



Universiteit
Leiden
The Netherlands

Disorder and interactions in high-temperature superconductors

Sulangi, M.A.

Citation

Sulangi, M. A. (2018, July 5). *Disorder and interactions in high-temperature superconductors*. *Casimir PhD Series*. Retrieved from <https://hdl.handle.net/1887/63332>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/63332>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/63332> holds various files of this Leiden University dissertation.

Author: Sulangi, M.A.

Title: Disorder and interactions in high-temperature superconductors

Issue Date: 2018-07-05

Disorder and Interactions in High-Temperature Superconductors

Proefschrift

ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van Rector Magnificus prof. mr. C. J. J. M. Stolker,
volgens besluit van het College voor Promoties
te verdedigen op donderdag 5 juli 2018
klokke 15.00 uur

door

Miguel Antonio Sulangi

geboren te Manilla, de Filipijnen
in 1989

Promotor: Prof. dr. J. Zaanen

Promotiecommissie: Prof. dr. P. J. Hirschfeld (University of Florida)
Prof. dr. ir. H. T. C. Stoof (Universiteit Utrecht)
Prof. dr. J. Aarts
Dr. M. P. Allan
Prof. dr. ir. T. H. Oosterkamp
Prof. dr. K. E. Schalm

Casimir PhD series, Delft-Leiden 2018-19

ISBN 978-90-8593-348-9

An electronic version of this thesis can be found at

<https://openaccess.leidenuniv.nl>.

The research described in this thesis is supported by the Netherlands Organisation for Scientific Research (NWO/OCW) as part of the Frontiers of Nanoscience (NanoFront) program.

The cover is a photograph, taken by the author, of the mouth of the Puerto Princesa Underground River from inside the cave.

To Elmer, Rhodora, and Thea Sulangi.

CONTENTS

1	INTRODUCTION	3
2	PHENOMENOLOGY OF THE CUPRATES	15
2.1	Angle-Resolved Photoemission Spectroscopy	17
2.2	Scanning Tunneling Spectroscopy	20
2.3	Superconductor	21
2.4	Pseudogap	29
2.5	Strange Metal	33
3	REVISITING QUASIPARTICLE SCATTERING INTERFERENCE IN HIGH-TEMPERATURE SUPERCONDUCTORS: THE PROBLEM OF NARROW PEAKS	43
3.1	Introduction	43
3.2	Model and Methods	50
3.2.1	Green's Functions and the Local Density of States	52
3.2.2	Modeling the Measurement Process	55
3.3	Pointlike Scatterers	57
3.3.1	Single Weak Pointlike Impurity	59
3.3.2	Multiple Weak Pointlike Impurities	61
3.3.3	Multiple Unitary Pointlike Impurities	66
3.3.4	Dependence of the Power Spectrum on the Im- purity Strength	68
3.4	Smooth Disorder	70
3.4.1	Single Smooth Scatterer	73
3.4.2	Multiple Smooth Scatterers	75
3.4.3	Quantifying the Range of the Potential	78
3.5	Spatially Random On-Site Energies	81
3.6	Spatially Random Superconducting Gap	85
3.7	Discussion and Conclusion	89
3.A	Appendix: Single Unitary Pointlike Scatterer	93

Contents

4	QUASIPARTICLE DENSITY OF STATES, LOCALIZATION, AND DISTRIBUTED DISORDER IN THE CUPRATE SUPERCONDUCTORS	97
4.1	Introduction	97
4.2	Methods	102
4.2.1	Quasiparticle Density of States	103
4.2.2	Specific Heat	108
4.2.3	Localization Length	109
4.3	Models of Disorder	111
4.3.1	Random-Potential Disorder	111
4.3.2	Multiple Unitary Scatterers	112
4.3.3	Smooth Disorder	114
4.4	Quasiparticle Density of States: An Overview	118
4.5	Correlation Between the LDOS and the Disorder Potential	137
4.6	Properties of the Density of States near $E = 0$	141
4.7	Low-Temperature Specific Heat	148
4.8	Quasiparticle Localization	150
4.9	Discussion and Conclusion	160
5	SELF-ENERGIES AND QUASIPARTICLE SCATTERING INTERFERENCE	169
5.1	Introduction	169
5.2	Self-Energies and Broadening	174
5.3	Methods	178
5.4	Self-Energies in the Superconducting State	181
5.5	Self-Energies in the Normal State	199
5.6	Discussion and Conclusion	216
6	CONCLUSIONS AND OUTLOOK	221
6.1	The Unreasonable Effectiveness of QPI	221
6.2	Disorder: Old Dog, New Tricks	224
6.3	Stretching QPI to Its Breaking Point	227
	Bibliography	233
	Samenvatting	261
	Acknowledgements	265

List of Publications 269
Curriculum Vitae 271

INTRODUCTION

Ever since their discovery over thirty years ago [18], the cuprate high- T_c superconductors have provided a unceasing torrent of mysteries that have bedeviled generations of physicists. The initial astonishment at the very high transition temperatures found for these materials—several times larger than the largest T_c hitherto known—was soon augmented by bewilderment about the fundamental nature of these materials. The very presence of superconductivity in these materials was itself surprising: the ground state of the parent compounds of these materials is known to be an antiferromagnetic Mott insulator, which disappears fairly quickly upon hole-doping and from which superconductivity emerges. The close proximity of the superconducting and magnetically-ordered phases to each other was a clear hint that strong correlations among the electrons in these materials play a central role, as conventional wisdom suggests that magnetism and superconductivity arise from different electron-electron interactions and are therefore unlikely to be found near each other in materials. The necessity of a strongly-interacting description of these materials became clearer upon further exploration of their phase diagram. At temperatures close to zero, with increasing doping one encounters the antiferromagnetic Mott insulator, a d -wave superconductor, and, eventually, a normal phase which appears to be well-described by Fermi-

liquid theory. On the underdoped side there are various charge- and spin-ordered phases which appear to coexist—or compete—with the d -wave superconductivity.

When temperature is raised, a veritable Pandora’s box of problems is unleashed. It is in fact at this finite-temperature regime where much of the most persistently confusing aspects of the cuprates are found. At low dopings, antiferromagnetic order persists at finite T . The d -wave superconducting state at low doping transitions, at T_c , into a mysterious phase called the “pseudogap,” which, unlike a normal metal, features an extremely pronounced depletion of electronic excitations near the Fermi energy. The superconducting state reaches its highest T_c at a special value of the doping—“optimal doping”—and, at temperatures higher than T_c near optimal doping, the superconductor gives way to an extremely unusual metal whose anomalous properties defy any description in terms of Fermi-liquid theory. This “strange metal” is seen to have a linear-in- T resistivity, unlike that of a conventional metal, for which the resistivity scales with a higher power. The pseudogap in turn is found to cross over to the strange metal even at very high temperatures, with the crossover scale set by a doping-dependent temperature T^* . The strange metal occupies a fan-shaped segment of the doping-temperature phase diagram; on the overdoped side of this fan a return to more conventional Fermi-liquid-like behavior occurs.

A consistent and unified description of *all* of these phases of the cuprates is at the moment not known. Even the superconductor, perhaps the most well-understood phase of the cuprates, is still unusual. For one, it is more stable against disorder than Bardeen-Cooper-Schrieffer theory suggests. The d -wave symmetry of the order parameter means that Anderson’s theorem [12] does not apply to this situation, and quenched disorder should rapidly kill the superconductivity. However, the inhomogeneities present in the cuprate materials happen to

coexist with a stable d -wave superconducting state, thus suggesting that some decidedly beyond-BCS mechanism ensures its survival in this disordered setting. Another instance in which the cuprate superconductors deviate from the BCS expectation is in how the superfluid density ρ_s —which quantifies the stiffness of the superconducting condensate against twists in the phase of the order parameter—behaves. The underdoped cuprates have a very small superfluid density which is found to scale with T_c [166]. This suggests that fluctuations of the phase of the superconducting order parameter acquire an outsize importance in the underdoped cuprates, even inside the superconducting state, with the possibility that a “preformed-pairing” picture—*i.e.*, pairs form below T^* , but acquire phase coherence only upon reaching T_c —may explain the pseudogap phase above T_c [44]. This is in stark contrast to the BCS picture, where the characteristic temperature scale of the phase fluctuations is much larger than T_c , leading to the relative unimportance of these fluctuations within the superconducting state. Finally, angle-resolved photoemission spectroscopy (ARPES) experiments suggest that the gap does not close at T_c , as BCS theory predicts; instead, it *fills* [140]. T_c appears to be set by the temperature at which the gap and the quasiparticle scattering rate energy scales cross over into each other, and the temperature at which the gap ultimately closes is higher than T_c . Nevertheless, despite these non-BCS-like features, the d -wave superconducting state is known to host coherent excitations which, as seen in ARPES [82, 176] and scanning tunneling spectroscopy (STS) experiments [70, 112], behave exactly as d -wave Bogoliubov quasiparticles do.

The contrast between the comparatively well-understood superconducting state and the proximate strange metal near optimal doping cannot be any more stark. Because the strange metal is not a Fermi liquid, there is no sense in which the normal-superconducting transi-

tion is describable by anything resembling BCS theory. However, it is possible to understand many features of this anomalous state via the marginal Fermi liquid theory proposed by Varma *et al.* [173]. In this phenomenological model, the plain-vanilla Fermi liquid is augmented by a momentum-independent but frequency-dependent self-energy whose imaginary part depends as T when the temperature is greater than the frequency. This has rather drastic consequences: long-lived quasiparticles, the backbone of Fermi-liquid theory, cease to exist, as the quasiparticle weight $Z \rightarrow 0$ at the Fermi surface as $T \rightarrow 0$. While the microscopic origin of this behavior is not known, on an effective-field-theory level this succeeds in reproducing the strange transport anomalies present in the cuprates. In addition, ARPES experiments find that the strange metal features a very incoherent spectrum whose behavior could be reasonably fit into the marginal-Fermi-liquid description [2]. However, it remains an open problem how coherent quasiparticles in the superconducting state form upon moving from the strange metal, where the excitations are incoherent, and how the pseudogap transitions into both the strange metal and the superconductor.

The landscape of the phases of the cuprates, as outlined above, is rich, complex, and, four decades on, still incredibly confusing. Despite this rather daunting state of affairs, this thesis will try to describe portions of the phase diagram of these materials. Given the immense difficulty of constructing a global theory of the cuprates, a more bare-bones and phenomenologically-minded approach can help illuminate which bits of physics are important and which are only of secondary importance. Much insight can be derived by considering fairly simple but well-understood models of these phases which can nevertheless be augmented by bells and whistles that account for deviations from the models we started out with. The marginal Fermi liquid theory of the

strange metal is perhaps the paradigmatic example of this approach: the Fermi-liquid starting point of the model is weakly coupled, but the addition of the self-energy incorporates the nontrivial effects of interactions, leading to the destruction of long-lived quasiparticles in the theory. Also, the evidence from spectroscopy suggesting that coherent quasiparticles are present in the superconductor allows considerable leeway in treating the superconducting state as a mean-field BCS superconductor with d -wave symmetry.

It is in this spirit that this thesis will examine both the superconducting and the normal state of the cuprates. A dominant theme underlying the work presented here is the nontrivial effect of disorder on various electronic properties of these materials. Disorder plays two distinct, almost antithetical roles in the cuprates, but it is often taken for granted how interrelated these two are. First, in the limit of weak disorder, it acts as a probe of the underlying electronic structure of these materials. In a normal metal, the presence of an impurity induces Friedel oscillations—which are simply modulations in the local density of states—whose spatial structure reveals details about the Fermi surface [31, 159, 71, 137]. The situation in the cuprates is no different. These Friedel oscillations have been observed in the superconducting state of the cuprates using STS, leading to the phenomenon dubbed “quasi-particle scattering interference” (QPI) [70, 112, 61, 90, 50]. The modulations found in the real-space differential conductance maps from STS can be Fourier-transformed, revealing a rich set of dispersive peaks in the power spectrum whose behavior can be used to reconstruct the band-structure details of the cuprates [182, 25]. Importantly, the quantum-mechanical-wave interference underlying this phenomenon illustrates how the Bogoliubov quasiparticles are well-defined and coherent excitations [189]. Second, under some circumstances it generates low-energy electronic states in the superconducting state, in the

process irrevocably altering the electronic spectrum of the clean case [68, 100]. The presence of a finite density of states at the Fermi energy deep in the superconducting phase has long been known throughout the cuprate family from specific heat experiments, and standard lore has it that these are generated by disorder [115, 116, 144]. In addition, doping these materials by zinc—a strong local scatterer—leads to impurity resonances being generated near the Fermi energy, seen vividly also by STS [129, 16, 17]. .

Given that we know with definiteness that the cuprates are macroscopically disordered materials, it becomes imperative to consider disorder both as a probe of electronic structure *and* as an origin of low-energy excitations seen in specific heat. The situation is complicated even further by subtleties present in the very nature of disorder in the cuprates. In particular, the copper-oxide planes—which host the physics of most relevance to experiments—are clean. Aside from rare defects which are thought to be Cu vacancies, no strong impurities are seen within the copper-oxide planes using STS, due to the strong copper-oxygen bonding present. It is possible to induce these strong impurities via chemical substitution of zinc or nickel atoms, but these would necessarily result in resonances near the Fermi energy which are not seen in cuprates without Zn or Ni dopants. What appears to be the case instead is that dopants located in the insulating layers adjacent to the copper-oxide planes are the source of disorder—but unlike the aforementioned Zn or Ni dopants, which act as pointlike scatterers, these would generate a smoother and longer-ranged disorder potential which affects the electrons in the copper-oxide planes [2, 126, 124, 125, 161]. However, unlike pointlike forms of disorder, these smoother disorder potentials are far less amenable to analytical treatment, and are accessible only with large-scale numerical methods.

