in the topography is a Bismuth atom. The bright atoms in the Bi-2201 topography are the Pb dopant atoms, which substitute for Bi. Both materials have approximately tetragonal symmetry, although their true symmetry is orthorhombic.




Figure 1.5: Topographic images of BSCCO, each 30nm wide. **(a)** Lead-doped Bi-2201. **(b)** Bi-2212

1.3.1 Datasets

Table 1.1 lists all of the datasets analyzed in this thesis. Altogether there are six different BSCCO samples imaged by three different instruments.

Table 1.1: **Table of datasets used in this thesis**

label	C	A	B	D	E	F
T_c	UD25K	UD32K	OPT35K	OD16K	OD15K	Bi-2212 OPT91K
material	Pb,Bi-2201	Pb,Bi-2201	Pb,Bi-2201	Pb,Bi-2201	Pb,Bi-2201	Bi-2212
nominal doping	0.101	0.128	0.16	0.241	0.243	0.16
size (nm)	52	68	72	36	56	50
size (pixels)	384	400	468	236	300	250
V_{set} (mV)	-150	-200	-100	100	-100	-120
R-map voltage	80	80	80	45	80	80
instrument	Hudson1	Hudson1	Hudson1	Hoffman1	Hoffman1	Hoffman2

1.4 STM Observations of Charge Order in BSCCO

Many Scanning Tunneling Microscopy (STM) experiments have observed charge modulations near $4a_0$ periodicity that may be due to fluctuating stripes [9], [10], [11], [34], [16], [14], [38], [39], [22]. An example from our own data is shown in Fig. 1.6, showing a “checkerboard” pattern made up of charge modulations in two orthogonal directions. We discount the possibility of static order, because the latter would cause a “giant resonance” in scattering experiments [2], which has not been seen in BSCCO, even near $1/8$ doping [5], [40]. Fluctuating stripes may be visible to STM, a static probe, when they are pinned by disorder. Several STM experiments on BSCCO have seen static or weakly-dispersive charge modulations near to $4a_0$ wavelength [9], [10], [11], [34], [38], [39], [22], [14]), that may be due to fluctuating stripes. This is very distinct from the highly dispersive low-energy modulations due to

quasiparticle interference (QPI). Other attributes of the stripes seen by STM are that they exist both above and below T_c [34],[22], they don't exist above T^* [22], and their incommensurate wavevector is linked to the size of the pseudogap [39]. Tantalizingly reminiscent of the static stripes in LSCO and LBCO, the amplitude becomes large near 1/8 doping [22].




Figure 1.6: **(a)** 30mV dI/dV map, 50nm wide, showing a clear checkerboard pattern. The image has been drift-corrected according to Lawler's algorithm [15]. **(b)** Fourier transform of (a). The stripe peaks Q^* are marked with blue arrows; their satellite reflections Q^{**} are marked with green. The atomic Bragg peaks Q_x and Q_y are shown with red.

This thesis will investigate stripes at both the usual $\approx 4a_0$ periodicity (we call this the Q^* peak), and a possibly related order of wavelength $\approx 4a_0/3$ known as Q^{**} . Both are labeled in the Fourier transform of Fig. 1.6. Although these modulations have been examined by STM before (the Q^{**} peak has been imaged in Refs. [13], [14], [15], [22], [18], while observations of Q^* were chronicled in the previous paragraph), we see two main motivations for studying them further. First, it an open question whether

Q^* and Q^{**} are ordering vectors of the same CDW. Second, there is more to learn from the disorder of these modulations. By further characterizing both the stripes and the disorder that plagues them, we hope to better understand the cuprates they inhabit.

1.4.1 Note about Nomenclature

The Q^* and Q^{**} peaks have been called many different things in the literature. Early on, it was thought that the order was commensurate with the lattice with period $4a_0$. So they are sometimes referred to as the “one-quarter” and “three-quarter” peaks. Kohsaka refers to them as q_1^* and q_5^* after their QPI peak cousins [14]. Lawler labels the Q^* peak as “Smectic” order with wavevector S [15], while Wise calls the Q^* order the “checkerboard” [39]. In this thesis, I follow Parker [22] and label them Q^* and Q^{**} . For convenience, I will sometimes refer to these modulations as stripes and sometimes as a CDW, although they do not have every characteristic of a CDW (for instance, the wavevector seems to weakly disperse with energy) and we have not demonstrated that they have every characteristic of a stripe, (e. g. we have not done spin-sensitive measurements.) Unless otherwise stated, I intend these labels only in the loosest sense. as periodic modulations in the density of states.

Table 1.2: **Other authors’ names for the Q^* and Q^{**} order**

Author	Q^* symbol	Q^{**} symbol	name of order	Reference
Kohsaka	q_1^*	q_5^*	Electronic cluster glass	[13],[14]
Lawler		S	Smectic	[15]
Mesaros		S	Smectic	[18]
Wise	(no symbol)		Checkerboard	[38],[39]
Parker	Q^*	Q^{**}	Fluctuating stripes	[22]

1.5 Outline

This dissertation will begin with a description, in Chapter 2, of the low-temperature, UHV, and potentially spin-polarized STM built by the author. Following this, Chapter 3 will describe new analytical techniques to determine the local wavevector of the stripe order. In Chapter 4, we use this local wavevector technique to elucidate the relationship between Q^* and Q^{**} , and to explore the nature of the disorder in these stripes. Chapter 5 describes a new algorithm allowing us to correct for non-ideal deviations to the shape of the STM tip, paving the way for Chapter 6 in which we examine the rotational symmetry-breaking characteristics of the stripe order. Finally, Chapter 7 summarizes and discusses the results presented in this dissertation.

Chapter 2

Design and Construction of a Potentially Spin-Polarized STM

This chapter describes the main design features of an STM system I built¹, intended to be a spin-polarized (SP), cryogenic, variable-temperature, UHV, high-magnetic-field, spectroscopic-imaging (SI) STM. At the present time, it functions as a low-noise, cryogenic, high-field SI-STM, capable of studying cleavable materials. Many of its other capabilities have been designed into the instrument, but require more work to bring them online. This chapter assumes that the reader has some basic knowledge of STM construction, however I will reference useful introductory material along the way.

¹With lots of help from Adam Pivonka and Ilija Zeljkovic!

2.1 System Overview

2.1.1 Requirements of a Spin-Polarized STM

A spin-polarized STM requires that the tip be magnetized. Specifically, the atom on the tip that is closest to the surface, the one doing the tunneling, needs to have a well-defined magnetization. This has been accomplished by a few other labs, with varying degrees of success, by making the tip out of bulk ferromagnetic or anti-ferromagnetic materials (reviewed in, e.g. [4]). But the most reliable way to get a spin-polarized tip with limited stray fields and a reproducible axis of magnetization is to coat the tip with magnetic films a few monolayers thick [4], [37]. This feat requires that the tip be transferred from a room-temperature UHV materials-preparation chamber, to the cryogenic STM, without breaking vacuum and without crashing the tip. Integrating in situ tip transfer with other constraints was the major design challenge in building this STM.

2.1.2 Other Design Features

The instrument has a highly constrained design, with many interacting systems. The most constraining features are the need for in situ tip and sample transfer; a 9T magnetic field, which requires that the scan head fit inside a small magnet bore and be built with nonmagnetic materials; the ability to operate at UHV at room temperature; and a cryogenic variable temperature fridge. All of these were accomplished with several levels of vibration isolation.

Figure 2.1 shows the experimental setup. The main vacuum chamber sits atop a




Figure 2.1: Photo of the experimental apparatus. The main vacuum chamber (“transfer chamber” sits atop a triangular table built on three air springs. Two UHV magnetic manipulators, used to access the cryostat, stick out from the top, while a load-lock and horizontal manipulator extends out the right side. Below the table hangs the variable temperature Janis cryostat, with two lower indium seals removed to expose the STM scan head (white.) A liquid helium dewar (not pictured) is raised around the cryostat before the instrument is cooled down. *Inset:* A schematic of the experimental table, with cryostat closed and dewar in place.

triangular table built on three air springs. Two UHV magnetic manipulators, used to access the cryostat, stick out from the top, while a load-lock and horizontal manipulator extends out the right side. Below the table hangs the variable temperature Janis cryostat, with two lower indium seals removed to expose the STM scan head (the small white object at bottom, center of Fig. 2.1.) A liquid helium dewar is raised

around the cryostat; the inset in Fig.2.1 shows the position of the dewar when an experiment is in progress. The table sits on air springs inside an acoustically-isolated room. A two axis superconducting solenoid (9T vertical, 2T horizontal) sits at the bottom of the dewar; the STM scan head fits in its bore.

2.2 STM Scan Head

A good STM needs to be very rigid, coupled to its surroundings in a very non-rigid way. The cartoon in Fig. 2.2 illustrates this principle. If you drive the grey block (“outside world”) at a frequency ω , and then solve for the relative amplitude of motion $z_{\text{tip}} - z_{\text{sample}}$, it will die off exponentially below $\omega_0 \propto \sqrt{k_{STM}}$ [26]. Our primary goal then, in designing a low-noise STM is to maximize the resonant frequency ω_0 . In reality, there will be many resonant frequencies in a complicated structure. We want to maximize the smallest of these.




Figure 2.2: Cartoon of the coupling between the STM and its surroundings, and between the STM tip and sample. For the best noise performance, the “spring constant” k_{STM} needs to be as large as possible, while the coupling to the outside world, k , needs to be as small as possible.

Our STM is loosely modeled on an STM from the Wiesendanger lab, a group with notable success at SP-STM (the system that inspired ours is described in Ref. [25]). Considering its complexity, the scan head is fairly compact, at 1.7" diameter and 6" tall, built to fit inside the 2.5" bore of our magnet². It utilizes a home-built version of the inertial walker designed by Pan [21], in order to slowly and reproducibly approach the tip to within tunneling range of the surface (the "Z-walker"). A description of the operation of this type of walker is given in Ref. [20]. It uses six stacks of shear piezos, two of which contact each edge of a triangular sapphire prism.

A photo of the scan head is shown in Fig. 2.3. The body is constructed primarily out of Macor, chosen because it is a UHV-safe insulator (i.e. it doesn't outgas much, and it doesn't conduct electricity) with thermal expansion closely matched to that of Titanium and sapphire. The z-walker is shown at (f): the sapphire beam moves up and down carrying the tip. The sample, mounted on a sapphire rotational stage ("rotor"), is shown at (c). Other features include capacitive position sensors on both the rotor and the z-walker, an x-y coarse motion stage that translates the tip relative to the sample, and a guide (a) for the grabber tool during tip transfer (described in a later section.) The scan head has a resonant frequency of greater than 3kHz, suitable for low-noise spectroscopy.

2.2.1 Rotor

An unusual feature of this instrument is the rotor, pictured in Fig. 2.4. The sapphire rotor holds the sample and uses the Pan walker mechanism to rotate about

²In this chapter, I routinely refer to lengths in inches, following standard machining practices in this country.




Figure 2.3: STM scan head, constructed primarily of Macor (white color). The figure is about 6 inches tall. Features are labeled by the letters at right: **(a)** Titanium guide, used to guide the grabber into place for sample and tip transfer. **(b)** Copper thermal stage, the main thermal anchor with the fridge. On front is a Cernox thermometer. Not pictured are a 500mW heater and a copper braid connecting this stage to the cryostat. **(c)** Sample tray, facing downwards, and attached with beryllium copper spring clips to a sapphire “rotor.” The coax that applies a bias to the sample can be seen exiting at right. **(d)** The STM tip, barely visible. It is mounted on an assembly attached to the movable sapphire prism described in (f). **(e)** Top of the scan tube is visible (gold color). On top, a carriage assembly is epoxied in place, which holds the tip, and allows horizontal coarse motion of the tip relative to the sample. **(f)** Beryllium copper spring plate, providing a normal force (via a small ruby ball) the the Pan inertial walker, and held in place by four titanium screws. The sapphire prism moves up and down with a full range of motion of more than half an inch. It is rigidly attached to the scan tube assembly that scans the tip back and forth across the sample. **(g)** Wires for the five scan tube electrodes, loose enough to allow motion of the sapphire beam. **(h)** (inside) Position sensor with two electrodes whose area of overlap changes when the sapphire beam moves, varying the capacitance between the electrodes.




Figure 2.4: *Top:* close-up of the rotor assembly in the scan-head. *Bottom left:* Schematic of the same system, not showing the body of the STM. *Bottom right:* Photos of the bridge, and one of the two piezo mounts placed inside the STM body. Two wired piezo stacks can be seen on each piece.

an axis that goes from left to right in the photo. A titanium plate is screwed to the sapphire rotor, and the contact with the sample bias is made here. The sample plate itself, also machined from Titanium, is held in place to the rotor with Beryllium copper spring clips.

There are several motivations for including this rotational stage. While the STM is taking measurements, the rotor is positioned so that the sample faces downward. If we want to look at single-atom magnetic impurities on the surface, we can rotate the sample to face upward, and then evaporate a small quantity of these atoms from above. The advantage to evaporating while the sample is at cryogenic temperatures is that the atoms stay in one location once they contact the surface, instead of self-assembling into islands. The second function of the rotor is as follows. If we have a

magnetic tip and want to flash-magnetize the tip and sample in different orientations from one another, we can rotate the sample and briefly turn on the magnetic field. But the primary design motivation for the rotor, and the only one presently in operation, is a part of tip and sample transfer. The rotor rotates by 90° so the sample faces forward (out of the page, in Fig. 2.4). This allows the sample to be removed from above, and it also gets the sample out of the way so that the tip can be removed from above. We will discuss sample and tip transfer in more detail later.