It cannot be emphasized enough that the aspects of disorder considered in this thesis remain central to some staggeringly persistent mysteries about the cuprates. One example of this is “QPI extinction” [90, 50]. STS experiments on underdoped cuprates in the superconducting state, across a fairly wide doping range, show the usual signatures of QPI—modulations in the differential conductance maps, prominent peaks in the power spectrum—right until the bias voltage is such that the tips of the contours of constant energy cross the antiferromagnetic zone boundary—*i.e.*, the portion of the Brillouin zone enclosed by the four lines connecting $(0, \pm\pi)$ and $(\pm\pi, 0)$. Beyond that voltage, most of the dispersing peaks seen in the power spectrum vanish, and what do remain are seen not to disperse as the voltage is further changed. Why this happens is not presently known, and numerous explanations using a mean-field free-fermion description—*e.g.*, unusual impurities [176], coexisting spin-density wave order [11]—appear unsatisfactory, or are simply infeasible, due to the lack of supporting experimental evidence for them. If one takes the results at face value, this phenomenon is a remarkably lucid demonstration of the breakdown of the quasiparticle description as one nears the antinodes—the famous “nodal-antinodal dichotomy” in action—but to bolster that interpretation, it first has to be understood what precisely is happening at the point where these peaks are extinguished, and here STS finds itself in disagreement with results from ARPES, for reasons that are not completely understood. Many aspects of this problem remain unclear even to this day.

Another example of this is a recent set of specific heat measurements on underdoped cuprates in the presence of magnetic fields [144]. The cuprate samples used in these experiments are some of the cleanest known—far more orderly than most other members of the cuprate family. Despite this, a residual linear-in-temperature term in the specific

heat at zero field is seen in the data, which is indicative of a nonzero density of states at the Fermi energy. It is not trivial to attribute this simply to disorder, as the coefficient of this T -linear term—which is proportional to the density of states at $E = 0$ —is larger than that seen in ostensibly dirtier cuprates. It has almost been taken for granted since the early years of the cuprates that disorder generates these low-lying excitations detected in specific heat. So, what is happening here? There is as yet no definitive answer, as the cuprates in question are not especially amenable to surface probes such as ARPES and STS, and thus one cannot examine the copper-oxide planes directly to see what the underlying source of the low-energy electronic excitations is. (NMR is a local probe which could in principle measure various local quantities directly, but one has to note the fact that a magnetic field is present in these experiments, and thus the decidedly nontrivial effects of this field on the superconducting state have to be carefully taken into account.) One idea that has been proposed is intra-unit-cell loop-current order which coexists with the d -wave superconducting state [21, 4, 88, 181], but this idea appears to run into difficulties when one tries to constrain the parameters of the model from the observed data. Disorder due to dopants within the buffer layers may be another way out—and as a matter of fact this is discussed in great detail in this thesis—as the dopants present in these cuprates could remain away from the copper-oxide planes while still disordering the electrons within the plane via a random screened Coulomb potential [162]. However, as we have only indirect means of probing disorder and any putative coexisting order, it is safe to say that there is still no firm resolution to this problem.

One admittedly heuristic takeaway from these two examples is that whenever disorder is involved in the cuprates, there is more than meets the eye. In the first case, QPI is Friedel oscillations, but it also is much

more than that. For one, the system is macroscopically disordered: there is no single isolated impurity, but rather the disorder—wherever it may come from—is of a distributed nature. Second, the tunneling process between the STM tip and the copper-oxide plane is highly nontrivial [109, 93]. Third, despite these two considerations, the peaks seen in experiment are somehow well-defined—almost miraculously so. In the second case, specific heat experiments provide a very beautiful probe of *all* low-energy excitations, but are blind to the precise mechanism through which these excitations are generated. In addition, in the case of the cuprates, what constitutes “disorder” can be ambiguous, thanks to the complex layered structure of these materials. The dopants reside off the plane, but for some cuprates the question of *where* exactly these reside is difficult to resolve. There are also non-stoichiometric alterations to the structure of the compounds, unrelated to hole doping, which induce further off-plane disorder [42]. All of these complicating factors mean that disorder is not a simple matter at all. There is only so much one can understand about its impact on the electronic properties of the cuprates without taking into account all these tricky caveats. One thus sees the need to thoroughly revisit disorder and to see the extent to which its effects manifest themselves on experimentally-measurable quantities.

These considerations motivate Chapters 3 and 4 of this thesis, which reexamine respectively the imprint of disorder on the local density of states (as seen by STS) and the quasiparticle density of states at the Fermi energy (as seen by specific heat experiments) in the *d*-wave superconducting state. For the first case, the vast majority of prior theoretical work on QPI has centered on the case of a single isolated pointlike impurity [182, 25, 125, 176, 93]. The extent to which the experimental evidence for QPI can be reproduced using distributed models of disorder is in general not clear. For the second case, the prevail-

ing explanation for the finite DOS at the Fermi energy in the cuprates is the so-called “dirty d -wave” theory, which implicitly assumes both a dirty copper-oxide plane and a local, pointlike model of disorder [54, 68, 100, 39, 14, 13], and as such fails to describe the realistic situation in which the copper-oxide planes are clean and a smooth disorder potential sourced by off-plane dopants is present.

Chapter 2 contains a review of the various phenomena seen in ARPES and STS, in addition to a brief overview of the basics of these two methods. We consider three phases of the cuprates individually—the superconductor, the pseudogap, and the strange metal—and exhaustively describe much of their known phenomenology. We will also discuss in detail some of the theories that have been put forth to explain these various phenomena, and point out regimes where these proposed theories fail.

In Chapter 3 we revisit quasiparticle interference in the cuprates and try to see if the strikingly sharp peaks seen in the experimental power spectra can be reproduced by an exhaustive array of models of distributed disorder—examples include an ensemble of weak pointlike scatterers, random chemical potential disorder, and smooth disorder—and the incorporation of a model of the STM tunneling process. What is found is that weak pointlike disorder and random chemical potential disorder best reproduce both the real-space and the Fourier-transformed spectra seen in experiment; that smooth disorder fails to fully reproduce the experimental power spectrum, as large-momentum scattering is suppressed in that particular case; and that the peaks in experiment are *sharper* than any of our simulations see, in a surprising reversal of what we usually expect.

Chapter 4 is devoted to a very thorough examination of various kinds of disorder and their impact on the quasiparticle density of states near the Fermi energy as the amount of disorder is increased.

While strong pointlike scatterers and random chemical-potential disorder do lead to a finite DOS, agreement with experiment is achieved only when the amount of disorder is unphysically large. Meanwhile, disorder due to off-plane dopants is found to lead to a realistic value of the DOS at the Fermi energy while leaving much of the d -wave state in intermediate and higher energies largely unaffected, and in addition sharp resonances *at* the Fermi energy are seen when the concentration of these smooth scatterers is very large. The localization length is also studied for various models of disorder, finding that for all three models of disorder the quasiparticles at the Fermi energy are localized, and that the dependence of the localization length on energy depends sensitively on the type of disorder present.

Chapter 5 spotlights QPI once more, this time focusing on the effects of self-energies on the LDOS power spectrum, with the advantage of knowing, from Chapter 2, that the single weak pointlike scatterer does lead to a phenomenologically accurate power spectrum. After discussing the effect of self-energies on the spectral function and the DOS, we proceed to analyze cases of interest to the cuprates. In the superconducting case, we study the “gap-closing/filling” phenomenology seen in ARPES experiments and attempt to analyze the extent to which STS measurements can also see this phenomenon, and contrast these with the BCS case, wherein the gap closes but does not simultaneously fill. The peaks seen in the superconducting power spectra are found to be rather sensitive to the amount of broadening present, with the peaks smearing and becoming incoherent at large self-energies. We also study the normal-state LDOS power spectra, assuming that the strange metal is well-described by a marginal Fermi liquid, and find that a key signature of this state is the presence of broad caustics in the power spectrum describing scattering wavevectors between points along the Fermi surface. The main difference between the marginal

INTRODUCTION

Fermi liquid and the ordinary Fermi liquid is found to be simply in the amount of broadening present in these caustics: the LDOS power spectra of a marginal Fermi liquid has much more broadening than that of the ordinary Fermi liquid.

Finally, a summary of our results is shown in Chapter 6, along with a lengthy discussion of potential future directions, both experimental and theoretical, in the study of the cuprates.

PHENOMENOLOGY OF THE CUPRATES

In this chapter we provide a fairly extensive summary of the basic facts known about the electronic excitations of the cuprate high-temperature superconductors. As this field is driven primarily by experiment, this chapter will feature mainly experimental results. A particular emphasis is placed on angle-resolved photoemission spectroscopy and scanning tunneling spectroscopy measurements, as these two experimental probes have been responsible for much of what we know about the momentum- and real-space structure of the electronic excitations in the cuprates. These experiments in fact provide much of the impetus for the theoretical work described in this thesis. Some discussion on the theories used to account for these experimental results is also included. Because of the vast amount of research performed using either probe, we will highlight only a fairly small number of results which illustrate how the cuprates deviate from and challenge both the BCS and Fermi-liquid paradigms [84]. It should be noted that the full phase diagram of the cuprates is very complex—by way of illustration, a phase diagram largely agreed upon by the high- T_c community is shown in Fig. 2.1—and we caution the reader right away that this review will not do justice to the remarkably diverse array of phenomena seen in the cuprate family.




Figure 2.1: Phase diagram of the copper-oxide high-temperature superconductors. The x - and y -axes correspond to the hole-doping level and the temperature, respectively. The antiferromagnetic Mott-insulating state (blue region, labeled “AF”) at low dopings transitions into d -wave superconductivity (green region, labeled “d-SC”) when hole-doping is increased. The pseudogap (yellow region) and strange metal (pink region) both appear at higher temperatures, with the onset of the pseudogap marked by the temperature T^* . The areas with green and red stripes show where spin-density-wave order and charge-density-wave order, respectively, have been detected. The dashed green and red lines demarcate where fluctuations corresponding to spin and charge order, respectively, first become apparent. Reprinted from Ref. [84].

We first provide a “theorist’s introduction” to ARPES and STS—more specifically, we discuss how these experiments are performed and what the quantities measured by either experiment are. Finally the numerous insights from either experiment are discussed, in order of increasing inscrutability: the d -wave superconductor, the pseudogap, and, finally, the strange metal.

2.1 ANGLE-RESOLVED PHOTOEMISSION SPECTROSCOPY




Figure 2.2: The layered, quasi-two-dimensional crystal structure of the cuprate high-temperature superconductors. The metallic CuO₂ planes are separated by insulating layers. The $d_{x^2-y^2}$ copper orbitals hybridize with the p_x and p_y oxygen orbitals, giving rise to the square-lattice structure of the CuO₂ planes. Reprinted from Ref. [84].

2.1 ANGLE-RESOLVED PHOTOEMISSION SPECTROSCOPY

Angle-resolved photoemission spectroscopy (ARPES) is a particularly revealing probe of the electronic structure of the cuprates. In a nutshell, this method takes advantage of the photoelectric effect to allow a direct look at the dispersion of the electronic excitations inside the cuprates. Much of what we now know about the cuprates—*e.g.*, the d -wave nature of the superconducting order parameter, the presence of so-called “Fermi arcs” inside the pseudogap, and marginal-Fermi-liquid-like behavior in the strange metal—can be traced back to pioneering ARPES experiments on a variety of cuprate materials. Perhaps the best-studied of these materials is Bi₂Sr₂CaCu₂O_{8+ δ} (Bi-2212), owing to the fact that it cleaves easily between layers and thus allows the physics occurring within its copper-oxide planes to be probed directly. The copper-oxide superconductors are known to have a quasi-two-dimensional layered structure, with the metallic CuO₂ planes sandwiched between insulating buffer layers—for an illustration, see Fig. 2.2.

Most phenomena of interest occur directly within the CuO_2 planes. ARPES is a particularly apt probe for understanding these phenomena, as it works best when used to study two-dimensional electron systems.

ARPES owes its existence to the photoelectric effect—the famous phenomenon wherein light incident on a material imparts energy to an electron, allowing it to escape [65, 41]. The quantum nature of light implies that the energy of a single photon is hf . Upon absorption of this energy, the electron can be dislodged from the material with kinetic energy $E_k = h\nu - \phi - |E_b|$, where ϕ is the work function of the surface of the material and E_b is the binding energy inside the solid. The absolute value of the momentum of the electron can in turn be calculated from the measured kinetic energy as $p = \sqrt{2mE_k}$, where m is the mass of the electron, and because the emission angles can be measured, the components of \mathbf{p} can also be obtained as well.

An ARPES experiment measures a quantity $I(\mathbf{k}, \omega)$, called the photoemission intensity, which on a crude level is simply the combined probability that an electron is excited by the photon; that the electron travels to the surface; and that the electron is finally liberated from the surface. (\mathbf{k} and ω here are the momentum parallel to the surface and the energy, respectively, of the electron.) The second and third steps in this process are surface-dependent, while the first step is sensitive to the electronic structure of the material, and thus contributes electronic-structure-dependent contributions to $I(\mathbf{k}, \omega)$. A discussion of how $I(\mathbf{k}, \omega)$ is calculated from the relevant transition probabilities is subtle and is discussed in thorough detail in a number of reviews [34, 33]. For our purposes it suffices to say that in the *sudden approximation*—in which the liberated electron does not interact with what

remains of the material upon escaping—the photoemission intensity $I(\mathbf{k}, \omega)$ can be written in the following way:

$$I(\mathbf{k}, \omega, T) = I_0(\mathbf{k}, \nu, \mathbf{A})f(\omega, T)A(\mathbf{k}, \omega) \otimes R(\delta\mathbf{k}, \delta\omega) \quad (2.1)$$

In this expression I_0 is proportional to matrix elements associated with the photon-absorption process; $f(\omega, T)$ is the Fermi function, given by

$$f(\omega, T) = \frac{1}{e^{\frac{\omega}{k_B T}} + 1}; \quad (2.2)$$

and $A(\mathbf{k}, \omega)$ is the spectral function, which is defined as

$$A(\mathbf{k}, \omega) = -\frac{1}{\pi} \text{Im} G(\mathbf{k}, \omega), \quad (2.3)$$

where $G(\mathbf{k}, \omega)$ is the translationally-invariant many-body retarded Green's function. $I(\mathbf{k}, \omega, T)$ is simply the product of these three factors convolved with $R(\delta\mathbf{k}, \delta\omega)$, which is a function describing the experimental resolution available. The Fermi function means that ARPES probes only the occupied states at temperature T . The main object of interest is $A(\mathbf{k}, \omega)$, which is simply the density of electronic excitations at energy ω and momentum \mathbf{k} and as such reveals much about the momentum-space structure of the electronic excitations of these materials.