2.2.2 Piezo Tube Assembly

The scan tube itself is potentially the lowest resonant frequency in the STM [26]. In order to scan the tip across the surface, the scan tube bends back and forth very slightly. There are four electrodes used to bend it in two directions, and fifth to stretch the tube towards the surface. The amount of bend you can get without depoling the piezoelectric scan tube determines how wide a field of view can be scanned. Applying voltages V and $-V$ to equal and opposite quadrants of the scan tube leads to a deflection

$$\Delta x = \frac{0.9d_{31}VL^2}{d_mt} \quad (2.1)$$

where L is the length of the tube, d_m the diameter, and t the wall thickness [1]. So the scan range is just twice this value, evaluated at V_{\max} . On the other hand, the resonant frequency ω_0 of the scan tube is proportional to d_mt/L^2 [?]. Increasing the scan range will automatically decrease the resonant frequency. The low resonant frequency that accompanies a large scan range means that an STM with a scan range of hundreds of microns will never be able to perform low-noise spectroscopy of the type

described in this thesis. The scan tube in our STM has length 0.9 inches, diameter 0.375", and wall thickness 0.3". And the nominal piezoelectric strain coefficient³ is $d_{31} = 0.95 \text{ \AA}/\text{V}$ leading to a scan range of about $1 \mu\text{m}$, at room temperature, if we allow a maximum of 150V. It is our experience that d_{31} shrinks by roughly a factor of 3 as you cool down from room temperature to 4K. The maximum scan range then remains unchanged between 300K and 4K, because we increase the maximum voltage applied to the tube from 150V to 450V (cryogenic vacuum much better than vacuum at room temperature, so we don't need to worry about arcing). The resonant frequency of our scan tube is about 3.3kHz, slightly lower than the simple model suggests because of the heavy carriage assembly mounted on the end.

It is common practice to have the scan tube be the same length as the sapphire prism that houses it. Our STM has an unusually long prism, because it must travel by more than half an inch to get out of the way of the rotor (in a typical Pan walker, the travel distance is less than a millimeter!). If our scan tube were as long as the prism, we would either have an unacceptably low resonant frequency, or we'd need to make the scan tube a larger diameter. The latter seems benign enough, except that the STM needs to fit inside the 2.5" magnet bore, and increasing the scan tube diameter means the entire STM diameter must increase too.

2.2.3 Set Screws

I would like to take this opportunity to advise anyone who wants to build an STM to avoid the use of set screws at all cost! They have no place in a rigid scan head.

³We use EBL#4, manufactured by EBL Products Inc.

They *will* loosen over time. Earlier models of this STM used set screws to hold the scan tube to the Sapphire prism. Normal motion of the STM caused the screws to loosen over time. If we didn't tighten them every few days the connection became wobbly. This lack of rigid connection between tip and sample, effectively a very low resonant STM frequency, made it impossible to get into stable tunneling feedback.

2.3 In Situ Sample and Tip transfer

It is common practice in many STM labs to transfer the sample in situ. Tip transfer is far less common. The task of inserting a small tip into the STM without crashing it is non-trivial. Accomplishing this feat took months of testing and troubleshooting. The sample and tip holders need to be tiny, because there is little space to spare inside the scan head, which must fit inside a 2.5" magnet bore. But the small size adds to the challenge of tip and sample transfer. They are pictured in Fig. 2.5.



Figure 2.5: From left: the sample tray, on which the sample is mounted, an STM tip in its holder, and a dime. These are the actual objects that we grab out of the STM with a 6 foot long manipulator and bring to room temperature (not the dime).

Sample and tip transfer are achieved using the same tool, shown in Fig. 2.6. To pick up the sample or tip, the grabber “swallows” the “wings” on top of the holders

(visible in Fig. 2.5), then rotates by 90° (sample) or 180° (tip). Figure 2.6 shows a schematic of the grabber tool carrying the sample. Once held by the grabber, the sample (or tip) is free to wobble slightly. The only reason they don't fall out is that there is a small shelf of height 0.01" inside the grabber tool, as seen in the top left of Fig. 2.6. This wiggle room allows the grabber to easily release the sample or tip without damaging the STM. The diagram in the lower-left is a proof of principle that the tip will not crash once it is being held by the grabber. In practice, we find that as long as the tip protrudes from its holder by less than 0.1", we can transfer it into the STM without crashing.




Figure 2.7: **(a)** Topographic image of Bi-2212, approximately 15nm wide, taken with the STM described in this chapter. **(b)** Checkerboard pattern seen in the 30mV layer of a DOS map of Bi-2212, 50nm wide. This map was taken with the same instrument but on a different sample. These images are not spin-polarized.

2.4 Atomic Resolution Tests

Atomic resolution spectroscopy has been performed successfully with this STM. A topography of Bi-2212 taken by this instrument is shown in Fig 2.7, as well as one layer from a 50nm DOS map of sample F (Bi-2212, not the same sample as the topography), showing the Q^* checkerboard pattern that I will focus on for the remainder of this dissertation. The instrument has not yet been used for spin-polarized tunneling, although the groundwork has been laid for this capability.

Chapter 3

Spatial Lockin: A New Algorithm for Mapping Local Wavevector

3.1 Background

Mapping disorder in the cuprates is nothing new. “Gapmaps” have been used for nearly a decade to illustrate the inhomogeneity of the superconducting gap and/or pseudogap. Here we characterize another facet of cuprate inhomogeneity: we shall map the local wavevector and amplitude of the checkerboard (Q^*) and the related Q^{**} order. We previously mapped Q^* showing that it varies from position to position on the surface and is correlated with the pseudogap, Δ_{PG} . In this chapter we present a new and more versatile algorithm to find the local amplitude and wavevector¹ of the CDW. The principle is simple: we find the local fourier amplitude everywhere for several different \mathbf{q} , then use these to determine the best local \mathbf{q} of the modulation.

¹and if desired, the phase

The algorithm is similar to methods that have been used in the literature to find the local amplitude of a disordered modulation. Parker [22], Lawler [15] and Mesaros [18] have all mapped local fourier components for a fixed \mathbf{q} in order to find amplitude (and phase, in the case of Mesaros) as a function of position. Wise mapped the local wavevector of the Q^* order using a Voronoi cell method [39]. The latter method has two limitations: first, it cannot access the wavevectors separately in the x- and y-directions. Secondly, if the pixel resolution is too low, (as in most DOS maps) the Voronoi cell method cannot be used to map the shorter-wavelength Q^{**} . The most obvious output of our new algorithm are “ α -maps”, which map the magnitude of the CDW wavevector.

This chapter first describes the Spatial Lockin algorithm mathematically and applies it to some idealized data. Following are some additional processing steps that improve performance on real data. After illustrating its performance on real data, we discuss several variations on the algorithm.

3.2 Mathematical Explanantion

3.2.1 Finding the local fourier component

Here we will describe the algorithm for a generic fourier component \mathbf{q} , although in this thesis the algorithm is always applied to $\mathbf{q} = Q^*$ or $\mathbf{q} = Q^{**}$.

The CDW can be viewed as an oscillation at \mathbf{q} with a local complex amplitude $\tilde{G}(\mathbf{q}, \mathbf{r})$. To find this local amplitude, we make an analogy to a lockin amplifier. We start with a conductance map $G(r)$. This is multiplied by a sine wave with wavevector

\mathbf{q} , then low-pass filtered. The local fourier component at wavevector \mathbf{q} evaluated at location \mathbf{r} is

$$\tilde{G}(\mathbf{q}, \mathbf{r}) = \int G(\mathbf{r}') e^{i\mathbf{q}\cdot\mathbf{r}'} F(\mathbf{r} - \mathbf{r}') d^2\mathbf{r}' \quad (3.1)$$

with

$$F(\mathbf{r}) = N e^{-\frac{r^2}{2L^2}}. \quad (3.2)$$

The normalization constant $N = 1/(2\pi L^2)$ chosen so $F(\mathbf{r})$ integrates to unity. A lockin amplifier such as SR830² contains a temporal low-pass filter. Here we instead use a spatial low-pass filter, which is contained in $\iint d^2\mathbf{r} F(\mathbf{r})$. We can write this more compactly as a convolution:

$$\tilde{G}(\mathbf{q}, \mathbf{r}) = (G(\mathbf{r}) e^{i\mathbf{q}\cdot\mathbf{r}}) \otimes F(\mathbf{r}) \quad (3.3)$$

It is illuminating to examine the same procedure in k-space. Starting with with the fourier transform of the data, $\tilde{G}(\mathbf{k})$, translate it in k-space by a vector \mathbf{q} , then multiply by the function $\tilde{F}(\mathbf{k})$ which is centered about the k-space origin,

$$\tilde{F}(\mathbf{k}) = N_k e^{-\frac{k^2}{2\Lambda^2}} \quad (3.4)$$

in order to low-pass filter. Since $\tilde{F}(\mathbf{k})$ all but the small spatial frequencies \mathbf{k} , the low-pass filtering is obvious in k-space. Note that $\tilde{F}(\mathbf{k})$ is just the fourier transform of $F(\mathbf{r})$. Since we are using Gaussian filters, the widths are related by

$$\Lambda = \frac{1}{(2\pi L)} \quad (3.5)$$

The result

$$\tilde{F}(\mathbf{k}) \tilde{G}(\mathbf{k} + \mathbf{q}). \quad (3.6)$$

²Made by Stanform Research Systems.

is the k-space envelope of oscillations at wavevector \mathbf{q} . Of course by the convolution theorem this equation is mathematically equivalent to the fourier transform of the local amplitude $\tilde{G}(\mathbf{q}, \mathbf{r})$ from Eq. 3.1 or 3.3

3.2.2 Choosing the Best Local Wavevector

We now have a procedure for computing the local fourier amplitude of a wavevector \mathbf{q} , $\tilde{G}(\mathbf{q}, \mathbf{r})$. To find the local wavevector, we build on this recipe. First we compute $\tilde{G}(\mathbf{q}, \mathbf{r})$ for a set of closely-spaced \mathbf{q} . Unless noted, in this dissertation I always take \mathbf{q} in the direction of the atomic Bragg peak Q_x (or Q_y), so that

$$\mathbf{q} = \alpha Q_x. \quad (3.7)$$

I often refer to α as “the wavevector”, although it is more accurately the *magnitude* of the CDW wavevector in units of $2\pi/a_0$. For example, $\alpha = 1$ refers to the k-space radius at the bragg peaks Q_x and Q_y , and $\alpha = 0$ is the k-space origin. For the oscillations studied in this dissertation, the Q^* peak has α between 0.08 and 0.30, while the Q^{**} peak has α between 0.70 and 0.92. When we find the best local wavevector for Q^* (or Q^{**}), we will call its magnitude α^* (or α^{**}).

Once we have $\tilde{G}(\alpha, \mathbf{r})$ for several different α , we go to each pixel \mathbf{r} and pick the value of $\alpha(\mathbf{r})$ that maximizes \tilde{G} at that pixel. A schematic of this process is shown in Fig. 3.1. Alternatively, we can pick $\alpha(\mathbf{r})$ using a weighted average. The two methods give similar results, and sometimes one works better than the other.

For the Q^* (or Q^{**}) peak, there are three different potential wavevector maps that we can make.

- **x-direction.** To map α_x^* , compute the amplitude $\tilde{G}(\alpha_x, \mathbf{r})$ only for the Q^* peak in one crystal direction, where $Q_x^* = \alpha_x^* Q_x$.
- **y-direction.** To map α_y^* , compute the amplitude $\tilde{G}(\alpha_y^*, \mathbf{r})$ only for the other crystal direction, $Q_y^* = \alpha_y^* Q_y$.
- **Combined x- and y-directions.** Compute the amplitude separately in each of the two crystal directions. Then at each point \mathbf{r} add them in quadrature:

$$\tilde{G}(\alpha, \mathbf{r}) = \sqrt{\tilde{G}_x(\alpha, \mathbf{r})^2 + \tilde{G}_y(\alpha, \mathbf{r})^2},$$

and find the α that maximizes this quantity. When we refer to α with no subscript, we shall assume that we are combining the directions.

Figure 3.2 demonstrates the spatial lockin algorithm on a set of ideal sine waves. Four perfect plane waves of different frequency, amplitude, and direction are shown in (a). The lower left quadrant contains the sum waves in two different directions. Figures 3.2 (b)-(d) show different outputs of the algorithm, which accurately picks the correct wavevector and amplitude. In Fig. 3.2(b) we see that, even for a perfect sine waves, the amplitude vs. α curves have a finite width around the correct value due to the finite Λ .

3.2.3 Choosing L

In choosing the gaussian half-width L (or in k-space, Λ) we face a tradeoff between local information (small L) and accurate information (large L , or small Λ). In practice, the algorithm does a good job down to $L \approx 0.6\lambda$, where λ is the average wavelength of the CDW. For the three-quarter peak, however, we seem to be limited




Figure 3.1: Schematic of the process to create an α -map. First the spatial lockin algorithm is run on the raw data $G(r)$. The result is a three-dimensional set of values $\tilde{G}(\alpha, \mathbf{r})$. Each layer gives the fourier amplitude of a different wavevector α . At each pixel (e.g. the pixel at $\mathbf{r} = r_0$) there is a curve of amplitude vs α . The value of α where the curve peaks is recorded in the alpha map.




Figure 3.2: Dummy data to illustrate the spatial lockin algorithm. (a) Sine waves of different frequency, amplitude and Q_x and Q_y direction. (b) Amplitude vs α for a test-pixel in the lower left quadrant, marked with an “x.” (c) Local alpha map corresponding to (a), generated with $L = 2a_0$. (d) Local amplitude map corresponding to (a).

by pixel resolution too. I find it best not to go below $L \approx 4$ pixels. One way to find the best L is to maximize the correlation with the gapmap, which is justified by the findings of Wise *et al* [39]. Fig. 3.3 plots the correlation coefficient R between the local wavevector map and the pseudogap (Δ_{PG}), for different values of L . These two quantities, Q^* and Δ_{PG} are known to correlate [39], so to choose L we maximize the correlation coefficient.




Figure 3.3: Correlation coefficient of α^* with the pseudogap, plotted as a function of spatial lockin gaussian halfwidth L , in units of a_0 .

3.3 Additional Steps

Now that we have demonstrated the main capabilities of the spatial lockin, we detail a few additional steps that help it run more smoothly on real data.

3.3.1 Preprocessing

Before running the spatial lockin algorithm, we do the following:

- **Drift correct.** Remove distortion due to thermal drift and piezo hysteresis using Lawler’s algorithm [15].
- **Stoll correct.** Remove artifacts due to a directional tip. This algorithm is described in Chapter 5 of this dissertation.
- **Fourier filter.** Remove signals not related to the Q^* and Q^{**} peaks.

The last point deserves some explanation. In order to believe the output of the spatial lockin, we need to make sure that all of our signal comes from the Q^* or Q^{**} peak. Figure 3.4 shows the k-space range swept out by filters in our spatial lockin method, illustrated for the Q^{**} peak in the Q_x direction. The radius of the red circle is Λ . Other bright spots in the Fourier transform—Bragg peaks Q_x or Q_y , any of the other CDW cloverleaf peaks, or supermodulation peaks—are often much brighter than the CDW peaks we care about. If any of these spots appear in the red region, or nearby (in the tails of the gaussian $\tilde{F}(\mathbf{k})$), they will dominate the spatial lockin signal. To be safe, we remove (by fourier filtering) everything outside the square formed by the Bragg peaks, including the Bragg peaks themselves. This filtering is the reason for the square cutoff in Fig. 3.4. Removing the Bragg peak is like a 60Hz notch filter in a lockin amplifier. In dataset F we have also removed spots due to supermodulation.

The problem is harder for the Q^* peak, which may overlap substantially with the central peak of the Fourier transform. First and foremost, we need to minimize




Figure 3.4: **(a)** Fourier transform of 80 mV R-map. The k-space range swept out by filters in our spatial lockin method is represented by area enclosed in the red circles. We want any signal appearing in this range to be due to Q^{**} and not to other physics. Each circle has radius Λ , the gaussian halfwidth of our fourier-space filter, corresponding to a real-space gaussian filter of halfwidth $L = 2a_0$. In this example, the Bragg peak Q_x and the four nearby CDW peaks will skew the output of the spatial lockin. **(b)** A fourier filtering solution to keep these peaks from contaminating the Q^{**} signal.

this overlap (and therefore maximize out signal-to-noise ratio) by picking a clever combination of dI/dV layers. For our OPT and UD Bi-2201 maps, we have found that energy-integrating a DOS map from -20mV to +20mV, where the Q^* signal is strongest, separates Q^* from the central peak sufficiently. We then divide by the -100mV layer, to remove the setup condition. In overdoped samples, separating Q^* from the central peak is not always possible. In these cases, it may be useful to try the angular fourier-transform variation described in section 3.5.

3.3.2 Postprocessing

After running the spatial lockin algorithm, a number of steps can improve the resulting α -maps. Even before these steps, correlation with the pseudogap is already well-established, with R approaching 0.50 in sample A (UD32K). Taken together, the following steps can increase R by a striking 0.05.

- Remove the 5% of pixels with the smallest amplitude and replace with the value of α at the nearest pixel. There are vortices in Q^{**} where α will not be accurate [18] but we can get a good approximation from the nearest neighbors.
- Smooth. A simple boxcar filter with a 5 pixel halfwidth (for Q^{**}) or 4-pixel halfwidth (for Q^*) gives good results.

3.4 Applying the Spatial Lockin

This section demonstrates the algorithm on real data. Figure 3.5 applies the spatial lockin to the three quarter peak of an 80 mV R-map of sample A, UD32K,

shown in (a). Curves of amplitude vs alpha taken at three test pixels are shown next in (c)-(e). In well-behaved datasets, e.g. the six datasets presented in this thesis (see Table 1.1), more than 97% of pixels show a Q^{**} peak similar to those illustrated. The Q^* peak is only well-behaved in datasets A and B. To pick the best local wavevector in each direction, we can either use the value of α at the peak, or take a weighted average. Empirically, the weighted average leads α to correlate much better with the pseudogap than the peak value. Figures 3.5f and g shown the resulting wavevector and amplitude maps. For this example, we have used the length $L=2a_0$.

3.5 Variations on the Spatial Lockin

- **Angular Fourier transform.** This variation provides a form of background subtraction, in case Q^* overlaps substantially with the central peak of the fourier transform. It is more computationally intensive than the basic spatial lockin. For each value of α , we evaluate the local fourier transform not just at the two points along Q_x and Q_y , but at 32 points around the circle, as in Fig. 3.6. This essentially gives us a fourier-transform in polar coordinates $\tilde{G}(\alpha, \theta)$, which we only evaluate at certain k-space radii α . The next step is to fourier transform *with respect to* θ , setting $\theta = 0$ along the Q_x direction. The Q^* amplitude at each α is then the real part of the frequency=4 component, that is the four-fold symmetric component, of this fourier-transform.

This AFT method is not capable of finding α separately in the x- and y- directions. But by taking the two-fold component of the transform, we can get a measure of the anisotropy of the modulation. If positive, the modulation




Figure 3.5: (a) Real space data on which we run the spatial lockin. An 80 mV R-map 68 nm wide, of sample A (UD32K Bi-2201). (b) Fourier transform of the image in (a). (c)-(e) Amplitude vs. α for three randomly-chosen pixels from (a). There are three curves: the Q_x -direction, Q_y -direction, and the overall CDW amplitude obtained by summing the Q_x and Q_y amplitudes in quadrature. (f) Map of α^{**} made using the spatial lockin method. (g) Map of the amplitude $|\tilde{G}(\mathbf{q}, \mathbf{r})|$ made using the spatial lockin method.




Figure 3.6: Schematic of the k-space coordinates used to evaluate an angular fourier transform.

is stronger in the x-direction; if negative, it is stronger in the y-direction.

Simpler algorithms have been used to subtract the amplitude of the central blob in a global FFT (see, e.g. Ref [22]). We tried finding the central blob from both a global and local FFT, then subtracting this amplitude from our local FFT amplitudes. In our case, neither of these simpler methods produced meaningful results.

Chapter 4

The Disordered Local Wavevector of the Q^* and Q^{**} Peaks

In the simplest picture, the Q^* and Q^{**} modulations are caused by the same physics, namely antinodal nesting of the Fermi surface. A cartoon of this nesting is shown in Fig. 4.1. In this picture, Q^* and Q^{**} are reflections of each other across the Brillouin zone boundary. If this cartoon is correct, then we should expect that as we increase the size of the hole pockets by doping, α^* will decrease and α^{**} increase. STM measurements have observed exactly this behavior by comparing samples of different doping for both the one quarter ([38], [39]) and three quarter ([14], [18]) peaks. The wavevector Q^* was shown to vary locally and correlate with the local value of the pseudogap energy scale Δ_{PG} [39], which can be viewed as a measure of local doping. We will show for the first time in this chapter that the Q^{**} also anticorrelates with Δ_{PG} .

The simple explanation shown in Fig. 4.1 is not universally accepted, even among




Figure 4.1: Cartoon of the BSCCO Fermi surface showing the antinodal nesting vectors. As hole doping increases, the blue hole pockets get larger; Q^* gets shorter while Q^{**} gets longer.

STM studies. The main reason is that the peaks have very different energy behavior. Q^* seems to have most of its amplitude near zero bias, while the Q^{**} has much more weight at high, positive bias, as observed by both Parker and Kohsaka (albeit with different normalizations) [22], [14]. While analyzing C_4 -symmetry-breaking, Lawler likened the Q^{**} to a liquid crystal smectic order. However, he chose to ignore Q^* , implying that it is perhaps caused by a different physical picture. While the strange energy dependence should definitely give us pause, we will assume for the purposes of this chapter that the two peaks are caused by the same physics. Indeed our findings later in this chapter on the doping and pseudogap dependence of these Fourier peaks will support this position.

Disorder complicates the picture. In the absence of disorder, a commensurate charge density wave (CDW) will have a spatially uniform wavevector. Not so for our disordered incommensurate CDWs in which the wavevectors Q^* and Q^{**} take on a wide range of values locally. This disorder is seen readily in the cuprate CDW.

Figure 4.2a shows a Fourier transform of the 10mV layer of a DOS map of sample A, UD32K. The circle shows the location of the Q^* peak and its radius gives the half width σ of the gaussian fit to the peak. The circle has radius $\sigma = 0.03(2\pi/a_0)$ in fourier-space, corresponding to a real-space disorder length scale (half width) of $5a_0$.¹ Figure 4.2b visualizes this disorder in real-space with a map Fourier filtered around the Q^* peak. Local amplitude variations were removed to clearly illustrate the peaks and troughs. The ordering length that we measure from the Fourier transform, on par with the wavelength itself, is consistent with our understanding that the modulations are fluctuating order made visible to our static probe by pinning due to quenched disorder (i.e. dopant disorder) [35], [12].

Charge ordered states in the cuprates may suffer from two different causes of disorder, that can be understood by analogy to water waves. In a shallow (flat-bottomed) pool, the wavelength of surface waves is set by the depth of the water. By contrast, in a shallow, rocky river the crests and troughs of waves are pinned by the locations of rocks and other subsurface features. In this extreme the local wavelength has very little to do with the depth of the water. We can imagine that our disordered charge density wave in the cuprates lies at one extreme or the other, or somewhere in between. Does the CDW respond to the local depth of the Fermi sea, or are its crests and troughs pinned by random point defects?

Here we come to the main motivation of this chapter: to characterize the type of disorder affecting Q^* and Q^{**} . Specifically, we will determine whether they are

¹We are confident that the width of the Fourier peak is due only to disorder and not to thermal drift or piezo hysteresis. The drift-correction routine straightens the atomic lattice so that the atomic Bragg peaks in the Fourier transform are only a single pixel wide. This routine does not appreciably affect the width of the Q^* or Q^{**} peaks.




Figure 4.2: **(a)** Fourier transform of the 10mV layer of sample A (UD32K Bi-2201). The circle shows the location of Q_x^* , while the arrow points to one of four atomic Bragg peaks. The radius of the circle gives the half width σ of a gaussian fit to the peak. The circle has radius $\sigma = 0.03(2\pi/a_0)$ in fourier-space. **(b)** Fourier filtered map of Q^* charge modulation (both crystal directions), with local amplitude variation removed. The small blue arrow is the real space disorder length scale (gaussian half width) $\sim 5a_0$ corresponding to the width of the fourier peak in (a).

affected most by local doping disorder, or by pinning. To accomplish this, we compare the two peaks to each other. Local doping disorder causes Q^* and Q^{**} to be related in a local way; pinning potentially disrupts this local relationship. If the local doping is the dominant form of disorder, then Q^* and Q^{**} will be simple reflections of one another. In this case, we expect their wavevectors to sum to the width of the Brillouin zone: $\alpha^* + \alpha^{**} = 1$. A glance at Ref. [14] suggests that this may be true globally, though the sum has not been explicitly, quantitatively checked. We ask whether the same relation is true of the local wavevectors at every point on the surface. Do the two wavevectors locally sum to one? If so, then we are in the flat-bottomed pool. If not, a shallow river.

The structure of this chapter is as follows. We first examine the global behavior of the Q^* and Q^{**} peaks to check that it matches reports from the literature. Next we use a new algorithm to map the local wavevectors of both Q^* and Q^{**} . We analyze the results in two different ways to see whether they sum to the Brillouin zone width locally as they do globally.

4.1 Materials

To answer these questions, we rely on two high-quality DOS maps of Pb-doped Bi-2201, samples A and B from Table 1.1. Due to the lead doping there is no supermodulation, which makes it easy to see the checkerboard by eye. The range of T_c is much smaller for these samples than for Bi-2212; T_c at optimal doping is only 35K. We focus primarily on the underdoped and optimally doped datasets for a simple reason: at higher doping, Q^* overlaps hugely with the long wavelength inhomogeneity (the central peak of the fourier transform.) However, we must be careful when comparing our results to those in the literature; all previous STM observations of Q^* and Q^{**} order were carried out on Bi-2212. The exceptions are Refs. [38] and [39] which use, among others, the same samples A and B used in this thesis.

We seldom see both the Q^* and Q^{**} peaks in the same energy layer; this makes it necessary to view them in different layers. Under the assumption that the order is characterized by one wavevector at each location on the sample, independent of energy, we combine the energy layers of the DOS map to increase the signal-to-noise ratio of the oscillations. Following Kohsaka, we image Q^{**} order in a high bias layer of map of $R = I(V)/I(-V)$ [14]. Since we're using $V = 80\text{mV}$ that is larger than




Figure 4.3: Approximate phase diagram of Bi-2201, showing the location of the three samples.

nearly all gap sizes present, we expect the images to be the same as those from Kohsaka's maps that have been scaled by gap. This map increases signal-to-noise for Q^{**} substantially over any single energy layer. No other normalization that we tried performed better (e.g. we tried $Z = \text{DOS}(V)/\text{DOS}(-V)$, unnormalized dI/dV layers at different energies). On the other hand, Q^* does not show up very well in an R-map. Its amplitude is highest near zero bias and the modulations are in phase across the Fermi level, so an R-map tends to extinguish the Q^* modulations. Instead, following the suggestion of Howald [10] we maximize signal-to-noise by integrating over a small but finite energy range near $V = 0$, where the oscillation is strongest (we

Table 4.1: **Samples analyzed in this chapter**

sample	A	B	D
T_c	UD32K	OPT35K	OD16K
approx hole concentration p	0.128	0.16	0.241
size	68nm	72nm	36nm
setup bias V_{set}	-200mV	-100mV	100mV
instrument	Hudson1	Hudson1	Hoffman1
global α^*	0.183	0.156	0.12
global α^{**}	0.792	0.830	0.850

use -20 to 20mV.) This will tend to cancel out dispersive, phase incoherent features such as QPI. To remove matrix element effects from the setup condition, we divide this integrated dI/dV by the conductance in the -100mV layer. We call the result an “integrated $Z_{(-)}$ map.” Figure 4.4 shows these maps along with their Fourier transforms for both sample A (UD32K), and sample B (OPT35K). As advertised, Q^* is prominent in the integrated $Z_{(-)}$ map (4.4a and 4.4c) while Q^{**} is much more visible in the R-map (4.4b and 4.4d).

The doping dependence of the three quarter and one quarter peaks behaves as we’d expect in the antinodal nesting picture. Figure 4.5 shows linecuts through the FFTs of the two maps. α^* decreases with doping while α^{**} decreases. The α values of the peaks were found by fitting the linecut peak to a gaussian in one dimension, illustrated in Fig. 4.5d. To remove the influence of long wavelength inhomogeneity, an angular Fourier transform method of background subtraction was employed to find α^* (see chapter 3). The wavevector magnitudes α for these three data sets can be found in Table 4.1. The wavevectors are summed and plotted in Fig. 4.5(c). The wavevector magnitudes α^* and α^{**} very nearly sum to one for all three dopings, although all three points have a sum that is slightly smaller than expected.




Figure 4.4: **(a)** Sample A (UD32K): Integrated $Z_{(-)}$ -map emphasizing Q^* , along with its Fourier transform. The real-space size is 68nm and 400 pixels wide. **(b)** 80mV R-map (also from sample A) highlighting Q^{**} , and Fourier transform. **(c)** and **(d)** are the same maps for sample B (OPT35K). The sample B field of view is 72nm and 468 pixels wide.




Figure 4.5: Linecuts of FFTs in Q_x direction for samples of three different dopings. **(a)** Integrated $Z_{(-)}$ map, with prominent Q^* peak. **(b)** R-map, with prominent Q^{**} peak **(c)** The global wavevectors sum close to one. **(d)** Fit of the Q^{**} peak for sample D (OD16K). The very small horizontal bar below the peak (under the axis) represents twice the standard error of the peak fit, $\delta\alpha = \pm 0.007$. This fit error, along with the error from the Q^* peak determines the error bars in (c).