In the ARPES literature, it is common to speak of “energy-distribution curves” (EDCs) and “momentum-distribution curves” (MDCs). EDCs are simply plots of the spectral function with binding energy at a fixed \mathbf{k} (for example, a momentum *at* the Fermi surface). MDCs on the other hand show the spectral function along a line in momentum space (for instance, along $k_x = k_y$) while holding the binding energy fixed.

2.2 SCANNING TUNNELING SPECTROSCOPY

Scanning tunneling spectroscopy (STS) is particularly advantageous as a probe for the cuprates because it enables the direct *real-space* visualization of the electronic structure of these materials, and because, unlike ARPES, both states below and above the Fermi level are accessible. In addition, it is also possible to examine the momentum-space details of these materials by employing the Fourier transform. A diverse panoply of phenomena has been visualized using STS such as inhomogeneous gaps, quasiparticle scattering interference, and static stripe phases—all phenomena whose real-space structure would have been less accessible to most other conventional probes. Like ARPES, STS is particularly optimized for layered two-dimensional systems such as the cuprates (see Fig. 2.2) and has studied Bi-2212 extensively thanks to the ease with which it can be cleaved.

STS relies on tunneling of electrons from a scanning tunneling microscope (STM) to probe the real-space structure of materials. An STM has a metallic tip which is put in proximity to the surface of the material of interest. A potential difference V is then applied between the tip and the material, and a tunneling current I is generated, the main quantity measured by these experiments. Assuming that the density of states of the metal in the tip is approximately constant, one can arrive at the following expression for I [51]:

$$I(\mathbf{r}, V) = m(\mathbf{r}) \int_0^{eV} \rho(\mathbf{r}, E) dE. \quad (2.4)$$

Here m is a position-dependent matrix element and $\rho(\mathbf{r}, E)$ is the local density of states at position \mathbf{r} and energy E . In terms of the many-

body retarded Green's function $G(\mathbf{r}, \omega)$ —here written in a real-space basis— $\rho(\mathbf{r}, E)$ is simply given by

$$\rho(\mathbf{r}, E) = -\frac{1}{\pi} \text{Im} G(\mathbf{r}, E). \quad (2.5)$$

Note that this definition is almost exactly the same as that for the spectral function $A(\mathbf{k}, \omega)$ in Eq. 2.3—only this time, instead of momentum space, one deals with real space instead.

At this moment the LDOS is hidden within the integral, but it can be obtained by taking the derivative of I with respect to V —the differential conductance g :

$$g(\mathbf{r}, E) = dI/dV|_{E=eV} \propto \rho(\mathbf{r}, E). \quad (2.6)$$

In real systems, however, the proportionality seen in the above expression is muddied by factors intrinsic to the experimental setup. To eliminate these factors, occasionally “Z-maps” are used instead. Here the proportionality factors are removed by taking the ratio of differential conductances taken at positive and negative bias voltages:

$$Z(\mathbf{r}, E) = \frac{g(\mathbf{r}, E)}{g(\mathbf{r}, -E)} = \frac{\rho(\mathbf{r}, E)}{\rho(\mathbf{r}, -E)}. \quad (2.7)$$

In any case, because STS probes the real-space density of states, it is particularly useful for visualizing phenomena arising from the breaking of translation symmetry due to, say, disorder or coexisting order.

2.3 SUPERCONDUCTOR

As we mentioned in the introduction, the superconducting state of the cuprates is the most well-understood of the many phases of these ma-




Figure 2.3: Left: Plot of the large Fermi surface seen in the normal state of the cuprates. The first Brillouin zone is shown. Because of the square-lattice structure of the copper-oxide planes, the BZ is a square. Shown here are the locations of the “nodes” and “antinodes.” Right: Plot of the absolute value of the d -wave gap function (thick blue line) along the Fermi surface (dashed red line) in the upper right-hand quadrant of the first Brillouin zone. The gap vanishes at the nodes and is largest at the antinodes.

terials. However, many aspects of this state remain unusual, which is not surprising as the phases to which it is proximate are even stranger. To begin with, the superconductor is an *unconventional* one, due to its d -wave pairing symmetry: the order parameter undergoes a sign change upon rotations by $\pi/2$. (For comparison, in a conventional s -wave superconductor, such as that predicted by BCS theory, the order parameter has the same symmetries as the underlying lattice.) The unusual symmetry of the order parameter can be seen in the momentum-space form of the gap function, which can be expressed as follows:

$$\Delta(\mathbf{k}) = 2\Delta_0(\cos k_x - \cos k_y). \quad (2.8)$$

The gap vanishes along the $k_x = k_y$ and $k_x = -k_y$ lines, and has its maximum absolute value near $(0, \pm\pi)$ and $(\pm\pi, 0)$. (This is il-

lustrated in Fig. 2.3.) This form of the order parameter implies that gapless quasiparticles exist at the “nodes,” which are the four points where the Fermi surface intersects the two lines along which the gap vanishes. At these points there are zero-energy quasiparticles with a linear Dirac-like dispersion at low energies. Already this implies that the thermodynamic signatures of the d -wave superconducting state are very different from those of an s -wave one, as in the latter case the quasiparticle spectrum is fully gapped and thus does not feature any low-energy excitations that can be seen in thermodynamic probes such as the specific heat. On the other hand, the quasiparticles near the antinodes—the regions in the vicinity of $(0, \pm\pi)$ and $(\pm\pi, 0)$ —are maximally gapped.

Nowadays the d -wave nature of the order parameter is a firmly established fact about the cuprates, but it is telling that in the early days of high- T_c superconductivity, the precise nature of the symmetry was a hotly debated topic. Here ARPES provides an unambiguous answer which has been confirmed again and again with increasing instrument precision. How would one detect this order-parameter symmetry? The dispersion of d -wave Bogoliubov quasiparticles is such that the excitations near the nodes live at the Fermi energy, while those at the antinodes are gapped. From measurements of spectral function within the nodal and antinodal regions, it was seen that the nodal spectrum shows no gap, while EDCs taken near the antinodes show a gap—the peaks of the EDCs are shifted relative to the Fermi level, suggesting the formation of a gap [157, 36]. Furthermore, ARPES finds that these superconducting quasiparticles are well-defined excitations—their peaks in the spectral function are very easily discerned [82, 46, 163, 102, 176]. One surprising aspect of these quasiparticles is that these become sharp as temperature is lowered past T_c . The normal-state spectrum features far less sharpness and no coherent quasiparticles can be seen

in the EDCs near optimal doping [82]. The precise mechanism underlying the manner in which sharp quasiparticles form below T_c is not known.

This picture, in which Bogoliubov quasiparticles with a d -wave dispersion propagate as coherent excitations within the superconductor, was bolstered by a number of complementary results obtained from STS. The first such result was the observation of very prominent resonances near the Fermi energy in zinc-doped Bi-2212 [129]. Zinc substitutes for copper within the copper-oxide planes, creating a very strong local scattering center. Such resonance states close to the Fermi level are consistent with theoretical predictions for d -wave superconductors featuring strong unitary scatterers [16, 17].

The second and perhaps far more consequential result is the observation of quasiparticle scattering interference (QPI) in the cuprates [70, 112, 61, 90, 50]. As mentioned in the introduction, differential conductance maps taken on the cuprates reveal energy-dependent modulations which are incommensurate. Taking the Fourier transform of these dI/dV maps shows an array of well-defined peaks whose position in “ \mathbf{q} -space” changes as bias voltage is altered. This suggests that these peaks do not originate from static charge or spin order, but arise instead from Friedel oscillations due to disorder intrinsic to the cuprates. But why peaks? It was realized that because the cuprates are d -wave superconductors, the scattering processes that give rise to these LDOS modulations are strongly influenced by the very unusual dispersion of d -wave Bogoliubov quasiparticles. When the energy is shifted away from the Fermi level, the contours of constant energy (CCEs) acquire a banana-like shape. Scattering occurs from a state lying on these contours to another, and when two points on these CCEs have a large joint density of states between them, the scattering wavevector connecting these has a strong intensity in the power spec-

trum of the differential conductance map. As it happens, any pair of the tips of these “bananas” has a large joint DOS, and it was seen that the peaks in the experimental power spectrum correspond perfectly with the scattering wavevectors describing tip-to-tip scattering. This is the simplest picture of the physics underlying the phenomenological “octet model” used to analyze differential conductance data from STS [182, 25].

QPI is important for two reasons. First, it acts as a momentum-space probe, allowing one to obtain information about the Fermi surface and the band structure of the cuprates. By tracking the position of the peaks in \mathbf{q} -space as a function of energy, the underlying band structure and momentum-dependent behavior of the Bogoliubov quasiparticles can be reconstructed. The remarkably sharp peaks and their particular dependence on energy confirm the d -wave nature of the superconducting state. Second, it confirms one key aspect of the superconductor which was already seen in ARPES: that the quasiparticles deep inside the superconducting state are coherent, well-defined excitations [189]. On a heuristic level, QPI can be understood simply as the interference of the quantum-mechanical waves corresponding to the Bogoliubov quasiparticles as they encounter quenched disorder. This description necessitates the coherence of these excitations, for otherwise they cannot propagate long enough to interfere with each other and produce modulations in the LDOS.

Having mentioned all the aspects in which the superconducting state of the cuprates behaves similarly to a d -wave BCS superconductor as seen by ARPES and STS, we now turn to some of its anomalous features. The first of these is the observation from STS experiments that the underdoped superconductor is quite spatially inhomogeneous [130, 95, 111, 9], inspiring the metaphor of “quantum mayonnaise” to describe the microscopic phase separation appearing in these materials

[188]. To be more specific, STS experiments suggest that two energy scales are at play here. Below the first, lower energy scale, the electronic structure is by and large spatially featureless, but above that scale there is an onset of heterogeneous features that are prominent at small hole doping. A second, higher energy scale is seen from the tunneling spectra, and the extracted values of the gap associated with this higher scale vary in space, forming domains at which a single gap value dominates. The disorder in the gap has been shown to be correlated with the positions of the off-plane dopants, and there is good reason to suspect that the latter causes the former, although the precise reason for this remains to be seen.

The second is the mysterious and hotly contested phenomenon of “QPI extinction” [90, 50], to which we had already alluded in the introduction. STS experiments on underdoped cuprates observe that many of the octet-model QPI peaks suddenly disappear once the bias voltage is raised past the point where the tips of the “bananas” intersect with the antiferromagnetic zone boundary—that is, the diagonal lines connecting the four points $(0, \pm\pi)$ and $(\pm\pi, 0)$. The octet-model peaks that *do* remain suddenly become dispersionless, with their positions in \mathbf{q} -space not varying appreciably once the bias voltage is increased further. This, in conjunction with the earlier observation of spatial inhomogeneity in the underdoped cuprates, has led to the interpretation that two classes of excitations are present—one class being delocalized, freely propagating excitations corresponding to the low-energy Bogoliubov quasiparticles, and another class being localized excitations that become more prominent as hole-doping decreases, and which are associated with the pseudogap phase emerging at higher temperatures. As will be clearer in the discussion on the pseudogap, this behavior well within the superconducting phase is also seen in the pseudogap,

and the results suggest that these high-energy antinodal excitations associated with the pseudogap do indeed persist below T_c .

The reason that this phenomenon remains the subject of much debate a decade after its discovery is due to how it directly conflicts with ARPES results. As mentioned earlier, ARPES sees coherent quasiparticles across the *entire* Fermi surface in the d -wave superconducting state, even at the antinodal regions [82, 46, 163, 102, 176]. According to ARPES, incoherent antinodal quasiparticles are characteristic only of above- T_c phases—the strange metal and the pseudogap—whereas no such “nodal-antinodal dichotomy” appears to be seen deep in the superconducting state. A curious fact also is that QPI extinction is seen even at moderate overdoping ($p \approx 0.19$), where any possible magnetic correlations due to the antiferromagnetic Mott insulator should be minimal at best. A number of proposals have been made to reconcile these two wildly different results. One line of reasoning argues that the QPI peaks are sensitive to the nature of disorder causing it, and that a proper accounting of the precise momentum-dependence of the T -matrix due to general forms of disorder could partially explain the extinction of the peaks [176]. While plausible, this does not appear to explain the onset of dispersionless peaks at higher energies, and it does not convincingly explain why some of the peaks are suddenly quenched at *that* particular energy. Another proposal puts forth that spin-density wave order coexisting with the d -wave superconductor can explain the partial extinction of these QPI peak [11]. In a nutshell, SDW order reconstructs the CCEs; thus, at the point where the tips of these “bananas” cross these lines, the CCEs undergo a change of topology, with a “banana” and its mirror joining together to form a closed pocket and leading to the diminishing of some of the octet-model peak intensities. However, no signatures of static or slowly fluctuating spin order have been detected in Bi-2212, making this explanation highly

limited. In any case, these STS results seem to suggest that physics beyond a mean-field-like d -wave superconductor plays a role in the cuprates as hole-doping is decreased, and that the superconducting and pseudogap phases are inextricably linked to each other.

The final anomalous observation in the superconducting state that we will discuss at length is the “filling” of the superconducting gap as temperature is increased towards T_c . This comes by way of fairly recent ARPES experiments on Bi-2212 over a wide doping range [141, 140, 138, 139]. It was found using near-nodal measurements of ARPES spectra that as T is increased towards T_c , the superconducting gap Δ_0 decreases, but at T_c , the gap is still nonzero—the gap closes at a higher temperature. In parallel with this, the quasiparticle scattering rate Γ rapidly increases as T_c is approached. It appears that T_c is set by the temperature at which the plots of Δ_0 and Γ as a function of temperature cross each other, with T_c being found to be near the point where $\Delta_0 \approx 3\Gamma$ —the origin of the numerical factor 3 is not understood. This gap-filling phenomenology is seen throughout a wide range of hole dopings, and is in stark contrast to what one expects from BCS theory, according to which the gap should fully close *at* T_c . These results are suggestive of the possibility that pairs indeed form at some temperature $T_p > T_c$, but with phase coherence of these pairs inhibited by the presence of strong pair-breaking at high temperatures (quantified by the quasiparticle scattering rate Γ) [44]. In this picture it is only when these pairs become sufficiently long-lived that they acquire phase coherence at T_c . The observed crossover of the two scales Δ_0 and Γ near T_c lends experimental support to the idea that phase fluctuations play an important role in the physics of the superconductor and the pseudogap, with preformed pairs existing above T_c which contribute to superconductivity only upon becoming phase-coherent as temperature is lowered past T_c .