Figure 4.6: **(a)** Map of the pseudogap in sample A (UD32K). **(b)** Average spectrum for gaps of different sizes. Each bin represents 10% of the pixels in the map.

As a first look at local alpha behavior, we analyze the data in a similar method to Wise [39]. We aim see whether Q^{**} anticorrelates with the pseudogap in the same way that Q^* is known to correlate [39]. First we create a map of the local pseudogap, shown for sample A (UD32K) in Fig.4.6. The data is divided into six overlapping bins, each containing about 40% of the pixels. Next we Fourier transform the points in each bin, masking out the rest. Finally, we fit the Fourier peaks to Gaussians as we did in Fig. 4.5d in order to find α . The resulting linecuts are shown in Fig. 4.7. Once again we rely on the angular fourier transform method of background subtraction from Chapter 3; this correction is reflected in the α^* values of Fig. 4.7c, though not in the linecuts of 4.7a. The values of α^* and α^{**} in each bin sum to nearly, but slightly less than 1. We believe this offset, as well as the small trend in the sum (Fig. 4.7e), is a systematic error due to an imperfect subtraction of the long-wavelength inhomogeneity background. Analyzing the wavevector in a global k-space sense seems to support our simple cartoon understanding. It remains to be seen whether $\alpha^* + \alpha^{**} = 1$ in a truly local sense.




Figure 4.7: Sample A (UD32K) Linecuts of FFTs in Q_x direction for six different overlapping gap bins. (a) Integrated $Z_{(-)}$ map, showing Q^* peak. (b) R-map, showing Q^{**} peak (c) The wavevector α^* plotted against gap value. (d) The wavevector α^{**} . (e) The global wavevectors sum to close to one. Horizontal error bars are the standard deviation of gaps in each bin, while vertical error bars reflect the accuracy of peak fits.

4.2 Local Wavevector

To see how Q^* and Q^{**} vary locally, we employ a new algorithm known as the “spatial lockin.” It gets its name from the similarity to a lockin amplifier: the data is multiplied by a sine wave and then low-pass filtered separately at each real-space pixel. The algorithm is essentially a series of local Fourier transforms, and was described in detail in the previous chapter. The main output is a map of the wavevector magnitude α . Figures 4.8 and 4.9 show α -maps for the fields of view in Fig. 4.4. The wavevector of the Q^{**} peak appears to change over a shorter length scale than that of Q^* , however the correlation length of the map is not meaningful, as it is set by the internal filtering parameter L from the algorithm (this is the “space constant” akin to “time constant” in a lockin amplifier.) The α^* maps were made using $L = 3a_0$ (sample A), $L = 3.5a_0$ (sample B), while the α^{**} maps were made using $L = 2a_0$ for both samples. Also shown are maps of the local pseudogap. The correlation between maps is visible by eye. The color scale in the map of α^{**} has been reversed to better illustrate its anticorrelation with the other two maps.

4.2.1 Correlation Coefficients

In the previous section we saw that the CDW wavevector correlates with the pseudogap in a global way. But how well do the maps correlate locally? One measure of local correlation is the correlation coefficient R between an α map and a gap map. A correlation of $R = 1$ (or $R = -1$) implies that two maps are perfectly correlated (or anticorrelated) and a scatterplot of their points lies along a straight line. A correlation of $R = 0$ means there is no discernible monotonic relationship between




Figure 4.8: α -maps corresponding to the maps of sample A (UD32K) in Fig. 4.4a and b, with units of $2\pi/a_0$. (a) Q^* peak (b) Q^{**} peak (c) Pseudogap. The color scales have been adjusted to highlight the correlation (and anticorrelation) between maps. Red regions tend to line up with other red regions, and blue with blue.

the two maps. For sample A (UD32K), our maps show a strong local relationship with the pseudogap, with $R = 0.55$ for Q^* and $R = -0.55$ for Q^{**} . For the peak in sample B (OPT35K), the values are $R = 0.40$ and $R = -0.35$. It is possible but not certain that the difference in R between the two maps is caused by a less accurate gap map for sample B; by contrast, the spectra in sample B showed a clear peak at positive energy making it easy to determine the gap. Tables 4.2 and 4.3 show the correlation coefficients between all possible alpha maps for these two samples. The directionally segregated maps α_x and α_y are not pictured. Recall from Chapter 3




Figure 4.9: α -maps corresponding to the maps of sample B in Fig. 4.4c and d, with units of $2\pi/a_0$. Combined X and Y directions. (a) Q^* peak (b) Q^{**} peak (c) Pseudogap. The color scales have been adjusted to highlight the correlation (and anticorrelation) between maps. Red regions tend to line up with other red regions, and blue with blue.

that α_x and α_y are computed using the amplitude from charge modulations in only one crystal direction, while the α -map that combines the directions is computed by adding in quadrature the amplitudes from the x- and y-directions.

We can see a few immediate results from Tables 4.2 and 4.3. First, α^* anticorrelates with α^{**} ; this should not be surprising given their separate relationships to the pseudogap. Secondly, the maps α_x (α_y), which include only information from one lattice direction Q_x (Q_y), show less anticorrelation between the Q^* and Q^{**} peaks

Table 4.2: Correlation coefficients between α -maps, sample A (UD32K)

	X-direction	Y-direction	Combined X and Y directions
$R(\alpha^*, \Delta_{\text{PG}})$	0.418	0.418	0.548
$R(\alpha^{**}, \Delta_{\text{PG}})$	-0.387	-0.384	-0.558
$R(\alpha^*, \alpha^{**})$	-0.359	-0.363	-0.417
$R(\alpha_x^*, \alpha_y^*)$		0.258	
$R(\alpha_x^{**}, \alpha_y^{**})$		0.128	
$R(\alpha_{\Sigma x}, \alpha_{\Sigma y})$		0.239	

Table 4.3: Correlation coefficients between α -maps, sample B (OPT35K)

	X-direction	Y-direction	Combined X and Y directions
$R(\alpha^*, \Delta_{\text{PG}})$	0.277	0.291	0.399
$R(\alpha^{**}, \Delta_{\text{PG}})$	-0.206	-0.293	-0.351
$R(\alpha^*, \alpha^{**})$	-0.242	-0.339	-0.345
$R(\alpha_x^*, \alpha_y^*)$		0.145	
$R(\alpha_x^{**}, \alpha_y^{**})$		0.081	
$R(\alpha_{\Sigma x}, \alpha_{\Sigma y})$		0.106	

than the α -maps that combine the two directions. That is

$$|R(\alpha_x^*, \alpha_x^{**})| < |R(\alpha_{x+y}^*, \alpha_{x+y}^{**})| \quad (4.1)$$

This is a surprising result. It seems to run counter to our hypothesis that the peaks sum to one locally, since we expect the sum to be independently satisfied in each lattice direction.

The local wavevectors Q^* and Q^{**} are not the only pair that we expect to be linked by local doping. Q_x^* and Q_y^* should also be linked by local doping, as should the pair Q_x^{**} and Q_y^{**} . After all, the local doping model says that the wavelength is determined by the local depth of the Fermi sea. But looking at Tables 4.2 and 4.3,

we see that these pairs have *smaller* correlation coefficients:

$$|R(\alpha_x, \alpha_y)| < |R(\alpha^*, \alpha^{**})| \quad (4.2)$$

Therefore the local doping can't be the only factor at play. It appears that pinning disrupts the correlation between wavevectors of different orientations more than between correlations of different wavelengths.

4.2.2 Standard Deviations

Now we sum the two maps. If we add together α^* and α^{**} for sample A (UD32K), the resulting map has a histogram centered at 0.9879 with a standard deviation of 0.0105. It is possible that the slight deviation from 1 is caused by an underestimate of α^* due to the long wavelength inhomogeneity (the central peak in the FFT). These maps were made without any form of background subtraction², so we expect α^* to be an underestimate, and to have a similar systematic error to the one in Fig. 4.7e. The fact that the two maps add up to 1 should not be surprising. After all, the average values $\langle \alpha^* \rangle$ and $\langle \alpha^{**} \rangle$ already sum to one. If we rotate one image by 90° before summing, in order to remove any correlation between the images, the sum still has the same average value as the original sum. However the real story happens in the standard deviations of the α -maps, σ^* and σ^{**} . Tables 4.4 and 4.5 show the mean and standard deviation of each alpha map, along with the standard deviation of the sums. Also shown are the values for the rotated sum, and the theoretical values for

²A candidate technique is the ‘‘Angular FFT’’ variation of the spatial lockin. We tried this but opted not to use the technique for local wavevector maps of underdoped samples because the resulting maps correlate far less with Δ_{PG} than the maps in this chapter made by the standard spatial lockin routine. A *global* version of the AFT was, however, used Fig. 4.7c.

uncorrelated maps, defined as

$$\langle \alpha_{\Sigma_{\text{random}}} \rangle = \langle \alpha^* \rangle + \langle \alpha^{**} \rangle$$

and

$$\sigma_{\text{random}} = \sqrt{\sigma^{*2} + \sigma^{**2}}$$

The theoretical random values agree very well with the rotated sums; the last two rows in each table are nearly identical. While the standard deviations 0.0105 (actual sum, sample A) and 0.0127 (rotated sum, sample A) may not seem very different, they show an unambiguous difference between correlation and randomness. This difference in histogram width is visualized in Fig. 4.10, which shows maps of $(\alpha^* + \alpha^{**} - \langle \alpha_{\Sigma} \rangle)^2$ (a) and the same value after rotating one map (b). The color scales are the same in (a) and (b). However, the elephant in the room is that the standard deviation of the sum is always still greater than either σ^* or σ^{**} . Summing the wavevectors does not get rid of the variation in wavevector as we'd expect if local doping were the whole story.

4.3 Conclusion

In summary, we have examined the behavior of the charge order wavevectors Q^* and Q^{**} both globally and locally in Bi-2201 of optimal and light underdoping. Globally, that is, looking at the fourier peaks from large subsets of the image, our finding echo those of previous studies. Both Q^* and Q^{**} display the qualitative doping dependence and pseudogap dependence that we'd expect from antinodal nesting. The sum of the global wavevectors is 1-2% less than the width of the Brillouin zone. While

Table 4.4: Mean $\langle\alpha\rangle$ and standard deviation σ of wavevector maps, sample A (UD32K)

	X-direction	Y-direction	Combined X and Y directions
$\langle\alpha^*\rangle \pm \sigma^*$	0.1808 ± 0.0133	0.1824 ± 0.0139	0.1820 ± 0.0110
$\langle\alpha^{**}\rangle \pm \sigma^{**}$	0.8062 ± 0.0087	0.8057 ± 0.0087	0.8059 ± 0.0064
$\langle\alpha_\Sigma\rangle \pm \sigma_\Sigma$	0.9870 ± 0.0134	0.9881 ± 0.0139	0.9879 ± 0.0105
$\langle\alpha_{\Sigma\text{random}}\rangle \pm \sigma_{\Sigma\text{random}}$	0.9870 ± 0.0159	0.9881 ± 0.0164	0.9879 ± 0.0127
$\langle\alpha_{\Sigma\text{rot}}\rangle \pm \sigma_{\Sigma\text{rot}}$	0.9870 ± 0.0160	0.9881 ± 0.0166	0.9879 ± 0.0130

Table 4.5: Mean $\langle\alpha\rangle$ and standard deviation σ of wavevector maps, sample B (OPT35K)

	X-direction	Y-direction	Combined X and Y directions
$\langle\alpha^*\rangle \pm \sigma^*$	0.1704 ± 0.0120	0.1697 ± 0.0120	0.1703 ± 0.0091
$\langle\alpha^{**}\rangle \pm \sigma^{**}$	0.8170 ± 0.0090	0.8168 ± 0.0088	0.8170 ± 0.0066
$\langle\alpha_\Sigma\rangle \pm \sigma_\Sigma$	0.9874 ± 0.0133	0.9865 ± 0.0123	0.9873 ± 0.0094
$\langle\alpha_{\Sigma\text{random}}\rangle \pm \sigma_{\Sigma\text{random}}$	0.9874 ± 0.0150	0.9865 ± 0.0149	0.9873 ± 0.0112
$\langle\alpha_{\Sigma\text{rot}}\rangle \pm \sigma_{\Sigma\text{rot}}$	0.9874 ± 0.0152	0.9865 ± 0.0147	0.9873 ± 0.0112

it is tempting to credit this small difference to a breakdown of our simple nesting model, we can't yet rule out a systematic error due to long wavelength inhomogeneity. Our global analysis establishes that local doping is an important effect for both Q^* and Q^{**} .

Local wavevector maps help us characterize the disorder underlying our charge modulations. We find we are somewhere between the two disorder extremes of local doping and strong pinning. On one hand both the Q^* and Q^{**} peaks correlate or anticorrelate well with Δ_{PG} , our proxy for local doping. On the other hand, they correlate less well with each other. When added together, the histogram of the sum is wider than either of the two α -map histograms, whether in the x-direction, the




Figure 4.10: **(a)** Sum $(\alpha^* + \alpha^{**} - \langle \alpha_\Sigma \rangle)^2$ for sample A (UD32K). **(b)** The same sum, performed after first rotating the α^{**} -map by 90° in order to simulate a random map with the same distribution.

y-direction, or combining information from the two directions. This is not what we'd expect in the local doping limit, so pinning must be important.

A further interesting facet to this pinning is that the CDW wavevectors in different lattice directions are less related than those in the same direction. Naively, these should all be related to each other equally through their common relationship with local doping. Thus our pinning disorder disrupts local doping coupling between different wavevector orientations more than between different wavevectors of the same orientation. This will come up again in Chapter 6, where we will see evidence that Q^* and Q^{**} are caused by the same physics.

Chapter 5

Correcting for Anisotropy of the STM Tip

5.1 Motivation

To understand the pseudogap state of BSCCO, there is a growing interest in electronic order that breaks underlying symmetries of the crystal. In particular, several recent STM studies have focused on broken local C_4 symmetry at different wavelengths [15], [18]. These papers rely on quantitative comparison of the amplitude of electronic order between different crystal directions, a comparison that is highly sensitive to STM tip geometry. For instance, to prove that long-range nematic order is present in underdoped Bi-2212, Lawler first needed to argue that the effect was not due to tip geometry[15]¹. Therefore, when studying broken C_4 symmetry, it is ideal

¹They did this by comparing the height of Bragg peaks in the topo. I feel that they adequately proved this in two of the four maps that appeared in the paper.

to have a tip with perfect rotational symmetry. However few tips are perfect, and even a slight anisotropy can create the appearance of broken crystal symmetry².