2.4 PSEUDOGAP

The pseudogap is perhaps the most complex phase of the cuprates, mainly for the sheer number of phenomena present—coexisting stripe order, Fermi arcs, and superconducting fluctuations—whose relationships with each other are not clear or understood with any certainty, and a definition that encompasses the phase in all its complexity is elusive. Contributing to the confusion surrounding this phase is the lack of any certainty as to whether the pseudogap can be understood via a conventional mean-field theory, or whether a very different, possibly exotic paradigm is necessary. A generally accepted, if rather anodyne, definition of the pseudogap regime is the following: it is the phase above T_c from the underdoped superconducting state which is characterized by a prominent suppression of the electronic spectral weight in the vicinity of the Fermi energy [164, 121].

Even this definition fails to encompass the highly unusual way in which this suppression of the DOS is organized in momentum space. The pseudogap can be best understood by looking at ARPES spectra across the Fermi surface, as one of the key aspects of this state is the rather severe degree to which the spectra seen in momentum space differ from what one would expect for a d -wave superconductor and a normal metal. In the pseudogap, the spectral weight at and near the antinodal regions show a pronounced gap. The common procedure is to symmetrize the EDCs, under the rather plausible condition of particle-hole symmetry, and what one sees from symmetrized spectra is that there are two peaks in the antinodal spectra located some distance away from the Fermi energy. These peaks in the antinodal EDCs of the pseudogap are unlike those of the d -wave superconductor in that they are relatively smoother and more suppressed in intensity. Once one moves from the antinodes to the nodes along the Fermi sur-

face, what one finds is that the gap shrinks and disappears suddenly at some point near the nodes, signaling the onset of “Fermi arcs”—finite sections of momentum space where electronic excitations at the Fermi energy can be found [119, 83].

It has to be emphasized that this behavior deviates very strongly from that of either a d -wave superconductor or a Fermi liquid. In the d -wave superconducting state below T_c , the symmetrized antinodal EDCs show sharp peaks about the Fermi energy, which get closer to each other as one moves towards the nodes, remaining well-defined until they collapse into a single peak at the node (where the gap is zero). For a Fermi liquid, the Fermi surface separates the occupied states from unoccupied ones and as a matter of principle is necessarily a closed manifold—it cannot have endpoints!

The Fermi arcs are a particularly tricky challenge for theorists to explain. A set of explanations has centered around the possibility that Fermi-surface reconstruction due to coexisting density-wave order is responsible for these arcs. In this scenario the large hole-like Fermi surface becomes replaced by a set of smaller pockets—but with the caveat that these pockets still remain closed. If one takes this seriously as an explanation, the Fermi arcs can only come from one side of these putative pockets [27]. It has been argued from models with coexisting density wave order that coherence factors could be responsible for the absence of spectral weight on the other, “invisible” side of the pocket, but no trace of this pocket has been seen in experiments to date.

A second explanation is that these Fermi arcs are simply d -wave nodes that are broadened by a large scattering rate [123, 120, 28]. In the pseudogap regime, a wide range of evidence has accumulated suggesting that the quasiparticle scattering rate is in fact fairly large in the pseudogap regime, and that pairing exists well above T_c . The origin of this large temperature-dependent scattering rate is not fully under-

stood, but once it is sufficiently large the d -wave gap starts to be filled in, generating a nonzero density of states at the Fermi energy. Because the d -wave gap is smallest near the nodes, the near-nodal region quickly fills as the scattering rate is increased, and the gaps seen in near-nodal EDCs disappear. Under this scenario the gapless region identified by ARPES is simply due to the induced low-energy states that arise from a large scattering rate. It is in fact not difficult to see how symmetrized EDC analyses may have misidentified a broadened d -wave node as a Fermi arc, as the gap is defined by the distance in energy from one peak to its mirror image across the Fermi energy, and increasing broadening smoothens out these peaks near the nodes to the point of incoherence once the scattering rate is large (*e.g.*, of the same order of magnitude as the superconducting gap) [175]. This explanation is consistent with the picture of the pseudogap as a phase-disordered d -wave superconductor [44], with pairs existing at high temperatures (the pseudogap) which then become phase-coherent below T_c .

Nevertheless, it appears that this preformed-pairing picture does not fully account for a plethora of other observations about the nature of the gaps in the pseudogap as a function of momentum and temperature. ARPES experiments see that the gap at the nodes has a different temperature dependence from that at the antinodes. Near the nodes, the gap shrinks fairly rapidly as temperature is raised, and while it remains finite at T_c it fully closes at a temperature not far off from T_c . In contrast, it appears that the gap near the antinodes shrinks with increasing temperature far more slowly: the antinodal gap is by and large unchanged as T_c is crossed, and only shrinks appreciably upon nearing a much higher temperature scale T^* [64, 175]. An instance of this is data on UD92 Bi-2212, which has $T_c = 92$ K. For these samples the near-nodal gap closes at $T \approx 97$ K, but the antinodal gap

goes to zero only when $T \approx 190$ K. It appears, on the face of these experimental results, that some strange sort of phase separation—but in momentum space—occurs for the electronic excitations of the pseudogap.

STS provides evidence supporting the phase-fluctuation picture, and in addition gives additional insight into the energy scales at play in this phase. In deeply underdoped cuprates, it was found that the octet-model peaks characteristic of QPI remain at temperatures above T_c —in fact they appear to persist to temperatures as high as $1.5T_c$ [99]. This suggests that in the pseudogap, d -wave pairing is still present. It is rather striking that the peaks do not appear to be sensitive to T_c ; these evolve smoothly as T is increased past T_c . The second is the observation, already seen in the superconducting phase, that at high energies, some of these octet-model peaks are suddenly quenched and replaced by nondispersive modulations. This result, if taken at face value, suggests that the excitations in the pseudogap living in the antinodes do not contribute to the scattering processes giving rise to QPI, and the lack of any dispersiveness is a sign that these are localized, as opposed to extended, states. Importantly, it is found that the energy at which some of these peaks disappear happens to coincide with the energy where the spatially homogeneous nature of the material is lost and where the inhomogeneities present in the “gap maps” become much more prominent. More to the point, low-gap regions—where the size of the gap is below the QPI extinction energy—exhibit sharp coherence peaks characteristic of the superconductor, while high-gap regions show gap-like features but do not have any prominent coherence peaks and have more of a pseudogap-like character. Finally, the presence of these high-energy nondispersive modulations is highly suggestive of charge order, and at these high energies STS finds signatures of broken spatial symmetries [96].

Taken together, these suggest that the states truly characteristic of the pseudogap—as opposed to the remnant Bogoliubov quasiparticles of the d -wave superconductor—are spatially localized, reside at a higher energy scale, and are associated with the antinodes. It is not altogether clear however how this real-space phase separation of the pseudogap-like and superconductor-like excitations relates to the momentum-space phase separation seen separately in ARPES and STS. In addition, many questions about the pseudogap remain. Two versions of the phase diagram of the cuprates circulate: one has the pseudogap crossover line at T^* entering the superconducting dome, terminating at $T \rightarrow 0$ near optimal doping. The other phase diagram features the pseudogap crossover line intersecting with the termination point of the superconducting dome as $T \rightarrow 0$. ARPES generally finds that a gap still persists above T_c even at optimal doping, supporting the latter picture [177, 141, 140]. Nevertheless this has been the subject of much debate, and a final resolution is still not within sight.

2.5 STRANGE METAL

We now come to the strange metal, which remains, without any doubt, the most perplexing of all the phases of the cuprates. It was recognized soon after the discovery of high- T_c superconductivity in these materials that the transport properties of the normal state are highly anomalous, at least as understood within the framework of Fermi-liquid theory. Perhaps the foremost marker of this anomalous strange metal is the behavior of the resistivity ρ as a function of temperature. For a normal metal described by Fermi-liquid theory, $\rho \sim T^2$. Instead what is seen in the cuprates is that $\rho \sim T$ [59]. In addition, this linear-in- T behavior persists up to very high temperatures, in striking contrast to the

behavior seen in a normal metal, for which the resistivity should saturate at such large temperatures. Other unusual transport signatures of the strange metal include the following: a mostly featureless (*i.e.*, temperature- and frequency-independent) Raman scattering intensity [160]; a constant thermal conductivity $\kappa(T)$ [55]; and a nuclear relaxation rate $1/T_1$ which has a temperature-independent component such that $1/T_1 \sim \alpha T + \beta$ (a normal metal would only have the T -linear part in the nuclear relaxation rate) [183, 180].

Despite these mysterious transport properties which hint at the fundamentally non-Fermi-liquid character of the normal state, it was recognized that one could formulate, under reasonable assumptions, an entirely phenomenological theory of this phase of matter [173]. Such a theory was developed early on by Varma and coworkers and was dubbed the “marginal Fermi liquid”—“marginal” for reasons we will explain in a short while. The basic assumption underlying the MFL is that the ordinary Fermi liquid is coupled to some set of excitations whose existence is taken as a given, and whose contribution to the density fluctuation spectrum has the following form:

$$\text{Im } \chi(\mathbf{q}, \omega, T) \sim \begin{cases} -\omega/T & \text{if } |\omega| < T \\ -\text{sgn}(\omega) & \text{if } |\omega| > T. \end{cases} \quad (2.9)$$

Note that $\text{Im } \chi$ is assumed to be momentum-independent. It can be shown from Eq. 2.9 that the self-energy Σ becomes

$$\Sigma(\omega, T) = \lambda \left(\omega \ln \frac{x}{\omega_c} - i \frac{\pi}{2} x \right). \quad (2.10)$$

Here $x = \max(|\omega|, T)$ —note that this could be represented by $x = \sqrt{\omega^2 + \pi^2 T^2}$, for ease of computation— ω_c is a cutoff frequency, and λ is a coupling constant.

Eq. 2.10, despite its rather compact form, contains a tremendous amount of information. First, the single-particle scattering rate, which is proportional to $\text{Im}\Sigma$, is momentum-independent, implying that the scattering rate as inferred from transport measurements such as optical conductivity should be the same as the single-particle scattering rate [1]. Second, the single-particle scattering rate is proportional to x , rather than to x^2 (which is the case for an ordinary Fermi liquid). Third, the quasiparticle weight Z , which is defined as

$$Z = \left(1 - \frac{\partial \text{Re}\Sigma}{\partial \omega}\right)^{-1}, \quad (2.11)$$

goes to zero logarithmically as $\omega \rightarrow 0$ (that is, as one scales towards the Fermi surface) at $T = 0$. This means that quasiparticles *do not* exist even at $T = 0$ for a marginal Fermi liquid. Fourth, $\text{Im}\Sigma$ is linear in ω at $T = 0$. This linearity implies that the quasiparticle width does not vanish faster than ω —a necessary criterion for the existence of quasiparticles—and $\text{Im}\Sigma \propto \omega$ in fact is the highest power for which the quasiparticle picture fails. Thus, $\text{Im}\Sigma \sim \omega$ is a “marginal” case. The logarithmic singularity in Z^{-1} is in fact the weakest such singularity possible.

How are transport measurements explained by this MFL self-energy? Much of the transport phenomenology is easy to explain because of the aforementioned momentum-independence of the self-energy, which leads to the equality of the single-particle scattering rate—which can actually be measured in ARPES—to the transport scattering rate. The linear-in- T resistivity can be explained by noting that $\rho = \frac{\Gamma_t}{\omega_p^2}$, where Γ_t is the transport scattering rate and ω_p is the plasma frequency. According to Eq. 2.10 $\Gamma_t \sim T$ at $\omega = 0$; this thus implies that $\rho \sim T$. The

constant thermal conductivity? The Wiedemann-Franz law implies that

$$\kappa(T) \propto T\sigma(T), \quad (2.12)$$

and, recalling that $\sigma(T) = 1/\rho(T) \sim 1/T$, leads to $\kappa(T) \sim \text{const.}$ For Raman scattering, the form of Eq. 2.9 directly leads to a featureless signal. Other transport anomalies can similarly be accounted for by Eqs. 2.9 and 2.10.

ARPES measurements taken in the strange-metal phase of the cuprates find considerable support for the MFL hypothesis. From Eqs. 2.3 and 2.10, a number of predictions could be made from MFL theory for momentum-distribution curves. (To remind the reader, MDCs are simply linecuts of the spectral function along a direction in momentum space at fixed frequency.) As the MFL self-energy is momentum-independent, the MDC profiles along cuts perpendicular to the Fermi surface should be of Lorentzian form. In particular its full width at half maximum should be $-\text{Im}\Sigma(\mathbf{k}_{\text{cut}}, \omega)$, where \mathbf{k}_{cut} are momenta along the chosen cut in momentum space. This implies that when one has $\omega \rightarrow 0$ and momenta along any cut perpendicular to the Fermi surface, the FWHM of the MDC should be proportional to T . Conversely, at fixed temperature, the MDC widths should scale linearly with x , with $x \approx \omega$.

These expectations were confirmed rather spectacularly by ARPES results from the Brookhaven and Argonne groups [169, 2, 168, 81]. However it was observed that the self-energies were anisotropic: the MDCs along antinodal directions were found to be broader than those taken along the nodal ones. Another feature was that the frequency- and temperature-dependence of the widths was found to be largely uniform across the Fermi surface, while the offset characterizing the momentum-space anisotropy in the self-energy was found to be frequency-

and temperature-independent. These groups primarily used MDC analysis to obtain fits of the extracted self-energy to the following form:

$$\text{Im } \Sigma(\mathbf{k}, \omega, T) = \Gamma(\mathbf{k}) - \lambda \frac{\pi}{2} x. \quad (2.13)$$

The momentum-dependent term $\Gamma(\mathbf{k})$ is free of any frequency- and temperature-dependence, and is largest at the antinodes and smallest near the nodal region. Its temperature-independence has allowed its identification as an elastic scattering rate, with the highly anisotropic form argued to arise from small-angle scattering from impurities located between the copper-oxide planes. Importantly, the fits taken from the MDCs data were also found to describe the EDCs reasonably well, with the antinodal EDCs being much broader than the nodal ones.

Before moving on to other aspects of the strange metal, a few things should be noted. First, the antinodal EDCs at and near optimal doping are often so incoherent that a peak is not discernable [82]. These should be contrasted with the antinodal EDCs in the normal state of the overdoped cuprates, which are generally seen to be fairly coherent, and with EDCs across the entire Fermi surface in the superconducting phase, which exhibit sharp quasiparticle peaks. As mentioned earlier, how these quasiparticles acquire coherence and become well-defined as temperature is lowered past T_c is still unsettled. Second, while peaks in the MDCs may be suggestive of quasiparticles, it is only when looking at EDCs that the truly non-quasiparticle nature of the strange metal becomes apparent. The MFL self-energy leads to the generation of a considerable amount of spectral weight away from the Fermi energy even at $T = 0$. This broadened spectral weight is seen clearly in EDCs even at low energy resolutions, but is however something to which an MDC analysis (for which the frequency is

held constant) would *not* see unless many such analyses are performed at different frequencies in order to ascertain the precise momentum-dependence of the widths. In fact this observed incoherence in the antinodal EDCs—in conjunction with the sharpness of the MDCs—has been put forth as evidence for spin-charge separation in the normal state, with a dimensional crossover from quasi-one-dimensional, Luttinger-liquid-like physics above T_c to two-dimensional physics below T_c proposed to occur [127, 26]. Third, while the differences between single-particle properties of the marginal Fermi liquid and those of the ordinary Fermi liquid are very sharp at $T = 0$ —*e.g.*, the non-analyticity of the self-energy, the broad spectral function EDCs, and the absence of quasiparticles—these differences become blurred at finite temperature, and the effects are subtle enough that one needs to take special care in attempting to distinguish these two phases from each other. These differences require that quantities such as the MDC widths be measured across a wide range of temperatures and frequencies in order to obtain the correct scaling.

One can only go so far with the marginal Fermi liquid, however. The microscopic origin of the strange metal is not definitively settled and remains an area of active research. Nevertheless one popular paradigm which has been used to explain the strange metal is quantum criticality [158, 146]. At $T = 0$, one can speak of distinct quantum phases which are accessed by manipulating some tuning parameter P . A quantum phase transition is simply a continuous transition between two proximate quantum phases occurring at a special value P_0 . A generic feature of quantum critical points is that, at finite temperature and frequency, correlation functions have the scaling form ω/T —a feature already built into the marginal Fermi liquid, as attested to by Eq. 2.9.

Further evidence for quantum criticality comes courtesy of the observation by Homes *et al.* that T_c is proportional to the product of

the superfluid density ρ_s and the normal-state resistivity $\rho(T_c)$ for a vast array of cuprates with varying doping levels and crystal structure [72]. By a simple argument using mainly dimensional analysis, it can be shown that this seemingly universal relation—dubbed “Homes’ law”—implies that the characteristic dissipation time in the strange metal appears to *saturate* the lowest possible bound for thermal fluctuations set by the uncertainty principle [190]. That is, in the strange metal the relaxation time $\tau \approx \frac{\hbar}{k_B T}$, suggesting that in this finite-temperature regime the strange metal appears to be controlled by the physics of a putative critical point located close to optimal doping.

Despite compelling evidence pointing to a quantum-critical origin of the strange metal, the nature of this quantum critical point—if it exists—remains shrouded in mystery, perhaps almost literally so by the superconducting dome. One version of the cuprate phase diagram has it that the pseudogap crossover line, set by T^* , terminates well within the superconducting dome. The quantum-critical point—which in this picture would lead to the pseudogap as a finite-temperature phase crossing over to the strange metal—is thus hidden by the superconducting state, perhaps suggesting that the $T = 0$ quantum-critical point is unstable to perturbations that lead superconducting order to develop. However, it is not known which of the many orders characterizing the pseudogap is responsible for this $T = 0$ critical point. One way of perhaps probing the physics of the strange metal is by quenching the superconductivity by means of an applied magnetic field. However, what is seen in experiments is a more conventional Fermi-liquid-like state with a reconstructed Fermi surface, instead of the large hole-like one seen in ARPES [37, 154, 153]. The relationship between this Fermi-liquid-like state in a magnetic field and the strange metal is not clear.

A theoretical understanding of the strange metal is complicated by the fact that there appears to be no way of imposing theoretical control over the physics of a Fermi liquid coupled to quantum-critical fluctuations [101, 113]. On an even more philosophical level, one is left questioning whether a Fermi liquid is even a valid starting point for the construction of a theory of the strange metal, considering the presence of strong correlations which preclude any means of understanding these highly anomalous finite-temperature regimes perturbatively. Unfortunately there are not many alternative theoretical paradigms available. The best-understood non-Fermi liquid is the Luttinger liquid, which is the theory of an interacting electron gas in one spatial dimension [60]. The one-dimensional nature of the problem allows exact solutions to be obtained by means of bosonization [165, 105, 110]. In addition, the Luttinger liquid is a quantum-critical *phase*—it needs no fine-tuning to reach criticality—and spin-charge separation occurs, with collective excitations forming which carry separately the charge and spin degrees of freedom. It has been shown that the ARPES and STS spectra of the Luttinger liquid behave very differently from that of a weakly-coupled Fermi liquid: two sets of dispersing features are present in these spectra which correspond to the spin- and charge-carrying excitations—not the underlying electrons—and which propagate at two different velocities, demonstrating explicitly spin-charge separation [87]. These are all specific to one dimension, of course. Extensions to two and three dimensions have been cooked up by coupling Luttinger liquids together along one or two transverse directions, but these still rely on the physics of the foundational one-dimensional electron gas from which these non-Fermi liquids are made [179, 43, 117]. Nevertheless, the Luttinger liquid remains an interesting metaphor for the higher-dimensional non-Fermi liquids of relevance to the cuprates, insofar as its very existence shows how collective excitations radically different

from the underlying electrons could emerge naturally from strong interactions, and a higher-dimensional generalization, if found, could prove to be of much use in explaining the strange metal.

Other recent approaches have included the use of holographic duality to describe strongly-interacting finite-density phases of matter [191, 63]. This method relies on the existence of a duality between a conformal field theory in d spacetime dimensions and a gravitational theory living in $d + 1$ -dimensional anti-de Sitter space [108, 185, 56]. As this is a weak-strong duality, difficult quantum field theory problems are mapped onto relatively tractable *classical* gravitational problems—under some special circumstances. This has been exploited to extract insights that would otherwise have been very difficult to obtain using a more conventional perturbative approach. One example of this is the use of holography to study finite-density fermionic systems that appear to be quantum-critical phases, such as “semi-local quantum liquids”—states of matter with an infinite correlation time but a finite correlation length—which are seen to emerge fairly naturally from holography [79, 78]. These mimic the features of the marginal Fermi liquid, which has manifest “local quantum criticality”—that is, the correlation functions have no momentum-dependence, but have a nontrivial dependence on the frequency. Similarly, insights from holography have led to an understanding that the anomalous transport properties in the strange metal could be explained by hydrodynamic considerations. Simply put, quantum-critical states of matter without quasiparticles at finite density have a small viscosity, and dissipation of momentum due to quenched disorder can be shown to directly lead to a resistivity that is linear in temperature [35]. These insights are far removed from the realm of well-defined quasiparticles; the nonperturbative nature of holographic methods allows these non-quasiparticle-like excitations to be tractably handled.

Complementary to this is the recent explosion of work on sign-free quantum Monte Carlo simulations of strongly interacting fermionic systems. One recent study takes as its starting point a simple lattice Hamiltonian with Ising degrees of freedom which can be tuned through a nematic phase transition [97]. Despite the simple features of the model, the Hamiltonian was found to lead to a surprisingly rich phase diagram. First, a superconducting dome forms above the quantum-critical point, with the point at which T_c is highest almost coincident with the Ising nematic critical point. Second, non-Fermi-liquid behavior is seen in the vicinity of the quantum-critical point—the spectral function becomes broader at certain portions of the Fermi surface. Perhaps more strikingly, there is an especially pronounced “nodal-antinodal dichotomy”—the imaginary part of the self-energy is much broader at the antinodes than at the nodes. It is surprising to see this feature emerge from an admittedly simple model of interacting electrons.

These two newfangled approaches highlight the necessity of fresh perspectives and stripped-down but nonperturbative models to illuminate the still-befuddling nature of the strange metal, and show their potential to lead to unexpected insights into even the much more well-understood neighboring phases. While these caricature the real cuprates to a rather severe degree, these can only help in providing paths towards understanding where a traditional perturbative approach fails.

REVISITING QUASIPARTICLE SCATTERING
INTERFERENCE IN HIGH-TEMPERATURE
SUPERCONDUCTORS: THE PROBLEM OF NARROW
PEAKS

3.1 INTRODUCTION

Scanning tunneling spectroscopy (STS) has matured into one of the most powerful techniques for studying complex electron systems. It has been most successful in the study of high- T_c superconductors, where it has revealed a spectacular array of new phenomena to be present in the cuprates [150]. Prominent examples of such phenomena include ordering in the pseudogap [174, 62, 89, 96], inhomogeneities in the superconducting gap and pseudogap [95, 45, 111], and quasiparticle interference (QPI) [70, 112].

Here we wish to revisit the interpretation of the QPI phenomenon. This was first observed in the cuprates when STS measurements done on superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ found that spatial modulations in the local density of states (LDOS) were present in the real-space maps. A particular category of these modulations is found to be *incommensurate* and, more importantly, *dispersive*—that is, the wavevector peaks in the Fourier power spectrum corresponding to these modula-

tions are found to be energy-dependent [70, 112, 90]. In the underdoped regime, these coexist with peaks which are non-dispersing and are attributed to the presence of “stripy” charge-density-wave order [74, 73] or an electronic glass [89]. In a remarkable advance, these were explained in a series of papers laying out the theory as understood for a single pointlike scatterer [70, 182, 25]. In essence, the effect can be understood in terms of interference fringes associated with the coherent Bogoliubov quasiparticles of the d -wave superconductor, which behave like quantum-mechanical waves that diffract in the presence of quenched disorder [189]. Given their quasi-relativistic dispersion, this scattering is strongly enhanced at wavevectors associated with the extrema of the dispersions at a given energy. This is illustrated in Figs. 3.1 and 3.2. With increasing energy, the contours of constant energy (CCEs) of the Bogoliubov excitations in momentum space change shape (Fig. 3.1). The scattering is strongly enhanced at the tips of the banana-shaped contours (Fig. 3.2), defining an octet of characteristic momenta. Upon Fourier-transforming the real-space STS maps, one finds peaks at these momenta, which disperse as function of energy (Figs. 3.3 and 3.4). This forms a set of data that allows one to reconstruct the dispersion relations of the Bogoliubov quasiparticles. These are strikingly consistent with results from ARPES, where these single-particle dispersions are measured directly in momentum space. It is beyond doubt that this “octet model” interpretation is correct for the cuprates, especially as additional evidence for QPI has also been obtained from $\text{Ca}_{2-x}\text{Na}_x\text{CuO}_2\text{Cl}_2$ [61]. The effect has also been observed in iron-based superconductors [8, 5, 6] and heavy-fermion materials [98, 151, 7, 171]. The success of the octet model has spurred a considerable amount of theoretical work on the signatures of QPI in related states of matter such as the pseudogap phase of the cuprates [133, 134, 114, 19], as well as in systems without a gap, such as graphite




Figure 3.1: Contours of constant energy for a d -wave superconductor for different energies E , in units where $t = 1$. Observe that energies from $E = 0.050$ to $E = 0.300$ feature closed, banana-shaped CCEs, while for higher energies such as $E = 0.350$ the CCE changes topology and becomes open.

[20] and the surface states of three-dimensional topological insulators [49, 145, 58]. The ubiquity of QPI in gapless systems is not surprising, as its signatures were in fact first imaged in conventional metals [31, 159, 71, 137].

The octet model is simply a kinematical picture describing the scattering of quasiparticles in the presence of disorder. It is another matter to explain how well-defined patterns of QPIs can arise under realistic conditions. This was intensely studied theoretically, at first starting from models describing d -wave fermions scattering from a single isolated impurity potential [182, 25, 196, 125, 176, 93]. In Section 3.3, we will reproduce a typical result involving a single point scatterer. One infers from the results that there is an overall similarity between these theoretical results and the experimental data. However, even on a qualitative level it is not completely satisfactory. In our numerically obtained Fourier-space maps, the “peaks” are actually associated with intensity enhancements of intersecting diffuse streaks and blurry regions. In contrast, the experimental QPI signals are remarkably well-defined *peaks*.




Figure 3.2: The octet model in k -space. Shown are the seven wavevectors connecting one tip of a “banana” to another when $E = 0.200$. Dashed arrows denote wavevectors connecting states where the superconducting gap has the same sign, while undashed ones connect states where the gap changes sign.

A caveat is that microscopic details do matter when taking into account the actual measurement process involved in STS experiments. This was anticipated early on by the observation that the mismatch between the s -wave orbital emanating from the tunneling tip and the microscopic $d_{x^2-y^2}$ copper-centred orbitals in the perovskite planes implies that the tunneling current enters the *nearest neighbors* of the copper site over which the tip is positioned [109]. This “fork mechanism” was recently confirmed by an impressive first-principles model of the tunneling process [93]. We will study the effects of this “fork” on the QPIs in Section 3.3. We will find that this is actually only a minor concern for the overall interpretation. Kreisel *et al.* also find that modifications coming from a realistic description of the tunneling process have the potential to resolve the apparent paradox that we will demonstrate. We will come back to this issue at the end of this chapter.