We present in this chapter an algorithm allowing us to study C_4 -symmetry-breaking for wavevectors along the tetragonal crystal axes (the copper-oxygen bond direction) with an STM tip that is not perfectly isotropic.

5.2 Quantifying Tip Anisotropy

A convenient way to gauge the anisotropy of a tip is to take the Fourier transform of a topography and compare the weight under the Bragg peaks in the two different crystal directions. This is especially useful because the electronic order this thesis focuses on has wavevectors in the same crystal directions as Q_x and Q_y . In BSCCO, the crystal has approximate tetragonal symmetry, so the Q_x and Q_y Bragg peaks in a topography should have equal weight. However even with a perfect STM tip, broken C_4 electronic symmetry could cause these peaks to have unequal weight. In other words long-range nematic order may cause a difference in amplitude between the two topographic directions. We show in this section that this contribution is small and that nearly all amplitude at the atomic wavelength in the topography is from structural contributions. We can therefore safely assume that any difference in amplitude of atomic corrugation between two lattice directions is due to an asymmetric tip.

First, some shorthand. I'll abbreviate the integrated density of states as

$$\text{IDOS}(V_{\text{set}}) \equiv \int_{-eV_{\text{set}}}^0 \rho(\epsilon) d\epsilon \quad (5.1)$$

²It is even possible for an anisotropic tip to make long range order appear less long-range, or disappear altogether.

and define a metric for topographic anisotropy, using the Fourier transform (FT) of the topo:

$$\beta \equiv \frac{\text{Weight under } Q_x \text{ Bragg peak of FT\{topo\}}}{\text{Weight under } Q_y \text{ Bragg peak of FT\{topo\}}}. \quad (5.2)$$

As we've defined it here, β is measured from an STM topography. Explicit instructions for calculating β will be given in Section 5.4. In this section we will show that β is also a good metric for the anisotropy of the STM tip itself.

Suppose the IDOS is perfectly flat except for a modulation at wavevector \mathbf{q} :

$$\text{IDOS}(V_{\text{set}}) = C + A \cos(\mathbf{q} \cdot \mathbf{r}). \quad (5.3)$$

C is the spatially-averaged IDOS. The exact energy dependence of this modulation doesn't matter for our present purposes. It could be centered all at one energy, or the modulation could add coherently across many energies between 0 and V_{set} . The IDOS modulation in Eq. 5.3 will appear in the topography as well as the DOS map. We want to find its amplitude in the topography.

Starting with the feedback equation for a topo,

$$I_{\text{set}} = M e^{-\kappa z} \text{IDOS}(V_{\text{set}}) \quad (5.4)$$

we can solve for z to get

$$z = \frac{1}{\kappa} \ln \left(\frac{M}{I_{\text{set}}} \text{IDOS}(V_{\text{set}}) \right) \quad (5.5)$$

$$= \frac{1}{\kappa} \ln \left(\frac{M}{I_{\text{set}}} \right) + \frac{1}{\kappa} \ln (C + A \cos(\mathbf{q} \cdot \mathbf{r})) \quad (5.6)$$

$$= \text{const} + \frac{1}{\kappa} \ln \left(1 + \frac{A}{C} \cos(\mathbf{q} \cdot \mathbf{r}) \right) \quad (5.7)$$

$$\approx \text{const} + \frac{1}{\kappa} \frac{A}{C} \cos(\mathbf{q} \cdot \mathbf{r}) \quad (5.8)$$

$$(5.9)$$

where the last equality holds if $\frac{A}{C} \ll 1$. We have ignored a weak spatial dependence in the decay constant κ . The amplitude of this modulation in the topo is therefore roughly

$$\text{amplitude} \approx \frac{1}{\kappa} \frac{A}{C}. \quad (5.10)$$

In practice, the amplitude is nearly always weak enough that Eq. 5.10 is a good approximation. Of course if $\frac{A}{C}$ is not small, then the oscillation is no longer sinusoidal. We can still estimate the amplitude as

$$\text{amplitude} \approx \frac{1}{2\kappa} \ln \left(\frac{1 + \frac{A}{C}}{1 - \frac{A}{C}} \right) \quad (5.11)$$

Equation 5.10 or 5.11 gives us the amplitude (in, say, picometers) of electronic order as it appears in the topography. To estimate this corrugation, all we need is the decay constant κ , and the fractional CDW amplitude $\frac{A}{C}$ in the IDOS. Let's assume a work function of $\phi = 3$ eV so that $\kappa = \sqrt{\frac{8m\phi}{\hbar^2}} = 1.8 \text{ \AA}^{-1}$. The only task then is to estimate $\frac{A}{C}$. For that, we turn to the literature.

Long range nematic order was observed by Lawler *et al* in underdoped Bi-2212 [15]. Lawler defines a nematic order parameter as

$$O_n^Q = \frac{\text{Re}\tilde{Z}(Q_x) - \text{Re}\tilde{Z}(Q_y)}{\bar{Z}} \quad (5.12)$$

where $\tilde{Z}(Q_x)$ is the Fourier transform of a Z-map evaluated at the atomic wavevector Q_x , and \bar{Z} is the spatial average of this map. The spatial average of O_n^Q is plotted versus energy in Lawler's Fig. 3c. The maximum magnitude of $O_n^Q \approx 0.01$ gives us an upper limit for $\frac{A_x}{C} - \frac{A_y}{C}$. Plugging this value in to Eq. 5.10 shows us that the difference in topographic corrugation amplitude between Q_x and Q_y is at most 0.5 picometers. A typical corrugation amplitude in a topography is 25 picometers, so

we should expect the Bragg peaks in a topography to have the same weight to better than 2%. We may ascribe differences greater than this to an anisotropic tip. Thus β can be used as an indicator of the anisotropy of the STM tip.

5.3 Stoll Formula

In STM experiments, a directional tip can create the appearance of broken crystal symmetry. Indeed we saw in the previous section that a directional tip is the *only* good explanation for a difference in amplitude between two lattice directions in BSCCO. If an STM tip is wider in one direction, the atomic corrugations of a topo will appear to have a smaller amplitude in that direction. This is also true of periodic oscillations at other wavelengths.

Stoll *et al* [31] have used a simple model to quantify the amount that a periodic oscillation is attenuated by a tip of finite size. The model is illustrated in Fig. 5.1. An STM tip terminating in a half-sphere of radius R , tunnels electrons through a barrier of width d into a conductor whose surface has geometric corrugation of wavelength λ and amplitude A . For large tip-sample separation ($d \gg 2/\kappa$) and large wavelength ($\lambda \gg \pi/\kappa$) there is a simple analytic expression for the amplitude h of the apparent corrugation seen in the STM image³:

$$\frac{h}{A} = \exp\left(-\frac{\pi^2(R+d)}{\kappa\lambda^2}\right) \quad (5.13)$$

As always, $\kappa = \sqrt{\frac{8m\phi}{\hbar^2}} \approx 1.8\text{\AA}^{-1}$ is the exponential falloff of the wavefunction. Reasonable values of these parameters This equation captures what we intuitively know:

³Also found on p. 28 in Chen [6].




Figure 5.1: Cartoon of the tunneling geometry described by the Stoll formula, Eq. 5.13.

the larger the radius of curvature, the more the amplitude is attenuated in our image, as the tip effectively averages over a larger fraction of the wavelength. For wavelengths much longer than the radius of curvature, this effect is tiny, while short wavelengths are completely washed out.

Table 5.1: Reasonable values of parameters in Eq. 5.13

π/κ	1.7\AA
$2/\kappa$	1.0\AA
d	2\AA
A	25pm
λ	$\gtrsim a_0$

Now we modify the Stoll formula, allowing the radius of curvature R to vary

with direction. Oscillations of the same wavelength λ and amplitude A , but with wavevectors in different direction may have different apparent amplitude h . Since we seldom have an accurate measure of the radius of curvature of the tip, we instead use this formula to describe a ratio of amplitude to some reference direction θ_{ref} :

$$\frac{h(\theta)}{h_{\text{ref}}} = \exp\left(-\frac{\pi^2}{\kappa\lambda^2}(R(\theta) - R_{\text{ref}})\right) \quad (5.14)$$

5.4 Recipe for Tip Anisotropy Correction

The Stoll formula suggests a simple recipe to correct anisotropic tips.

- We model our tip with an ellipsoidal shape such that the radius of curvature depends on angle according to

$$R(\theta) = R_{\text{ref}} + \Delta R \sin^2(\theta - \theta_{\text{ref}}). \quad (5.15)$$

Furthermore, we will assume the axes of the ellipsoidal tip are aligned with the crystal axes. We will revisit this assumption later. Finally, we choose the reference direction to be along a crystal axis as well. Let $\Delta R = R_2 - R_{\text{ref}}$ be the difference in radii between the two principal axes, and θ_{ref} be the angle at our reference direction, at which the tip has its shortest radius of curvature, R_{ref} .

Making the θ dependence explicit, Eq. 5.14 becomes

$$\frac{h(\theta)}{h_{\text{ref}}} = \exp\left(-\frac{\pi^2}{\kappa\lambda^2}\Delta R \sin^2(\theta - \theta_{\text{ref}})\right). \quad (5.16)$$

In the direction where $\theta = \theta_{\text{ref}}$, of course $h(\theta)/h_{\text{ref}} = 1$.

- Quantify the Bragg peak anisotropy β according to Eq. 6.3. To get an accurate measurement, we drift-correct so that the Bragg peak weight falls into a single

pixel. Then we Fourier transform the drift-corrected topo using a sine window; this prevents “leakage” of the spectral weight, which is now contained in the pixel at Q_x and its nearest neighbors to better than 1 part in 10^4 . To get the numerator (denominator) in Eq. 6.3, simply add in quadrature the values of $\text{Abs}(\text{Fourier transform})$ at the pixel Q_x (Q_y) and its nearest neighbors.

- To simplify the algebra going forward, we set $\theta_{\text{ref}} = 0$. Choose polar coordinates such that the line $\theta = 0$ aligns with the crystal axis with the highest Bragg peak amplitude in the topo. Call this Bragg peak Q_1 . In other words, if $\beta > 1$ then let $Q_1 = Q_x$ and choose $\theta = 0$ along the Q_x -direction, while if $\beta < 1$ then $Q_1 = Q_y$ and $\theta = 0$ in the Q_y -direction. With this convention,

$$\frac{h(\theta)}{h_{\text{ref}}} = \exp\left(-\frac{\pi^2}{\kappa\lambda^2}\Delta R \sin^2\theta\right). \quad (5.17)$$

- Now we apply β . Because it is defined at the atomic periodicity $\lambda = a_0$, and because the lattice directions are 90° apart, we can write

$$\beta = \frac{h(Q_x)}{h(Q_y)} = \exp\left(-\frac{\pi^2}{\kappa a_0^2}\Delta R \sin^2(90^\circ)\right) = \exp\left(-\frac{\pi^2}{\kappa a_0^2}\Delta R\right) \quad (5.18)$$

so that Eq. 5.17 becomes

$$\frac{h(\theta)}{h_{\text{ref}}} = \exp\left(-\frac{\pi^2}{\kappa a_0^2}\left(\frac{a_0^2}{\lambda^2}\right)\Delta R \sin^2\theta\right) = \beta^u \quad (5.19)$$

with

$$u = \left(\frac{a_0 k \sin\theta}{2\pi}\right)^2.$$

- We use this formalism to make a k-space correction factor $F_{\text{Stoll}}(\mathbf{k})$ that corrects for amplitude attenuation at wavevector $\mathbf{k} \equiv k(\cos\theta, \sin\theta)$ by comparing it to

the attenuation at Q_1 :

$$F_{\text{Stoll}}(\mathbf{k}) = \frac{h_{\text{ref}}}{h(\mathbf{k})} = \beta^{-u(\mathbf{k})} \quad (5.20)$$

or in terms of the magnitude $\alpha = \frac{k}{2\pi/a_0}$ it is just

$$F_{\text{Stoll}}(\alpha, \theta) = \beta^{\alpha^2 \sin^2 \theta} \quad (5.21)$$

An image $G(\mathbf{r})$, with Fourier transform $G(\mathbf{k})$, is corrected as follows:

$$G_{\text{corrected}}(\mathbf{r}) = \mathcal{F}^{-1}\{F_{\text{Stoll}}(\mathbf{k})G_{\text{uncorrected}}(\mathbf{k})\} \quad (5.22)$$

In words, you Fourier transform the image, multiply by the correction factor $F_{\text{Stoll}}(\mathbf{k})$ and inverse transform to get the corrected image. An example correction factor is shown in Fig. 5.2, next to the Fourier transform of the R-map to be corrected. The correction changes all Fourier amplitudes into the values they would have had if the STM tip were isotropic with radius R_{ref} .

- It is illuminating to estimate ΔR using Eq. 5.18. Using $a_0 = 3.83\text{\AA}$ and the same value $\kappa = 1.8\text{\AA}^{-1}$ from earlier in this chapter, we find for all of our tips that ΔR is less than 5\AA . Figure 5.3 shows the mathematical identity between ΔR and β with these assumptions. Of course when $\beta = 1$ the tip is isotropic with respect to the lattice directions, so $\Delta R = 0$ and no correction is necessary.

5.4.1 Tip Angle

The elongated direction of a directional tip will not generally align with one of the crystal axes. If we somehow knew the angle, it would be straightforward to generalize




Figure 5.2: **(a)** Fourier transform of an image to be corrected. The $\mathbf{k} = 0$ point lies in the center of this and all Fourier transforms in this dissertation. **(b)** k-space Stoll correction factor $F_{\text{Stoll}}(\mathbf{k})$, to be multiplied by the Fourier transform at left. $\theta_{\text{ref}}=0$ is the direction at which the correction equals 1. For this data set, $\beta = 1.4$.




Figure 5.3: Mathematical relationship between tip anisotropy β and Stoll length ΔR from Eq. 5.18, assuming $a_0 = 3.83\text{\AA}$ and $\kappa = 1.8\text{\AA}^{-1}$. The tip is isotropic at $\beta = 1$, so ΔR goes to zero.