The serious problem with the pointlike scatterer model lies in its inconsistency with the actual chemistry of the cuprates. Pointlike impurities are naturally explained in terms of substitutional defects in the cuprate planes. However the CuO_2 planes are well-established to be




Figure 3.3: Locations of the special \mathbf{q}_i wavevectors in extended \mathbf{q} -space. The energy is $E = 0.200$, same as in Fig. 3.2. The octet model predicts that peaks in the Fourier-transformed LDOS will be present at these locations. A square demarcating the boundary of the first Brillouin zone (*i.e.*, $-\pi \leq q_x, q_y \leq \pi$) is shown. Note that certain wavevectors (in this particular case, \mathbf{q}_4 and \mathbf{q}_5) may extend beyond the first Brillouin zone. In our lattice simulations these peaks will be folded back into the first Brillouin zone.

very clean with regard to their stoichiometry. In fact, zinc and nickel can be substituted for copper in the CuO_2 planes. Since such chemical defects correspond to strong potentials, this gives rise to a major modification of the electronic structure at the impurity core. This is indeed seen in STS, as the zinc impurities show up very prominently in the LDOS maps of zinc-doped BSCCO [129, 17]. The details of these core states were in fact instrumental in identifying the “fork” mechanism [109, 93]. Nickel impurities were found to be similarly visible in the case of nickel-doped BSCCO, the difference in this case being that nickel impurities are magnetic scatterers [77]. On the other hand, the STS spectra of pristine cuprates do not show any of these localized impurity states.

Instead, it appears that disorder in the cuprates should be of a more distributed and smooth kind. Doping occurs away from the CuO_2 planes. These are charged impurities, and given the poor screening along the c -axis, one then expects smooth, Coulombic disorder, simi-

lar to what is realized in modulation doping of semiconductors [130]. Such off-plane dopants have indeed been imaged in STS experiments on BSCCO [111]. Similarly, dopants might modulate the tilting patterns in the CuO_2 planes, resulting in a similar form of distributed disorder [42]. This involves *inherently many-impurity* effects that are not easy to study using the standard single-impurity T -matrix method. We note that multiple pointlike impurities have indeed been considered before in the literature [196, 25, 14]. However, the most general many-impurity problem is technically very demanding, especially when one tries to consider forms of disorder other than point impurities, or when one tries to scale up the system size.

Given these difficulties, we take advantage of an alternative numerical method to directly compute the LDOS, inspired by methods heavily in use in the quantum transport community. This is outlined in Section 3.2. Our point of departure is a tight-binding Hamiltonian on a square lattice describing a d -wave superconductor. Instead of diagonalizing this real-space Hamiltonian, we compute the Green's function directly by *inverting* the Hamiltonian, which can be done efficiently, and from the Green's function we obtain the LDOS. Superconducting gap functions and even full self-energies can be straightforwardly incorporated. Any form of spatial inhomogeneities can be modeled efficiently using this method, and our system sizes can be made very large—for instance, LDOS maps of systems with size 1000×120 , which we use, can be obtained in a matter of minutes—the better to approach the same large field of view as current experiments have. We originally aimed to use this to study more complex phenomena such as the gap inhomogeneities (“quantum mayonnaise”) found in the pseudogap regime, as well as the effects of the electronic self-energies on STS results [32]. However, we found out that issues arise already on the most funda-

mental level of the theory of QPI deep in the superconducting state of the cuprates, which is the subject of this chapter.

Using this method, we can address any conceivable form of spatial disorder and study its effects on the QPI spectra. We set the stage in Section 3.3, focusing on the case of a single weak pointlike impurity. We then insert a large number of such weak pointlike impurities at random positions and examine QPI with and without the filter effect. We then examine in detail the related case where many *unitary* scatterers are present. We next turn our attention to a single Coulombic impurity and subsequently to a densely distributed random ensemble of such smooth scatterers. Although the real-space patterns appear to be suggestively similar to the stripe-like textures seen in experiment, this runs into a very serious problem: the peaks in the power spectra involving large momenta disappear very rapidly, and this holds even if the range of the potential is shortened. We consider then the case of a random on-site potential, similar to Anderson's model of disorder. Although the effects of quasiparticle scattering interference can indeed be seen in the real-space and Fourier-transformed maps, this form of disorder results in power spectra which show considerable fuzziness, in contrast to the well-defined peaks seen in experiment. We end by considering a simple model of superconducting gap disorder. Although this works quite well for the simplified case we consider, the problem is that, for more realistic smooth gap inhomogeneities, large-momenta peaks will be suppressed.

By eyeballing the numerous plots present in this chapter, the reader may already have convinced himself or herself that there is a serious problem with the standard explanation of QPIs. By making the model of disorder more and more realistic, the correspondence with experiment deteriorates. As we will discuss in the final section, it is an

interesting open challenge to explain the sharpness of the QPI peaks as seen in STS measurements.

3.2 MODEL AND METHODS

Two important requirements in theoretically reproducing results from STS experiments are large system sizes and the ability to model general forms of inhomogeneities. Modern STS experiments feature a very large field of view, which allows large-scale inhomogeneities present in materials to be visualized. Replicating this large field of view numerically is a challenge because simulations with large system sizes require sizable amounts of computational effort. Most numerical work on disordered high-temperature superconductors has centered around two methods: the T -matrix method and exact diagonalization. The T -matrix approach has the advantage of being exact for the case of pointlike impurities and requires minimal numerical effort, even for large system sizes, but is restricted in its applicability—smooth potential scatterers, for instance, are not accessible in this formalism. On the other hand, exact diagonalization allows any form of disorder to be modeled, but at the expense of being restricted to relatively small system sizes.

In this chapter we utilize a method—a novel one as far as its application to both disordered d -wave superconductors and the modeling of STS experiments is concerned—that is formally *exact*, allows any form of disorder to be modeled, gives access to very large system sizes, and is computationally efficient. In addition, since it is based on Green's functions, it is straightforward to include the effects of self-energies; this will be the subject of Chapter 5 of this thesis. Before introducing the method, we will first discuss the lattice model of the cuprates that

we will use in this chapter. Our starting point is the following tight-binding Hamiltonian for a d -wave superconductor on a square lattice:

$$H = \sum_{\langle i,j \rangle} \left[- \sum_{\sigma} t_{ij} c_{i\sigma}^{\dagger} c_{j\sigma} + \Delta_{ij} c_{i\uparrow}^{\dagger} c_{j\downarrow}^{\dagger} + \Delta_{ij}^* c_{i\uparrow} c_{j\downarrow} \right]. \quad (3.1)$$

We include nearest-neighbor and next-nearest-neighbor hopping (specified by the amplitudes t and t' , respectively) and a chemical potential μ . d -wave pairing is incorporated by ensuring that the gap function has the form $\Delta_{ij} = \pm\Delta_0$, where (i, j) are two nearest-neighbor sites and the positive and negative values of Δ_{ij} are chosen for pairs of sites along the x - and y -directions, respectively. This is a mean-field Hamiltonian for the d -wave superconducting state of the cuprates. We set the lattice spacing $a = 1$ and the nearest-neighbor hopping $t = 1$ —*i.e.*, we will thus measure all energies in units of t .

In the clean limit, the Hamiltonian can be diagonalized by going to momentum space. The quasiparticle energies are given by

$$E(\mathbf{k}) = \sqrt{\epsilon_k^2 + \Delta_k^2}, \quad (3.2)$$

where

$$\epsilon_k = -2t(\cos k_x + \cos k_y) - 4t' \cos k_x \cos k_y - \mu \quad (3.3)$$

and

$$\Delta_k = 2\Delta_0(\cos k_x - \cos k_y). \quad (3.4)$$

Eq. 3.2 describes the dispersion of the Bogoliubov quasiparticles of a d -wave superconductor. At $E = 0$ there are four points in momentum space at which zero-energy excitations exist. For the purposes of our calculations we take the band-structure and pairing parameters relative to $t = 1$ as $t' = -0.3$, $\mu = -0.8$, and $\Delta_0 = 0.08$ throughout this chapter. We note that while these band-structure parameters cover

hoppings only up to the next-nearest-neighbor level, we selected them to be close to the phenomenological values obtained by Norman *et al.* for optimally-doped BSCCO [122]. Our results will turn out not to depend sensitively on band-structure details.

3.2.1 *Green's Functions and the Local Density of States*

The central quantity of interest in our study is the local density of states (LDOS) of a superconductor in the presence of disorder. The LDOS at position \mathbf{r} and energy E can be expressed as

$$\rho(\mathbf{r}, E) = -\frac{1}{\pi} \text{Im} G_{11}(\mathbf{r}, \mathbf{r}, E + i0^+), \quad (3.5)$$

where G is simply the full Green's function corresponding to H in Nambu space, given by

$$G = (\omega\mathbb{1} - H)^{-1}, \quad (3.6)$$

and G_{11} is the particle Green's function. One can observe from Eq. 4.8 that to obtain the LDOS we do not need all the elements of G —the bare LDOS can be obtained from just the diagonal elements of G . (Note however that when we will come to include nontrivial tunneling processes, more elements of G will be needed; this will be described in detail in the next subsection.) Here we do not determine the gap function self-consistently.

We proceed by noting that H , in a real-space basis, can be written as a block tridiagonal matrix—*without any approximations*—when periodic boundary conditions are imposed along the y -direction and open

boundary conditions are placed along the x -direction. H exhibits the following structure:

$$\mathbf{H} = \begin{pmatrix} \mathbf{a}_1 & \mathbf{b}_1 & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{b}_1^\dagger & \mathbf{a}_2 & \mathbf{b}_2 & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{b}_2^\dagger & \mathbf{a}_3 & \mathbf{b}_3 & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{b}_3^\dagger & \mathbf{a}_4 & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \mathbf{b}_{N_x-2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{b}_{N_x-2}^\dagger & \mathbf{a}_{N_x-1} & \mathbf{b}_{N_x-1} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{b}_{N_x-1}^\dagger & \mathbf{a}_{N_x} \end{pmatrix}. \quad (3.7)$$

N_x and N_y denote the number of sites in the x - and y -directions, respectively. \mathbf{a}_i is a $2N_y \times 2N_y$ block containing all hoppings, pairings, and on-site energies along the y -direction at the i th column. \mathbf{b}_i meanwhile is a $2N_y \times 2N_y$ block that contains hopping and pairing terms along the x -direction between the i th and $(i+1)$ th columns.

By construction the inverse Green's function $G^{-1} = \omega\mathbb{1} - H$ is block tridiagonal as well. A well-known result states that one can obtain the diagonal blocks of G , and hence the LDOS, using the following block-by-block algorithm: [52, 143, 69]

$$\mathbf{G}_{ii} = [\omega\mathbf{1} - \mathbf{a}_i - \mathbf{C}_i - \mathbf{D}_i]^{-1}. \quad (3.8)$$

\mathbf{C}_i and \mathbf{D}_i are calculated from the following expressions:

$$\mathbf{C}_i = \begin{cases} \mathbf{0} & \text{if } i = 1 \\ \mathbf{b}_{i-1}^\dagger [\omega\mathbf{1} - \mathbf{a}_{i-1} - \mathbf{C}_{i-1}]^{-1} \mathbf{b}_{i-1} & \text{if } 1 < i \leq N_x \end{cases} \quad (3.9)$$

and

$$\mathbf{D}_i = \begin{cases} \mathbf{0} & \text{if } i = N_x \\ \mathbf{b}_i[\omega\mathbf{1} - \mathbf{a}_{i+1} - \mathbf{D}_{i+1}]^{-1}\mathbf{b}_i^\dagger & \text{if } 1 \leq i < N_x. \end{cases} \quad (3.10)$$

This algorithm is very fast compared to full exact diagonalization. Taking into account the block matrix inversions needed, the computational complexity of this algorithm is $O(N_x N_y^3)$. This allows us to make N_x very large without significantly impacting performance, and this results in reducing finite-size effects in that direction considerably. In contrast, because the complexity scales as the cube of the length along the y -direction, N_y is taken to be considerably smaller than N_x . However, even in that case the scaling of the complexity with N_y is still very favorable compared to other methods. N_y in turn can be made much larger than the typical length of the system in exact diagonalization studies. We again reiterate that this procedure is exact—no approximations or truncations have been performed at any stage of the computation. Recursive techniques such as this, which make use of the sparsity of the Hamiltonian matrix, are very widely used in the quantum transport community to compute Green's functions [38, 136, 184, 103, 104, 94].

We then obtain the LDOS of the full system from the diagonal blocks \mathbf{G}_{ii} using Eq. 4.8. For our computations we took $N_x = 1000$ and $N_y = 120$. The LDOS maps were then extracted from the middle 100×100 subsection of the system. We note that this 100×100 field of view is similar to what present-day STS measurements are capable of. While minor artifacts from the open boundary condition along the x -direction remain, the very large value of N_x and taking the LDOS maps from the middlemost segment of the system combine to ensure that these effects are minimized. In obtaining the LDOS we used a

small finite inverse quasiparticle lifetime given by $\eta = 0.01$, expressed in units of t .

The power spectrum can then be straightforwardly computed by performing a fast Fourier transform on the real-space maps. The quantity we are interested in is the amplitude of the Fourier-transformed maps, $|\rho(\mathbf{q}, E)|$.

3.2.2 Modeling the Measurement Process

Our discussion beforehand neglected the specifics of the tunneling process between the tip and the CuO_2 plane. Here we will discuss how to incorporate the “fork mechanism,” an effective description of the tunneling process, in our computations. This mechanism was first proposed as an attempt to account for some inconsistencies between experimentally- and theoretically-obtained maps for zinc-doped BSCCO [129]. The motivation was the observation that, for zinc-doped BSCCO, the LDOS maps show no suppression *at* the impurity site, whereas theory predicts that maximal suppression should occur precisely there. One possibility is that some kind of filtering mechanism occurs when an electron tunnels from the STM tip to the copper-oxide plane. Martin *et al.* argued that the tunneling matrix element is actually of a d -wave nature [109]. Because the electron would have to tunnel through an insulating BiO layer before reaching the CuO_2 layer, the most dominant tunneling process involves *nearest-neighbor* $3d_{x^2-y^2}$ orbitals. The filtered LDOS at a site thus consists of a sum of the LDOS at the four nearest-neighbor sites *and* multiple pairwise interference factors. Such a filtering mechanism has been put on rigorous footing in recent first-principles work [93].