Eq. 5.21. Then the correction process in Eq. 5.22 would be accurate in all directions θ . But it is hard, if not impossible, to measure this exact elongation angle. Order in the sample, especially atomic periodicity, tends to skew this measurement. But even without the exact tip angle, we can tell the relative attenuation between wavevectors

along the crystal axes. Fortunately, in this thesis, we focus on two such wavevectors, Q^* and Q^{**} .

5.5 Limitations

Before showing how amazingly well this algorithm works, I will mention a few of its limitations.

Stoll tip correction is not the same as tip deconvolution. The corrected image is not a “true” representation of the crystal. Any STM image was taken with a tip of finite size R , and therefore the higher spatial frequencies are attenuated. The correction only transforms the image into one taken with an isotropic tip (of finite radius R_{ref}). It does not give us access to spatial frequencies beyond those we have already sampled. And due to the condition $\lambda \gg \pi/\kappa \approx 1.7\text{\AA}$ (see Table 5.1), it will not generally be accurate for wavelengths $\lambda \lesssim a_0$. It will not recreate topographic features with substantial three-dimensional structure. Nor will it fix double tips or correct for other kinds of funny-shaped tips. It is usually obvious from an autocorrelation of the topography whether you have such a wonky tip. None of the datasets in this thesis were taken with uncorrectable tips, however we excluded some datasets that were.

Even for simple directional tips, there are some data sets that are beyond correction. If a tip is *too* directional, that is if $\beta > 2$ (or $\beta < 0.5$) the Stoll correction tends to overcorrect. Similarly, if the tip is elongated along a direction θ_{ref} that is more than 30° from a lattice direction, a large correction will be required to make up for any observed Bragg peak anisotropy β . In these cases too, it tends to overcorrect.

Rather, Stoll correction is optimal for studies on atomically flat surfaces will small

amplitude electronic oscillations at a variety of wavelengths $\lambda \leq a_0$. The image must have atomic resolution. For the formulas in this paper to apply, the crystal structure must have C_4 symmetry, although there is no obvious reason why they could not be generalized to crystals with other symmetry. Although the Bi-O plane of BSCCO has a small orthorhombic distortion, the effect is small ($\sim 1\%$ of a unit cell) and equally affects the weight of the atomic peaks in the x- and y-directions [41]. Stoll correction enables quantitative comparison between the amplitude of electronic order in separate crystal directions.

5.6 Results

Here we demonstrate the Stoll correction method on a data set of underdoped Bi-2201 (sample A). Its initial tip anisotropy is $\beta = 1.38$, corresponding to a Stoll length of $\Delta R \approx 0.88\text{\AA}$. This initial anisotropy is visualized in Fig. 5.4a, which shows linecuts through the Bragg peaks in the Fourier transform of the topography. After Stoll correcting, there is no Bragg peak anisotropy remaining in the topography (Fig. 5.4b). The reference direction is along the Q_x direction. Notice that the autocorrelation is smeared-out in the perpendicular crystal direction Q_y (see Fig. 5.6a), which indicates that the tip is elongated along in the Q_y -direction. As mentioned earlier, if the angle is within 30° of this axis, the correction factor remains practically the same.

Figure 5.5 shows a drift-corrected topography before and after correcting. Although the high frequency components have been slightly enhanced by Stoll correction in (b), the difference is very subtle. The difference becomes apparent in the autocorrelation. Figure 5.6 shows autocorrelations of the topographies in Fig. 5.5,




Figure 5.4: Linecuts through the Fourier transform of the sample A (UD32K) topography ($\beta=1.38$) in the Q_x (red) and Q_y (blue) directions near the Bragg peaks. **(a)** Before correcting, and **(b)** after correcting.

with resolution enhanced by a simple fourier technique to illustrate the atomic corrugation. Only a zoomed-in area of the autocorrelation around the central peak is shown. Linecuts through the x- and y-directions of the autocorrelations illustrate how Stoll correction corrects the amplitude of atomic corrugation between the two directions.

So far, we have only applied our correction to topographic images. The real utility of Stoll correction lies in looking at electronic order in DOS maps. In a typical STM experiment, a topography is taken at the same time as the DOS map, allowing us to use the topo to Stoll-correct the entire map. To do this, we use the topography to make the correction factor $F_{\text{Stoll}}(\mathbf{k})$ in Eq. 5.21. Then we use this to correct any simultaneously- acquired map that we like, using Eq. 5.22. Figure 5.7 shows the result of Stoll-correcting an 80mV R-map in the same underdoped sample described above, in order to look at the Q^{**} peak. The Fourier transform before and after correcting is shown in Fig. 5.7. Linecuts in the Q_x and Q_y directions clearly show the Q^{**} peak. Initially, it appears that Q^{**} has much higher amplitude in the Q_x direction. After




Figure 5.5: **(a)** Topography (Sample A, UD32K) before correcting and **(b)** after correcting. The difference is subtle. The field of view is 68nm wide.

Stoll correcting, this difference has largely disappeared.

Not surprisingly, tip anisotropy affects real space measures of Q^{**} as well. In Chapter 6, we will define a global order parameter n , equal to the fraction of pixels in which the Q^{**} amplitude is larger in the x-direction than the y-direction. Figure 5.8 shows that n is sensitive to tip anisotropy but that Stoll correction removes this sensitivity. In Fig. 5.8a, n is plotted against the tip anisotropy β , for six different samples. Without correcting, these two quantities are highly related(a). The same figure after correcting is shown in (b). We expect this fraction to be close but not exactly equal to $n = 0.5$. The spread is much smaller after Stoll correcting. The exception is the outlying point on the left, with initial anisotropy $\beta = 0.2$. For extremely directional tips such as this one, the algorithm overcorrects, as it has done here.




Figure 5.6: **(a)** Autocorrelation of a drift-corrected topography before Stoll correction. Only a width $7a_0$ from the center of the autocorrelation is shown. The resolution has been enhanced by padding the Fourier transform with zeros. **(b)** Linecuts through the central peak of the autocorrelation along the X (red) and Y (blue) directions, clearly showing the difference in amplitude before correcting. **(c)** and **(d)** Autocorrelation and linecuts through the same topography, after Stoll correcting.

5.7 Conclusion

We have presented an algorithm for correcting for small amounts of tip anisotropy in STM images. Stoll's formula has allowed us to take a measurable tip anisotropy of atomic periodicity and correct for it at all wavelengths. It enables quantitative study of symmetry breaking, for order with wavevector along a crystal axis direction, and wavelengths equal to or larger than a_0 . This includes many charge ordering phenomena in the antinodal direction in BSCCO including nematic order, and the




Figure 5.7: **(a)** Fourier transform of the sample A (UD32K) 80mV R-map examined in Chapter 4, with no correction. **(b)** Linecuts in the Fourier transform in (a), in the direction of the atomic Bragg peaks. **(c)** Fourier transform of the R-map in (a) after Stoll correcting. **(d)** Linecuts through the Bragg peaks in (c).




Figure 5.8: **(a)** Fraction of points n (“ Q_x fraction”) where Q^{**} is stronger in Q_x direction, plotted vs tip anisotropy β for 7 different tips and samples. **(b)** The same plot, after correcting tip anisotropy, plotted vs initial (uncorrected) tip asymmetry. **(c)** Final Q_x fraction vs initial Q_x fraction. The spread is much larger before correcting than after. The lines at $n = 0.5$ show the value we expect very near to the critical point, with a perfectly isotropic tip.

Q^* and Q^{**} modulations.

Without this algorithm, nothing in Chapter 6 would be possible. Studies of the Q^{**} peak are especially sensitive to a directional tip. By virtue of its longer wavelength, Q^* is hardly affected by tip anisotropy. We anticipate that Stoll correction will make many more STM data sets available to analyze C_4 symmetry breaking in the cuprates.

Chapter 6

Local Anisotropy of the Q^{**} Peak

6.1 Background

In Chapter 4 we examined the degree to which the two CDW peaks Q^* and Q^{**} were related to one another. We found that, while they likely have the same physical cause, they respond independently to pinning disorder. In this chapter, we further characterize the CDW disorder by looking at the difference in amplitude between order in x- and y-directions. Authors (e.g. Kivelson [12] and Vojta [35]) routinely assume that CDW disorder in BSCCO is “random field disorder,” as opposed to “random mass” (aka random bond) disorder. And indeed it seems intuitive that randomly-placed dopant atoms should lead to random field disorder. The location of interstitial oxygen dopants was shown to correlate with the locations of troughs of quasiparticle interference (QPI) modulations [17]. In this chapter we will test experimentally the assumption that the CDW experiences random field type disorder, by mapping the preferred orientation of the CDW in our data onto an Ising variable,

and then using the tools of disordered statistical mechanics.

We analyze only the Q^{**} peak: its shorter wavelength gives us good spatial resolution. Our starting point in each sample is the same R-map that we used in Chapter 4 to highlight Q^{**} . The map has first been drift-corrected and Stoll-corrected. This time we include all of the samples in Table 1.1, which includes two overdoped samples. The Q^* peak is hard to see in overdoped samples because its wavelengths overlap almost completely with the central peak, and because the checkerboard amplitude gets smaller as we get closer to the T^* line [22]. However, since we will not analyze Q^* in this chapter, there is no barrier to using overdoped samples. We have, however, excluded data sets taken with extremely elongated tips, or other funny-shaped tips, which cannot be corrected with the Stoll procedure in Chapter 5.

Figure 6.1 shows a 30nm subfield of the R-map for sample B (OPT35K) along with its simultaneous topography and a Fourier transform of the entire field of view. The Q^{**} peak is clearly visible in the Fourier transform, part of a quartet of CDW peaks surrounding each reciprocal lattice vector.

Outline of this chapter

In this chapter we start with a brief discussion of whether Q^* and Q^{**} can be described as stripes or checkerboard. Then we describe our procedure for mapping the local CDW orientation. Then we interpret this orientation as an Ising map, and describe the procedure to calculate two critical exponents from the clusters of Ising spins. Critical exponents from samples of several dopings are then compared to theoretical predictions for different classes of a disordered Ising model. The 3D random field class is found to agree most closely with the measured critical exponents.




Figure 6.1: (a) Topographic STM image of sample B (OPT35K), taken at 6K. The field-of-view is 30nm wide. (b) An R-map of the same field of view as (a), created by dividing $I(+80\text{mV})/I(-80\text{mV})$. Shown are the correlation lengths computed according to Eqs. 4.5 and 4.6 in Robertson, *et al* (2006)[28]. (c) Fourier-transform of the R-map (from the entire 72nm field of view of sample B), clearly showing a quartet of CDW peaks surrounding each reciprocal lattice vector. (d) Correlation lengths of five samples of Pb,Bi-2201 plotted as a function of doping. The doping=0.16 points corresponds to the R-map in (b). The doping values are assigned using the Presland formula [27], which may be inaccurate for Bi-2201.

6.2 Stripes or Checkerboard?

By what justification do we treat the x- and y-directions separately? We postulate that the charge order is inherently unidirectional in our samples. Doping causes the cuprate superconductors to be inherently disordered. In the absence of disorder,

supposing we could image this fluctuating order, would the CDW preserve the C_4 symmetry of the lattice and be a “checkerboard”, or would it break this symmetry to form “stripes”? In the fluctuating regime and in the presence of disorder, it is extremely hard to tell the difference between stripes and checkerboard [28].

We postulate that Q^* and Q^{**} are disordered stripes in all of our data sets. To tell the difference between the two scenarios, Robertson *et al* [28] defined the two quantities

$$\xi_{\text{orient}} = \sqrt{\frac{|\int d\mathbf{r} [|A_x|^2 - |A_y|^2]|^2}{\int d\mathbf{r} (|A_x|^2 + |A_y|^2)^2}}, \quad (6.1)$$

and

$$\xi_{\text{CDW}} = \sqrt{\frac{|\int d\mathbf{r} A_x|^2}{\int d\mathbf{r} |A_x|^2} + \frac{|\int d\mathbf{r} A_y|^2}{\int d\mathbf{r} |A_y|^2}}. \quad (6.2)$$

where A_x and A_y are the complex local fourier amplitude of the CDW in the x- and y-direction. The condition $\xi_{\text{orient}} > \xi_{\text{CDW}}$ indicates the presence of an “orientational glass,” which implies stripe order, although $\xi_{\text{orient}} < \xi_{\text{CDW}}$ is not sufficient to show checkerboard order. Similar ideas are put forth in Ref. [7]. Indeed, in one of our samples (sample B, OPT35K) the correlation length of the Q^{**} stripe orientation is almost a factor of ten longer than the CDW correlation length. There is no question that in this particular OPT35K sample the charge density wave has a stripe form. For the other samples it is more ambiguous. The values of these correlation lengths ξ_{CDW} and ξ_{orient} are shown for our data in Fig. 6.1. One possible cause for the anomalously long ξ_{orient} in the optimally doped sample is that the hole concentration could be close to $p = 1/8$. Although the Presland formula stipulates that hole concentration at optimal doping is $p = 0.16$, there is a feeling in the community that this formula is not accurate for Bi-2201, and that the superconducting dome is shifted towards

lower T_c and lower doping relative to other BSCCO compounds¹. On the other hand, the wavevector of the checkerboard in this sample is less than $0.20(2\pi/a_0)$, not the $\alpha = 0.25$ that we'd expect for charge ordering at the magic 1/8 doping.

6.3 Mapping local anisotropy

Our goal in this section is to map the preferred direction of the stripes. Our method is simple: first, we use the Spatial Lockin algorithm (Chapter 3) to compute the local CDW amplitude in the x- and y-directions. The local CDW amplitude is computed by finding the local amplitude at wavevectors ranging from $\alpha = 0.65$ to 0.95 and taking the average of these amplitudes. Our map changes only slightly when we use different methods to find the local amplitude. We define the local CDW anisotropy in terms of the local CDW Fourier amplitudes $\tilde{G}(Q_x^{**}, \mathbf{r})$ and $\tilde{G}(Q_y^{**}, \mathbf{r})$:

$$\text{CDW anisotropy} = \tilde{G}(Q_x^{**}, \mathbf{r}) - \tilde{G}(Q_y^{**}, \mathbf{r}). \quad (6.3)$$

This is similar to the order parameter used by Lawler to describe “smectic order” of the Q^{**} peak[15]. Where this quantity is greater than zero, the stripes lie primarily in the Q_x -direction.