Here we adopt the simplest form of the fork mechanism and recast this into the Green's function formalism we use in our computations. We introduce a filter function $f(\mathbf{r}, \mathbf{r}')$ which incorporates the tunneling matrix elements between the STM tip and the the CuO₂ plane. The filtered LDOS, $\rho_f(\mathbf{r})$, can therefore be expressed as a generalized convolution between the two-point Green's function G and f :

$$\rho_f(\mathbf{r}, E) = -\frac{1}{\pi} \text{Im} \sum_{\mathbf{r}_1, \mathbf{r}_2} f(\mathbf{r} - \mathbf{r}_1, \mathbf{r} - \mathbf{r}_2) \quad (3.11)$$

$$\times G_{11}(\mathbf{r}_1, \mathbf{r}_2, E + i0^+). \quad (3.12)$$

The filtering mechanism can be incorporated by a suitable choice of f . For instance, to have s -wave filtering (*i.e.*, direct tunneling, which should result in the *bare* LDOS), the filter function is simply given by

$$f(\mathbf{r}, \mathbf{r}') = \delta_{\mathbf{r},0} \delta_{\mathbf{r}',0}, \quad (3.13)$$

which would simply result in Eq. 4.8. To have the desired d -wave fork effect, the following choice of f is needed:

$$f(\mathbf{r}, \mathbf{r}') = (\delta_{\mathbf{r},\hat{x}} + \delta_{\mathbf{r},-\hat{x}} - \delta_{\mathbf{r},\hat{y}} - \delta_{\mathbf{r},-\hat{y}}) \times (\delta_{\mathbf{r}',\hat{x}} + \delta_{\mathbf{r}',-\hat{x}} - \delta_{\mathbf{r}',\hat{y}} - \delta_{\mathbf{r}',-\hat{y}}). \quad (3.14)$$

Here \hat{x} and \hat{y} are unit vectors in the x - and y -directions, respectively.

Now we discuss how this is implemented in our computations. Observe that Eq. 3.12 with a d -wave filter has sixteen terms. This presents a complication in our block-by-block algorithm, because now we will have to obtain the first and second block diagonals above and below the main block diagonal. To be more precise, in addition to \mathbf{G}_{ii} , we will need the following eight other blocks to calculate $\rho_f(\mathbf{r}, E)$: $\mathbf{G}_{i-1,i-1}$, $\mathbf{G}_{i-1,i}$, $\mathbf{G}_{i-1,i+1}$, $\mathbf{G}_{i,i-1}$, $\mathbf{G}_{i,i+1}$, $\mathbf{G}_{i+1,i-1}$, $\mathbf{G}_{i+1,i}$, and $\mathbf{G}_{i+1,i+1}$. Fortunately




Figure 3.4: Plots of the magnitudes of the various \mathbf{q}_i wavevectors as a function of energy E . Lines denote the expected dispersions of the \mathbf{q}_i wavevectors as predicted by the octet model. Points show observed peaks for the case of a single weak pointlike impurity with $V = 0.5$ at selected energies. Note that the dispersions for the large-wavevector peaks are shown without backfolding. We do not show peaks associated with \mathbf{q}_4 and \mathbf{q}_5 , as these cannot be discerned clearly from the numerically-obtained power spectrum for a weak impurity. These dispersions are consistent with the behavior of peaks as observed in experiment.

all off-diagonal blocks are calculable recursively using the following expressions: [52, 143]

$$\mathbf{G}_{ij} = \begin{cases} -[\omega\mathbf{1} - \mathbf{a}_i - \mathbf{D}_i]^{-1}\mathbf{b}_{i-1}^\dagger\mathbf{G}_{i-1,j} & \text{if } i > j, \\ -[\omega\mathbf{1} - \mathbf{a}_i - \mathbf{C}_i]^{-1}\mathbf{b}_i\mathbf{G}_{i+1,j} & \text{if } i < j. \end{cases} \quad (3.15)$$

Here, \mathbf{a}_i , \mathbf{b}_i , \mathbf{C}_i , and \mathbf{D}_i are defined in the same way as before.

3.3 POINTLIKE SCATTERERS

We first consider QPI arising from pointlike impurities. This is by far the most comprehensively studied form of disorder in the cuprates. QPI was first understood theoretically by considering the effect of a single isolated impurity on the LDOS of the cuprates [182, 25]. We revisit this single-impurity case first in order to lay down a reference

template in the form of this well-known case to facilitate comparisons with new results. We will then turn to the case of many pointlike impurities distributed randomly on the plane.

The phenomenological octet model is an empirical success—in experiment one can clearly identify a set of seven dispersing peaks in the Fourier transform of the LDOS maps. Given the knowledge of the dispersion of the d -wave Bogoliubov quasiparticles, one can construct, for a given bias voltage, contours of constant energy (CCEs) in the first Brillouin zone, which are given by solutions to Eq. 3.2 for a given energy E . These CCEs are closed banana-shaped contours until E is such that their tips reach the Brillouin zone boundary. Each of these four “bananas” is centered around a *node*—*i.e.*, one of four points along the normal-state Fermi surface where $\Delta_{\mathbf{k}}$ vanishes. Plots of these CCEs with the parameters we set are shown in Fig. 3.1. Within the octet model, scattering processes from one tip of a banana to another become dominant, owing to the large joint density of states between any two such points. These dominant scattering processes manifest themselves in a set of visible peaks at seven characteristic momenta \mathbf{q}_i , with $i = 1, 2, \dots, 7$ in the power spectrum. These momenta are shown in Fig. 3.2.

Because the banana-shaped contours change their shape as E changes, these \mathbf{q}_i 's should disperse; $|\mathbf{q}_7|$, for instance, should increase with increasing $|E|$. In Fig. 3.4 we reproduce the dispersions of the various \mathbf{q}_i wavevectors as predicted by the octet model and compare them with peaks obtained from exact numerical calculations involving a single weak pointlike scatterer. The expected dispersions are easily calculated from Eq. 3.2, making use of the fact that the density of states at energy E is strongly enhanced by contributions at points in momentum space where $|\nabla_{\mathbf{k}}E|$ is a minimum, which are precisely at the tips of the “bananas” [112]. Here it can be seen that most of the peaks




Figure 3.5: Real-space LDOS maps for the single weak pointlike scatterer case. Here an isolated pointlike impurity ($V = 0.5$) is placed in the middle of the sample. The field of view is 100×100 . Shown are maps corresponding to energies $E = \pm 0.100$ and $E = \pm 0.250$. Inset: a close-up view of the impurity.

from our numerics match quite well with the predictions of the octet model. The behavior of the peaks as one varies the energy matches very closely with what is seen in experiment.

3.3.1 Single Weak Pointlike Impurity

We first start with the best-case scenario as far as reproducing the phenomenology of the octet model is concerned: the case of a single pointlike scatterer. To examine this more clearly, we add an on-site energy of $V = 0.5$ to a single site in the middle of the field of view. This is a weak, *non-unitary* potential, so this would not induce resonances at zero energy. The LDOS maps results are shown in Fig. 3.5. In the real-space images, one can see clear, energy-dependent oscillations in the LDOS which emanate from the impurity core. Despite the weakness of the potential, these oscillations dominate the signal at all energies, and the isolated impurity itself can be easily seen. It should be noted that at the impurity site the LDOS is not suppressed, but instead has a finite value for the energies we considered.




Figure 3.6: Fourier-transformed maps for the single weak pointlike scatterer case, with $V = 0.5$. Power spectra for both positive and negative bias voltages are shown for energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$. Arrows indicate where the peaks corresponding to the characteristic momenta of the octet model show up in the upper-right quadrant. The color scaling varies linearly with energy.

In contrast to the rather limited information conveyed by the real-space maps, the Fourier-transformed maps, shown in Fig. 3.6, display considerably more information. These are identical to the Fourier maps computed using the standard single-impurity T -matrix method—as it should, since that is a different method of solving the same problem. These show peaks with positions that are indeed consistent with the octet model. However, one also sees that these peaks are little more than enhanced regions in a more diffuse background. Even when the potential is weak, the spectra are dominated by momenta that connect different segments of the bananas, giving rise to patterns consisting of diffuse streaks, blurry regions, and propeller-shaped sections. The special momenta of the octet model merely correspond to points at which the spectral weight is enhanced relative to the background. That

is, these points coexist alongside these background patterns that arise from other scattering processes. A noteworthy feature of the power spectra of the case of a weak point potential is that \mathbf{q}_4 and \mathbf{q}_5 are not discernable at all. The most dominant peaks are \mathbf{q}_2 , \mathbf{q}_3 , \mathbf{q}_6 , and \mathbf{q}_7 , which become even more pronounced at higher energies. It is quite telling that, even at the idealized single point-impurity level, the correspondence between the full numerics and the expectations from the octet model is not fully realized—we remind the reader yet again that experimental Fourier maps show *all* seven peaks.

As we have emphasized before, impurity cores are not seen in the data, which excludes the possibility that QPI is caused by strong local impurity potentials. However our real-space results suggest that even a weak impurity gives rise to telltale patterns in the LDOS that point to its existence, and that these weak impurities can be easily identified in real space. The Fourier-transformed maps featuring a single weak impurity also show rather imperfect correspondence with experiment—power spectra from STS show far sharper peaks than our theoretically-obtained maps display. As we will subsequently argue, the addition of any realistic details to this idealized case will have the effect of further blurring the sharp features in the Fourier spectra. The presence of these complicating factors compounds the difficulty of explaining the sharpness of the octet model QPI peaks as seen in experiments.

3.3.2 *Multiple Weak Pointlike Impurities*

The many-impurity case is the next case we will consider. This has in fact been considered before using either a multiple-scattering T -matrix approach [196] or exact diagonalization of the Bogoliubov-de Gennes




Figure 3.7: Real-space LDOS maps for a d -wave superconductor with a 0.5% concentration of weak pointlike scatterers ($V = 0.5$) distributed randomly across the CuO_2 plane. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

Hamiltonian for small system sizes [14]. Here we take advantage of the flexibility of the numerical method we use and obtain *exact* results for large system sizes. We randomly distribute many weak pointlike scatterers in our system, and to optimize the correspondence with experimental results, we take the concentration of such weak scatterers to be low, with only 0.5% of lattice sites possessing such an impurity. As in the isolated-impurity case, we take the strength of each impurity to be $V = 0.5$.

As in the single-impurity case, the impurities are easily visible in the real-space images, but in addition we also see stripe-like patterns covering the entire field of view, which are seen to depend on the energy (Fig. 3.7). At first glance these look strikingly similar to the real-space patterns due to QPI seen in the raw experimental data. It is worth noting that the original real-space QPI results were initially misidentified as stripy charge-density waves. On closer inspection, novel multiple-scattering effects are seen when impurities get close together, as already discussed in the literature [15, 195, 196]. For instance, when two impurities line up such that their diagonal streaks overlap each other




Figure 3.8: Fourier-transformed maps for a system with a 0.5% concentration of weak pointlike scatterers ($V = 0.5$). Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. The color scaling varies linearly with energy.

neatly, the streaks constructively interfere and have the effect that they become more intense.

The Fourier-transformed maps are themselves quite illuminating. The consequence of the randomness of the impurity positions is that the Fourier maps show *speckle patterns*, as demonstrated in Fig. 3.8. This is just in line with our expectations: the familiar speckle patterns produced by laser light scattering against a random medium have precisely the same origin. Not surprisingly, one sees very similar speckle in the *experimental* Fourier maps. At these low impurity concentrations, the outcome is a speckled version of the single-impurity results. This looks much more like the real data, and the special momenta of the octet model are by and large still discernible in this case. The peaks that are prominent in the single-impurity case are similarly visible, with the difference that there is much more fuzziness present in these




Figure 3.9: Filtered real-space LDOS maps for a d -wave superconductor with a 0.5% concentration of weak pointlike scatterers ($V = 0.5$) distributed randomly across the CuO_2 plane. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

regions. However, because this is simply a many-impurity version of the single weak-scatterer case, this inherits the fact that no large spectral weight is associated with the \mathbf{q}_4 and \mathbf{q}_5 wavevectors.

To complete the discussion of the multiple weak-impurity case, we will include the fork effect, discussed earlier in Section 3.2, and see whether this leads to dramatic differences in the observed real-space and Fourier-space maps. In Figs. 3.9 and 3.10 we show plots with the filtered LDOS for the weak-impurity case. It can be seen that the impurities are considerably more visible in the filtered real-space maps than in the unfiltered ones. The patterns in the filtered real-space maps resemble those found in the bare cases. One takeaway from this case is that for weak impurities the individual impurities remain visible whether the fork effect is present or not.

The Fourier transforms of the filtered maps have a number of interesting features. Most of the momenta predicted by the octet model do show up in the power spectrum, and, notably, the locations of the peaks are *not* altered relative to the unfiltered case. This is not surprising, as the fork effect does not alter the dispersion of the Bogoli-




Figure 3.10: Fourier-transformed filtered maps for a system with a 0.5% concentration of weak pointlike scatterers ($V = 0.5$). Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. The color scaling varies with energy.

ubov quasiparticles, so the basic physics of the octet model remains in place. The main qualitative effect of the fork mechanism is the shifting of spectral weight from one part of momentum space to another, resulting in some differences from the unfiltered case—but nothing that results in the complete suppression of peaks expected from the octet model. The fork effect preserves the special momenta of the octet model. The shifting of the spectral weight however results in fuzzier peaks than in the unfiltered case.

The overall effect of the fork mechanism, at least in our simple treatment, is to amplify or suppress portions of the power spectrum without altering the presence of peaks that the octet model predicts will be present. In this sense the fork mechanism, while indeed a crucial phenomenon that one must ultimately incorporate in any description of the tunneling process, plays only a minor role in the overall descrip-




Figure 3.11: Real-space LDOS maps for a d -wave superconductor with a 0.5% concentration of unitary pointlike scatterers ($V = 10$) distributed randomly across the CuO_2 plane. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

tion of quasiparticle interference in BSCCO. The issues associated with the pointlike impurity case sans the fork effect—that the impurities are visible in real space and that the peaks seen in experiment are sharper than seen in numerical simulations—remain even when the fork effect is taken into account. In this sense the issues we discussed require a resolution beyond simply accounting for filter effects, and require examining whether the form of disorder we had used—namely, weak pointlike scatterers—is indeed correct.