Figure 6.2 shows the result of this binarization process, on sample B (OPT35K). The red regions have stripes primarily in the x-direction, while the blue regions are primarily in the y-direction. The image has been Fourier filtered around the Q^{**} peak

¹The closest thing to a measurement of hole doping in Bi-2201 of which I am aware is Ref. [3], which measured the related compound BSLCO, which like Bi-2201 has a maximum T_c of $\sim 35\text{K}$. The curve of T_c vs. hole concentration was found to peak near $p = 0.16$, although it otherwise did not follow the Presland curve. There have been many earlier attempts to determine the phase boundary, but according to Ando they suffer from ambiguities from the Bi valence.

to make it easier to see the stripe direction by eye. Fourier transforms of the red and blue regions are shown in (b) and (c) respectively. The Q^{**} peak is clearly stronger in the x-direction in (b) and in the y-direction in (c). This shows us that we have successfully masked the image by stripe direction.




Figure 6.2: **(a)** The R-map from Fig. 6.1b, colored to indicate stripe orientation. The full 70nm field of view is shown. Stripes in the red regions prefer to align along the Q_x -direction (the wavevector points up and to the right in this image) while those in the blue region point in the orthogonal Q_y -direction. For visual clarity, the R-map has been fourier-filtered to include only the Q^{**} peak. 61% of the field of view is red (and the rest blue.) **(b)** Fourier transform of the red regions of the R-map (without fourier-filtering.) **(c)** Fourier transform of the blue regions of the R-map.

When we apply the same masking process to an underdoped data set, we get the map in Fig. 6.3. As before, the Fourier transform of the red region shows a strong

Q^{**} peak in the x-direction (b) while the blue regions have a strong y-direction Q^{**} peak(c). This time, we additionally apply the same red-blue mask to an integrated $Z_{(-)}$ map, in order to highlight the Q^* peak. The resulting Fourier transforms of the red regions (d) and blue regions (e) show strong Q^* peaks in, respectively, the x- and y-directions. In other words, masking by the orientation of Q^{**} stripes also effectively separates the Q^* peak into its separate orientations. This is a very interesting result, which strongly implies that Q^* and Q^{**} reflect the same physics.

The stripes in the x-direction occupy a nearly equal fraction of the image as those in the y-direction. This is true for all five of the Bi-2201 data sets we examined. This lack of long range order in the orientation of the Q^{**} peak agrees with earlier studies in underdoped Bi-2212 [15]. Sample B (OPT35K, seen in Fig. 6.2) is a possible exception to this rule. In sample B, 61% of the area is red. Let us define a global orientation order parameter

$$n = \frac{\text{area of red region}}{\text{total area}} \quad (6.4)$$

to look for trends among our five samples. Since the choice of the x- and y-axes shouldn't matter, we plot in Fig. 6.4 the quantity $(n - 0.5)^2$ vs the image size(a), the tip anisotropy β before correcting, and doping (c). There is no obvious systematic error in n due to image size or Stoll correction. Neither is there any obvious trend in n with doping. The importance of Stoll correction in obtaining this result cannot be overstated. If we don't Stoll correct, n is highly influenced by tip anisotropy (as we mention at the end of the preceding chapter.) We see that there no long range Q^{**} order at any doping, except possibly in the optimally-doped sample B.




Figure 6.3: **(a)** 80mV R-map of sample A (UD32K, nominal doping $p=0.128$), colored to indicate stripe orientation. The full 65nm field of view is shown. Stripes in the red regions prefer to align along the Q_x -direction (the wavevector points up and to the right in this image) while those in the blue region point in the orthogonal Q_y -direction. For clarity, the R-map has been fourier-filtered to include only the Q^{**} peak. **(b)** Fourier transform of the red regions of the R-map (without fourier-filtering) showing strong Q^{**} peaks in the x-direction **(c)** Fourier transform of the blue regions of the R-map showing strong Q^{**} peaks in the y-direction. **(d)** Fourier transform of an integrated $Z_{(-)}$ map masked by the red regions showing strong x-direction Q^* peaks. **(e)** Fourier transform of an integrated $Z_{(-)}$ map masked by the blue regions showing strong y-direction Q^* peaks.




Figure 6.4: Difference between Q_x fraction and 0.5 (ideal value), squared, plotted against three different parameters, for five different images. There are no obvious trends. The x-axes are: (a) Image size. (b) tip anisotropy β before correcting. (c) Nominal doping.

6.4 Finding critical exponents

In this section we view the local preferred orientation of the CDW as an Ising variable, associating one Ising spin with each pixel of our map. From the areas and perimeters of the clusters in our Ising map, we can determine two critical exponents that we can use to characterize the disorder of the CDW. For the five Bi-2201 data sets examined in this chapter, Ising Q maps are shown in Fig. 6.5. We will next define the critical exponents and describe in detail how they are computed. The sizes of the maps in this figure scales with the sizes of the maps themselves.

We employ a type of analysis known as “Ising Cluster Analysis,” described in




Figure 6.5: Ising maps from the five Bi-2201 maps we examined in this Chapter, shown to scale. (a) Sample C (UD25K), 52nm, 101 clusters. (b) Sample A (UD32K), 68nm, 124 clusters. (c) Sample B (OPT35K), 72nm, 160 clusters. (d) Sample E (OD15K), 56nm, 130 clusters. (e) Sample D (OD16K), 36nm, 36 clusters.

detail in, e.g. Ref. [30]. To compute the critical exponents, we view each Ising map as a collection of “clusters,” contiguous regions of the map with the same spin. Not all clusters are created equal. Clusters that appear inside other clusters are known as “subclusters,” and clusters that touch the edge are “edge clusters.” There is typically one big cluster known as the “spanning cluster” of which all other clusters are subclusters (except edge clusters.)

The critical exponents we can extract from Ising clusters are τ and d_c/d . They

are defined in terms of the perimeter p and area A of the clusters:

$$D(A) \propto A^{-\tau} \quad (6.5)$$

and

$$p \propto A^{d_c/d}. \quad (6.6)$$

$D(A)$ is the number of clusters of a certain area, per octave of area. The quantity d_c is the fractal dimension of the surface of the spanning cluster, while d is the dimensionality of the order being studied.

Finding the area of clusters in a pixellated image is straightforward: you simply count the pixels that belong to a cluster. There is one subtlety, however: the cluster area used to find d_c/d counts the area of a subcluster toward the area of the cluster containing it, while the area used to find τ does not include subclusters. The perimeter of a cluster is the number of edges of that cluster. For example, a cluster of one pixel area always has a perimeter of 4, unless it is an edge cluster. Similarly, a cluster of two pixels always has perimeter of 6. Edges of the image are not counted toward cluster perimeter. For small clusters, then, there is a well-defined relationship between cluster area and perimeter. For this reason, when fitting to find the exponent d_c/d , we do not include clusters of size less than 10 pixels. There are, however, few of these clusters to begin with, due to the smoothing length scale L needed to find the CDW amplitude in the Spatial Lockin algorithm.

There is a third exponent that in principle we could compute. It relates the spin-spin correlation function $G(r) = \langle S(r)S(r') \rangle - \langle S(r) \rangle \langle S(r') \rangle$ to distance as $G(r) \propto |r|^{-(d-2-\eta)}$. We attempted to find this third exponent, but there is no range of areas over which this power law behavior is obeyed. Instead, we observe the correlation

length behavior $G(r) \propto e^{-r/\xi}$ over all length scales that we could observe. This may mean that we are not close enough to the critical point to see power law scaling in $G(r)$.

Plots used to determine the critical exponents τ and d_c/d are shown in Fig. 6.6. The clusters have been logarithmically binned by area to produce these plots. We expect plots of $D(A)$ to be under-populated for small cluster sizes, because the spatial lockin has a finite coarsening length scale of $L = 2a_0 \approx 4\text{-}5$ pixels. We therefore do not fit to clusters with area less than 25 pixels. Similarly there are not many large clusters because of the finite size of the image; we exclude clusters with area more than 6% the area of the entire field of view. All areas and perimeters are reported in units of pixels.

One further check of accuracy is needed. In the process of binarizing the CDW anisotropy (Eq. 6.3) to make an Ising map we chose to binarize by cutting off at zero. It is possible that the Stoll correction algorithm has left behind a small amount of residual tip anisotropy, skewing the histogram. How much does this binarization cutoff affect value of the critical exponents? There is no change within error bars for τ , while d_c/d can vary by as much as $\pm 5\%$. To see this, we vary the cutoff and measure the resulting critical exponents, in Fig. 6.7. The process is like artificially tuning the anisotropy of our STM tip.

Fig. 6.7d shows a histogram of the CDW anisotropy for sample A (UD32K, our $p = 0.128$ data point.) Its mean lies 0.04σ away from zero, on the negative side. In Fig. 6.7a,b we plot τ and d_c/d for cutoff values ranging from nearly $-\sigma$ to $+\sigma$. For all other figures in this chapter, we have taken the cutoff at exactly zero. The mean




Figure 6.6: (a),(c) Plots of $D(A)$ vs A , to obtain the exponent τ . (b),(d) Plots of Perimeter vs Area, to obtain the critical exponent d_c/d . Clusters have been logarithmically binned by area to produce these plots. The choice of points to fit is described in the text. Error bars correspond to the standard deviation of the mean of the clusters in each bin. Points with no error bars occur when only one cluster is in the bin, and are not included in the fits.




Figure 6.7: **(a)** Critical exponent τ plotted vs the binarization cutoff for sample A (UD32K), used to turn the CDW amplitude into an Ising map. **(b)** Critical exponent d_c/d plotted for the same cutoff as (a). The points where the cutoff equals zero are the ones we plot in Fig. 6.8. The point directly to the left of zero cutoff (cutoff = -0.0015) corresponds to the mean value of the anisotropy before binarizing. **(c)** Fraction of pixels n in which Q_x is stronger than Q_y , plotted vs cutoff. **(d)** Histogram showing the range of values of the anisotropy before binarizing. The mean value in the histogram is 0.04σ to the left of zero.

of the histogram in 6.7d corresponds to the plot point immediately to the left of zero in plots 6.7a,b. Values far from the center are extremely unlikely to be the correct exponents, and are shown only to reflect the spread. Fig. 6.7c shows how n varies with cutoff; only the points very close to the center represent n that have been observed in our data sets after Stoll correction. We find from this exercise τ is insensitive to

cutoff within its error bars, for reasonable values of the cutoff. The critical exponent d_c/d is more sensitive however, with values ranging from 0.60 to 0.67 for this sample. Luckily, we see in the next section that this spread is much smaller than the difference in d_c/d values between different theories, so our conclusion is independent of cutoff.

6.5 Results: Discussion of critical exponents

The clusters in Fig. 6.5 show the power law scaling behavior expected in proximity to a critical point. In some cases, the scaling is observed over nearly three decades, a decent value for experimental data. In this section, we compare the values of the exponents we found in Fig. 6.6 to one another and to theoretical models for Ising systems with different kinds of disorder.

Figure 6.8 shows the values of τ and d_c for samples of different doping. The value of the fractal dimension d_c was obtained by taking $d = 2$ because we are imaging a 2D surface. τ seems to decrease with doping, while d_c shows no clear trend. It is worth noting that sample B (OPT35K), the same one to have a large orientational correlation length, also has a much larger fractal dimension than the other samples.

The exponent τ is theoretically constrained to $2 < \tau < 3$ [30]. However, we should not be troubled that our values of τ uniformly violate the theoretical constraint. There is a known systematic error that causes underestimates of τ on Ising maps with small overall size [23]. If we had several more decades of scaling we might avoid this problem. Furthermore, in all of the models examined, τ takes values ranging from only 2.0 to 2.25. Clearly the critical exponent τ will not be very useful to us in determining which class of disorder expressed by CDW orientation. As we see shortly,




Figure 6.8: **(a)** Critical exponent τ plotted vs nominal doping for six samples of Pb,Bi-2201. **(b)** Fractal dimension d_c plotted for the same data sets as (a). Doping values are assigned using the Presland formula, which may be inaccurate for Bi-2201.

d_c/d is much more valuable in constraining the model.

Theoretical values for the fractal dimension d_c have been carefully collected by Phillabaum *et al* for Ising models with different types of disorder, in two and three dimensions [24]. In the 3D limit, the theoretical dimension shown is the fractal dimension of a 2D cross section. The values are shown by the solid circles in Fig. 6.9. The theoretical values consider disorder of both random bond and random field type in the zero temperature limit. This disorder is due, whether directly or indirectly, to the random locations of dopant atoms within the sample. It competes with the ferromagnetic coupling between “spins” (the local CDW orientation at each pixel.) Since cuprate superconductors are highly layered materials, we allow the Hamiltonian to have different coupling strength between spins within the ab plane (J^{\parallel}) and between spins connected in the c axis direction (J^{\perp}). The closer the ratio J^{\perp}/J^{\parallel} is to one, the more the material takes on a three-dimensional character. The disordered

Hamiltonian is

$$H = - \sum_{\langle ij \rangle_{\parallel}} J^{\parallel} (1 + \delta J_{ij}^{\parallel}) \sigma_i \sigma_j - \sum_{\langle ij \rangle_{\perp}} J^{\perp} (1 + \delta J_{ij}^{\perp}) \sigma_i \sigma_j - \sum_i h_i \sigma_i \quad (6.7)$$

The term h_i is the random field at each lattice site.

We turn now to Figure 6.9 to compare the range of fractal dimensions we saw in Fig. 6.8b to the theoretical values collected by Phillabaum [24]. The figure shows a smooth interpolation between the 2D value for d_c and one half of the 3D value for d_c (because, if you take a cross section of a 3D fractal, the fractal dimension of the cross section is one half the 3D value.) The interpolation between 2D and 3D values is achieved by tuning the ratio J^{\perp}/J^{\parallel} . The range of d_c values that we observe (light blue band in Fig. 6.9) is most consistent with an Ising model of Random Field disorder near 3D. This Ising map procedure has been repeated in one sample of optimally-doped Bi-2212, which has two copper oxide layers in each unit cell. The Bi-2212 fractal dimension d_c is $d_c = 1.34 \pm 0.02$. There appears to be no difference in fractal dimension between the two BSCCO compounds.

6.6 Conclusion

In this chapter we have analyzed the orientation of the CDW stripes in Bi-2201 using methods from disordered statistical mechanics, thus forging a new link between theory and experiment. In all samples we observe power law scaling behavior in perimeter and area over 2–3 decades, indicating that our samples are near a critical point. In all but one sample, long range orientational order is absent: at small length scales the stripes break C_4 symmetry, although globally they do not. This is consistent




Figure 6.9: Comparison of the measured fractal dimension d_c with different disorder classes of the Ising model. Blue circles are the theoretical values in 2D. Purple circles are the expected fractal dimension of a 2D cross-section taken from the 3D Ising map. The dark blue shaded lines interpolate theoretical d_c values between the 2D and 3D regimes by tuning the coupling strength between spins in adjacent relative to those in the same layer. Figure adapted from Phillabaum *et al* (2011) [24]. The range of experimental d_c (light blue band) includes the Bi-2201 values plotted in Fig. 6.8b, as well as one sample of OPT Bi-2212.

with the fluctuating side of a stripe-ordering phase transition. Critical exponents d_c/d and τ were extracted from these power law regimes. The fractal dimension d_c was compared to theoretical values for disordered Ising maps with different types of disorder and different dimensionality; the values we find indicate that the stripe orientation is subject to 3D Random Field disorder. Perhaps this reflects the pattern of dopant atoms distributed randomly in 3D throughout the sample. This is the same class of disorder reached by a similar analysis of nematic order in Dy,Bi-2212 [24]. We found no difference in fractal dimension d_c between the CDW order in Bi-2201 and Bi-2212.

Along the way, we also found that the Q^* and Q^{**} stripes locally favor the same

crystal orientation. This supports the argument that the two wavevectors are ordering vectors of the same CDW. It also tells us that the critical exponents disorder class we measured for Q^{**} likely apply to Q^* as well.

Chapter 7

Conclusion

This thesis drew on data from three STMs and Bi-2201 samples of several different dopings, to make three main observations about the modulations at two different wavevectors Q^* and Q^{**} .

We have seen strong new evidence about the relationship between Q^* and Q^{**} modulations. At a particular spot on the surface, if the Q^{**} stripes are pointing in the x-direction, it is likely that the Q^* stripes are pointing in the same direction. This is true even though the local stripe direction is highly disordered. Combined with the finding that the two peaks respond to doping and local doping according to the same nesting cartoon, this result paints a picture where Q^* and Q^{**} have the same physical cause. A challenge for the cuprate community is to reconcile this relationship with the very different energy behaviors of these peaks.

The next two results focus on the disorder itself. The disorder comes in two flavors, both of which are important. First is the disorder of the optimal local wavelength. This wavelength is set by the local doping—the local depth of the Fermi sea—and

is measured by the pseudogap. The optimal wavelength is obeyed on average, as we see by looking at samples of different doping, or masking by similar pseudogap. Charge order at both wavevectors Q^* and Q^{**} responds to doping on average in the opposite direction, such that the two vectors sum to the width of the Brillouin zone for any doping. But locally the picture is different: the CDW only chooses the optimal wavelength in locations where defect pinning is absent. Those regions are few and far between, and they aren't even the same regions for the Q^* and Q^{**} charge orders. Therefore we can tell pinning is important because the Q^* and Q^{**} wavevectors don't sum locally to the width of the Brillouin zone. Although Q^* and Q^{**} are caused by the same physics, pinning destroys the local relationship between them.

The local stripe orientation is highly disordered. Our second result says this disorder belongs to a particular class: random field disorder. While the cuprate community has long believed this to be true, there is now data to back up this belief. Our data showed that in all of our samples, the stripe orientation bears the signature fractal dimension of random field disorder. Surprisingly, this fractal dimension agrees more with random field disorder in its 3D limit. This is surprising because the layered crystal structure of the cuprates seems more in keeping with two dimensional physics. We speculate that this result reflects the 3D nature of the distribution of dopant atoms throughout the crystal, and that dopant sites in one layer are not independent of those in neighboring layers.

Two new algorithms made these results possible. The first, the Spatial Lockin algorithm, uses a series of local fourier transforms to find the local wavevector. The second, Stoll correction, corrects for slight deviations from an isotropic STM tip.

Without this correction, a slightly anisotropic tip would make it impossible to gauge whether electronic order in the sample breaks C_4 symmetry, whether locally or globally. I believe these two algorithms will be useful analytical tools in future STM studies of the cuprates. The Stoll correction algorithm may have an even broader impact, as it is useful for STM studies of any material displaying broken rotational symmetry.

Bibliography

- [1] Website of EBL products, manufacturer of piezoelectric components, 2011.
- [2] P. Abbamonte, A. Rusydi, S. Smadici, G. D. Gu, G. A. Sawatzky, and D. L. Feng. Spatially modulated 'Mottness' in $\text{La}_{2-x}\text{Ba}_x\text{CuO}_4$. *Nature Physics*, 1(3):155–158, December 2005.
- [3] Y. Ando, Y. Hanaki, S. Ono, T. Murayama, K. Segawa, N. Miyamoto, and S. Komiya. Carrier concentrations in $\text{Bi}_{-2}\text{Sr}_{-2-z}\text{La}_{-z}\text{CuO}_{-6+\delta}$ single crystals and their relation to the Hall coefficient and thermopower. *Physical Review B*, 61:R14956, 2000.
- [4] M Bode. Spin-polarized scanning tunnelling microscopy. *Reports on Progress in Physics*, 66(4):523–582, April 2003.
- [5] U Chatterjee, M Shi, A Kaminski, A Kanigel, H M Fretwell, K Terashima, T Takahashi, S Rosenkranz, Z Z Li, H Raffy, A Santander-Syro, K Kadowaki, M R Norman, M Randeria, and J C Campuzano. Nondispersive fermi arcs and the absence of charge ordering in the pseudogap phase of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Physical Review Letters*, 96(10):107006, 2006.

-
- [6] C. Julian Chen. *Introduction to Scanning Tunneling Microscopy*. Oxford University Press, Oxford, U.K., 2 edition, 1993.
- [7] Adrian Del Maestro, Bernd Rosenow, and Subir Sachdev. From stripe to checkerboard ordering of charge-density waves on the square lattice in the presence of quenched disorder. *Physical Review B*, 74(2):1–10, July 2006.
- [8] M. Fujita, H. Goka, K. Yamada, and M. Matsuda. Competition between Charge- and Spin-Density-Wave Order and Superconductivity in $\text{La}_{1.875}\text{Ba}_{0.125-x}\text{Sr}_x\text{CuO}_4$. *Physical Review Letters*, 88:167008, 2002.
- [9] J. E. Hoffman, E. W. Hudson, K. M. Lang, V. Madhavan, H. Eisaki, S. Uchida, and J. C. Davis. A four unit cell periodic pattern of quasi-particle states surrounding vortex cores in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 295(5554):466–9, January 2002.
- [10] C. Howald, H. Eisaki, N Kaneko, and A. Kapitulnik. Coexistence of periodic modulation of quasiparticle states and superconductivity in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Proceedings of the National Academy of Sciences*, 100(17):9705–9, August 2003.
- [11] C. Howald, N. Kaneko, M. Greven, and A Kapitulnik. Periodic density-of-states modulations in superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Physical Review B*, 67:014533, 2003.
- [12] S. A. Kivelson, I. P. Bindloss, V. Oganessian, J. M. Tranquada, A. Kapitul-

- nik, and C. Howald. How to detect fluctuating stripes in the high-temperature superconductors. *Reviews of Modern Physics*, 75(4):1201–1241, October 2003.
- [13] Y. Kohsaka, C. Taylor, K. Fujita, A. R. Schmidt, C. Lupien, T. Hanaguri, M. Azuma, M. Takano, H. Eisaki, H. Takagi, S. Uchida, and J. C. Davis. An intrinsic bond-centered electronic glass with unidirectional domains in underdoped cuprates. *Science*, 315(5817):1380–5, March 2007.
- [14] Y. Kohsaka, C. Taylor, P. Wahl, A. R. Schmidt, Jinhwan Lee, K. Fujita, J. W. Alldredge, K. McElroy, Jinho Lee, H. Eisaki, S. Uchida, D.-H. Lee, and J. C. Davis. How Cooper pairs vanish approaching the Mott insulator in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 454(7208):1072–1078, August 2008.
- [15] M. J. Lawler, K. Fujita, Jinhwan Lee, A. R. Schmidt, Y. Kohsaka, Chung Koo Kim, H. Eisaki, S. Uchida, J. C. Davis, J. P. Sethna, and Eun-Ah Kim. Intra-unit-cell electronic nematicity of the high- T_c copper-oxide pseudogap states. *Nature*, 466(7304):347–351, July 2010.
- [16] K. McElroy, D.-H. Lee, J. Hoffman, K. Lang, J. Lee, E. Hudson, H. Eisaki, S. Uchida, and J. Davis. Coincidence of Checkerboard Charge Order and Antinodal State Decoherence in Strongly Underdoped Superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Physical Review Letters*, 94(19):197005, May 2005.
- [17] K. McElroy, Jinho Lee, J. A. Slezak, D.-H. Lee, H. Eisaki, S. Uchida, and J. C. Davis. Atomic-Scale Sources and Mechanism of Nanoscale Electronic Disorder in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 309(5737):1048–52, August 2005.

- [18] a. Mesaros, K. Fujita, H. Eisaki, S. Uchida, J. C. Davis, S. Sachdev, J. Zaanen, M. J. Lawler, and E.-a. Kim. Topological Defects Coupling Smectic Modulations to Intra-Unit-Cell Nematicity in Cuprates. *Science*, 333(6041):426–430, July 2011.
- [19] A. R. Moodenbaugh, Youwen Xu, M. Suenaga, T. J. Folkerts, and R. N. Shelton. Superconducting properties of $\text{La}_{2-x}\text{Ba}_x\text{CuO}_4$. *Physical Review B*, 38:4596, 1988.
- [20] S. H. Pan, E. W. Hudson, and J. C. Davis. [³He refrigerator based very low temperature scanning tunneling microscope. *Review of Scientific Instruments*, 70(2):1459, February 1999.
- [21] SH Pan. No Title, 1993.
- [22] Colin V. Parker, Pegor Aynajian, Eduardo H. da Silva Neto, Aakash Pushp, Shimpei Ono, Jinsheng Wen, Zhijun Xu, Genda Gu, and Ali Yazdani. Fluctuating stripes at the onset of the pseudogap in the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+x}$. *Nature*, 468(7324):677–680, December 2010.
- [23] Olga Perkovic, Karin Dahmen, and James P. Sethna. Avalanches, Barkhausen Noise, and Plain Old Criticality. *Physical Review Letters*, 75:4528, 1995.
- [24] B Phillabaum, E W Carlson, and K A Dahmen. Spatial Complexity Due to Incipient Electronic Nematicity in Cuprates. page (under review), 2011.
- [25] O. Pietzsch, a. Kubetzka, D. Haude, M. Bode, and R. Wiesendanger. A low-temperature ultrahigh vacuum scanning tunneling microscope with a split-coil

- magnet and a rotary motion stepper motor for high spatial resolution studies of surface magnetism. *Review of Scientific Instruments*, 71(2):424, 2000.
- [26] Dieter W. Pohl. Some design criteria in scanning tunneling microscopy. *IBM Journal of Research and Development*, 30(4):417–427, July 1986.
- [27] M Presland, J Tallon, R Buckley, R Liu, and N Flower. General trends in oxygen stoichiometry effects on T_c in Bi and Tl superconductors. *Physica C: Superconductivity*, 176(1-3):95–105, May 1991.
- [28] John Robertson, Steven Kivelson, Eduardo Fradkin, Alan Fang, and Aharon Kapitulnik. Distinguishing patterns of charge order: Stripes or checkerboards. *Physical Review B*, 74(13):1–10, October 2006.
- [29] Subir Sachdev. Colloquium: Order and quantum phase transitions in the cuprate superconductors. *Reviews of Modern Physics*, 75(3):913–932, July 2003.
- [30] Dietrich Stauffer and Amnon Aharony. *Introduction to Percolation Theory*. Taylor & Francis, Ltd., Philadelphia, PA, revised se edition, 1994.
- [31] E. Stoll. Resolution of the scanning tunnel microscope. *Surface Science Letters*, 143(2):L411–416, 1984.
- [32] Joseph A. Stroscio and William U. Kaiser. *Scanning Tunneling Microscopy*. Academic Press, Inc., San Diego, CA, 1993.
- [33] J. M. Tranquada, B. J. Sternlieb, J. D. Axe, Y. Nakamura, and S. Uchida. Evidence for Stripe Correlations of Spins and Holes in Copper Oxide Superconductors. *Nature*, 375:561–3, 1995.

-
- [34] Michael Vershinin, Shashank Misra, S. Ono, Y. Abe, Yoichi Ando, and Ali Yazdani. Local ordering in the pseudogap state of the high-T_c superconductor Bi₂Sr₂CaCu₂O_(8+delta). *Science*, 303(5666):1995–1998, March 2004.
- [35] Matthias Vojta. Lattice symmetry breaking in cuprate superconductors: stripes, nematics, and superconductivity. *Advances in Physics*, 58(6):699–820, November 2009.
- [36] Roland Wiesendanger. *Scanning Probe Microscopy and Spectroscopy*. Cambridge University Press, Cambridge, UK, 1994.
- [37] Roland Wiesendanger. Spin mapping at the nanoscale and atomic scale. *Reviews of Modern Physics*, 81(4):1495–1550, November 2009.
- [38] W. D. Wise, M. C. Boyer, Kamallesh Chatterjee, Takeshi Kondo, T. Takeuchi, H. Ikuta, Yayu Wang, and E. W. Hudson. Charge-density-wave origin of cuprate checkerboard visualized by scanning tunnelling microscopy. *Nature Physics*, 4(9):696–699, July 2008.
- [39] W. D. Wise, Kamallesh Chatterjee, M. C. Boyer, Takeshi Kondo, T. Takeuchi, H. Ikuta, Zhijun Xu, J. S. Wen, G. D. Gu, Yayu Wang, and E. W. Hudson. Imaging nanoscale Fermi-surface variations in an inhomogeneous superconductor. *Nature Physics*, 5(3):213–216, January 2009.
- [40] Guangyong Xu, G. D. Gu, M. Hücker, B. Fauqué, T. G. Perring, L. P. Regnault, and J. M. Tranquada. Testing the itinerancy of spin dynamics in superconducting Bi₂Sr₂CaCu₂O_{8+δ}. *Nature Physics*, 5(9):642–646, August 2009.

-
- [41] Ilija Zeljkovic, Elizabeth J Main, Tess L Williams, M C Boyer, Kamalesh Chatterjee, W D Wise, T Takeuchi, Hiroshi Ikuta, G D Gu, E W Hudson, and Jennifer E Hoffman. STM imaging of inversion-symmetry-breaking structural distortion in the Bi-based cuprate superconductors. *Arxiv Preprint*, page 1104.4342, 2011.