3.3.3 Multiple Unitary Pointlike Impurities

For completeness we discuss the case where many *unitary* pointlike scatterers are present, especially in relation to the weak-potential case we previously tackled. Plots are shown in Figs. 3.11 and 3.12. In these plots we took the many-impurity disorder configuration to be the same as in the weak case, and we set $V = 10$. This form of disorder provides a realistic description of zinc-doped BSCCO, as zinc impurities are known to behave as unitary scatterers [129].




Figure 3.12: Fourier-transformed maps for a system with a 0.5% concentration of unitary pointlike scatterers ($V = 10$). Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. The color scale is the same for all energies

It is worth noting the similarities and differences between the weak-impurity and unitary-impurity cases. The real-space pictures for both cases are similar in that the individual impurities themselves can be easily detected. There is a difference, however: in the unitary case, the LDOS is heavily suppressed at the impurity site, whereas in the weak-impurity case it is generally not so. Real-space maps from both weak- and unitary-scatterer cases feature long-ranged diagonal streaks, but the modulations for the unitary-scatterer case are much more pronounced than in the weak case. The power spectra of the unitary-impurity case also display considerable differences from those of the weak case. While peaks at the same locations and with similar dispersive behavior can be observed in both cases, the weights of those peaks are different. In particular, \mathbf{q}_1 , \mathbf{q}_4 , and \mathbf{q}_5 are much stronger than in the weak case, and in fact become the most prominent wavevectors in




Figure 3.13: Plot of the impurity weight s (defined in Eq. 3.16) versus the the impurity strength V for $E = 0.100$ and $E = 0.250$. Here we consider a single pointlike impurity located in the center of the sample.

the power spectrum as energies increase. That said, the Fourier maps are far noisier than in the weak case, and as a consequence of strong scattering due to the large size of V , the main feature of the power spectrum is a series of diffuse streaks originating from scattering between points on CCEs. In a manner similar to that of the weak-impurity case, the peaks corresponding to the octet momenta emerge as the special points along these streaks with the highest spectral weight. These streaks in the unitary case are a considerably more prominent feature of the power spectrum than in the weak-impurity case.

3.3.4 Dependence of the Power Spectrum on the Impurity Strength

While we have restricted ourselves to the case of pointlike impurities, our results for weak and unitary impurities suggest that even within the pointlike model of disorder, qualitatively different behavior can be observed by varying the impurity strength. One could then ask if it is possible to identify whether the QPI observed in experiment is due primarily to unitary or weak scatterers. We will attempt to provide a

measure that quantifies the impact of the impurity strength V on the power spectrum.

Our main measurable of interest will be a quantity s , which we dub the *impurity weight* and define in the following way:

$$s(V, E) = \frac{\sum_{\mathbf{q} \in \text{BZ}} |\rho(\mathbf{q}, V, E)| - |\rho(\mathbf{q} = \mathbf{0}, V, E)|}{\sum_{\mathbf{q} \in \text{BZ}} |\rho(\mathbf{q}, V, E)|}. \quad (3.16)$$

Here $\rho(\mathbf{q}, V, E)$ is the Fourier transform of the LDOS map at energy E of a d -wave superconductor with a single pointlike impurity with strength V positioned in the middle of the field of view. As Eq. 3.16 shows, the impurity weight is simply the ratio of the integrated power spectrum *without* the $\mathbf{q} = \mathbf{0}$ contribution to the *total* integrated power spectrum (*i.e.*, *with* the $\mathbf{q} = \mathbf{0}$ contribution). $\rho(\mathbf{q} = \mathbf{0}, V, E)$ is removed from the numerator because that contribution is what one obtains when Fourier-transforming an LDOS map of a spatially homogeneous d -wave superconductor. The numerator of Eq. 3.16 thus describes only the contributions of the inhomogeneities to the power spectrum. One then expects that in the limit where the impurity is very weak, the power spectrum is dominated by the $\mathbf{q} = \mathbf{0}$ contribution and hence the impurity weight s is very small. We note that because of the underlying lattice the power spectrum is backfolded into the first Brillouin zone. We consider only unfiltered LDOS maps and their Fourier transforms.

We plot s as a function of V for two representative energies in Fig. 3.13. We let V vary from $V = 0.25$ to $V = 10$, covering the unitary- and weak-scatterer cases discussed in depth earlier, and consider $E = 0.100$ and $E = 0.250$. It can be seen that when the impurity is weak, s is a small quantity that depends approximately linearly on V . There is a broad crossover region around $V \approx 2$ where s begins to




Figure 3.14: Real-space LDOS maps for a d -wave superconductor with a single smooth-potential scatterer ($V_{sm} = 0.5$, $L = 4$, $z = 2$) located at the center of the field of view and off the CuO_2 plane. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

increase more slowly with V . For larger values of V corresponding to unitary scatterers, s does not show any dependence on V and saturates to a fixed value.

As a tool for potentially identifying the nature of pointlike scatterers in experiment, the impurity weight is admittedly limited, unless one already knows this for cuprates that are already firmly identified as hosting unitary scatterers, such as zinc-doped BSCCO. The main take-away from these results is that for weak scatterers the impurity weight is less than for unitary ones. In this light it would be interesting to revisit data from BSCCO with and without zinc impurities and calculate the impurity weight for various bias voltages. One identifying signal that QPI in BSCCO is caused by weak impurities is an s -value that is less than that obtained from zinc-doped BSCCO.

3.4 SMOOTH DISORDER

When one takes into account the chemistry of intrinsic disorder in the cuprates, it is difficult to justify pointlike disorder as a possible source




Figure 3.15: Fourier-transformed maps for a system with a single smooth-potential scatterer ($V_{sm} = 0.5$, $L = 4$, $z = 2$). Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. Only \mathbf{q}_7 is visible at low energies, while both \mathbf{q}_1 and \mathbf{q}_7 can be seen at higher energies. The color scale is the same for all energies.

of the effects we study. Since the early history of the high- T_c field, the prevailing understanding is that the doping mechanism is more closely related to modulation doping. The cuprate planes are widely assumed to be chemically very clean. The dopants are located in the ionic/insulating buffer layers some distance away from the metallic planes. The dopants are charged impurities which act as sources for poorly screened Coulomb potentials, which in turn affect the physics on the CuO_2 planes. The overall result is a smooth disorder potential which is characterized by small scattering wavevectors [130]. In a similar way, these dopants can also affect the tilting patterns inside the cuprate planes, and the elastic strain will result in a smooth form of disorder as well [42]. Smooth disorder potentials have been invoked in explaining the apparent discrepancy between the magnitude of the trans-

port and single-particle lifetimes; the former, which depends heavily on large-momentum scattering, is much smaller than the latter, and hence any scattering that occurs is argued to be forward (*i.e.*, small-momentum) scattering due to impurities located off the CuO_2 planes [2, 172, 3].

Previous theoretical treatments of smooth disorder have been motivated by bulk measurements [76, 10], but there has been a good amount of work motivated by STS studies as well [196, 124, 125]. In particular, Nunner *et al.* provide a comprehensive treatment of the Fourier spectra of a single isolated weak smooth scatterer [125]. However, in general, work on this form of disorder has not been as extensive as that on pointlike disorder, especially in the limit where a very large number of smooth scatterers are present. Following our treatment of pointlike scatterers, we will first revisit the case of a single smooth scatterer, first studied by Nunner *et al.*, to provide a picture of which scattering processes dominate. We will then discuss the consequences on the LDOS and the power spectrum when one has a large number of these impurities in the sample. We will also look at the sensitivity of the power spectrum to changes in the screening length of Coulomb potentials, especially as such details are not microscopically known.

Smooth potential scatterers in d -wave superconductors are not quite as easy to model as pointlike scatterers, due to the fact that one cannot apply the T -matrix formalism to this form of disorder to obtain the LDOS. The typical method involves extracting the LDOS directly by diagonalizing the Bogoliubov-de Gennes Hamiltonian. This has the restriction that only small systems can usually be accessed. However, smooth scatterers are easily treated by the numerical method that we use, with the advantage that we can scale up the system size to better visualize the LDOS. The flexibility of our method allows us to realisti-

cally model a smooth disorder potential *a la* modulation doping. We model smooth disorder using a screened Coulomb potential arising from a source located *outside* the copper-oxygen plane:

$$V(\mathbf{r}) = V_{sm} \frac{e^{-\frac{\sqrt{(\mathbf{r}-\mathbf{r}_i)^2+z^2}}{L}}}{\sqrt{(\mathbf{r}-\mathbf{r}_i)^2+z^2}} \quad (3.17)$$

Here \mathbf{r}_i is the location of the impurity projected onto the CuO_2 plane, z is the distance along the z -axis from the CuO_2 plane to the impurity, and L is the screening length. We will take the potential to be weak, with $V_{sm} = 0.5$, and as a typical case we set $z = 2$ and $L = 4$ in units of lattice constants, so the length scales are small relative to the system size.

3.4.1 Single Smooth Scatterer

At the single-impurity level, there are already rather drastic differences between the maps for the smooth scatterer and those for the pointlike one. Fig. 3.14 shows real-space LDOS maps for a smooth scatterer for various energies. Note the scale that we used to make the image clearer—the modulations are *much* smaller than in the pointlike impurity case. The LDOS is not suppressed above the impurity site, but is reduced from the clean-limit value only by a small amount. There is a pattern of crisscrossing diagonal streaks with four-fold rotational symmetry centered about the impurity site. When one uses the same scale as we used in the pointlike case to visualize this, these patterns are quite hard to see.

When one takes the Fourier transform of these LDOS maps, the differences from the pointlike case are even more pronounced, as one can see in Fig. 3.15. Unlike in the case of a pointlike scatterer, the Fourier-

transformed maps show that only *small-momenta* scattering processes contribute to the LDOS modulations. Large-momenta processes are almost completely suppressed. A closer examination reveals that only *intranodal* scattering processes occur in the presence of smooth potentials at low energies. That is, scattering occurs only between states lying on the same “banana.” This can be seen by looking at the surviving peaks. For a broad range of energies, only \mathbf{q}_7 —the peak corresponding to diagonal tip-to-tip scattering along the same “banana”—survives.

With increasing energy, even \mathbf{q}_7 becomes suppressed. A faint peak corresponding to \mathbf{q}_1 begins to appear in the power spectrum, but it is much less visible than \mathbf{q}_7 was at lower energies. The spectrum shows mostly streaks corresponding to small-momenta intranodal processes, as well as peaks in the horizontal and vertical directions where these streaks overlap. The mostly incoherent momenta seen in the power spectrum and the absence of any prominent peaks explain why the real-space picture is largely featureless. There are no longer any processes corresponding to \mathbf{q}_7 that will give rise to periodic modulations along the diagonal directions. As in the previous real-space picture, there is no suppression of the LDOS above the impurity; instead there is only a small reduction of the LDOS.

The takeaway from the single-impurity case is that impurity-induced modulations in the LDOS do occur for smooth scatterers, as they do for pointlike scatterers. The crucial difference is that large-momentum scattering is absent, thanks to the smoothness of the potential—even when $V(\mathbf{r})$ is reasonably short-ranged, with a screening length on the order of a few lattice constants.




Figure 3.16: Real-space LDOS maps for a d -wave superconductor with a 20% concentration of smooth scatterers ($V_{sm} = 0.5$, $L = 4$, $z = 2$) distributed randomly over the buffer planes adjacent to the CuO_2 plane. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

3.4.2 Multiple Smooth Scatterers

Now that we have intuition about the single smooth scatterer, we can discuss the extension to the case with a very large number of such impurities. We will take the number of smooth scatterers to be 20% of the total number of lattice sites and randomly place them across the sample. Real-space and Fourier-transformed plots of one realization of disorder are plotted in Figs. 3.16 and 3.17.

At low energies the real-space map display stripe-like patterns, featuring modulations in the diagonal directions, which display striking similarities to STS measurements of BSCCO. One could form the impression that they look even more akin to the stripy QPI patterns in the experimental data than what we found for pointlike disorder. Moreover, there is now no discernable sign of the impurity cores. Insofar as these cores were present (albeit difficult to discern) for the single smooth-impurity case, they are now washed away by multi-impurity interference effects at these high concentrations. The absence of clear-cut indications of the precise locations of the off-plane impurities is




Figure 3.17: Fourier-transformed maps for a system with a 20% concentration of smooth scatterers ($V_{sm} = 0.5$, $L = 4$, $z = 2$). Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. Only \mathbf{q}_7 is visible at low energies, while both \mathbf{q}_1 and \mathbf{q}_7 can be seen at higher energies. The color scale is the same for all energies.

consistent with the experimental results obtained by McElroy *et al.*, who find that while the positions of the impurities are indeed correlated with the LDOS in that the areas with LDOS suppression at low energies ($|E| < 60$ meV) are likely to be found near the impurities, this correlation is not by any means perfect [111]. The suppression of the LDOS does not imply that an interstitial impurity is above that site; indeed, experiment shows that many regions where the LDOS is suppressed also occur away from impurity sites.

This similarity to real-space experimental images is deceiving, however. Like in the case of the single smooth scatterer, the power spectrum of the many-impurity map here shows suppression of large-momentum internodal scattering processes. The main feature of the power spectrum is a band of wavevectors in the diagonal directions




Figure 3.18: Plot of the Fourier transforms of smooth potentials (given by Eq. 3.17) with various screening lengths L , shown for momenta in the range $0 \leq |\mathbf{k}| \leq 2\pi$. Inset: semi-log plot of the same quantities. Note that V_{sm} is adjusted so that $V(\mathbf{r} = \mathbf{0})$ is the same for all L .

forming a cross in the center of the first Brillouin zone. These arise from intranodal scattering processes between states on one “banana.” These diagonal streaks have a length that is set by \mathbf{q}_7 . At low energies, no peaks in the spectrum arise from internodal scattering.

As in the single-scatterer case, when energies increase, the diagonal wavevectors become less pronounced in the power spectrum, while wavevectors in the horizontal and vertical directions become more visible. It can be seen that instead of a diagonal cross, one now has a regular cross, with a broad range of wavevectors in the horizontal and vertical directions now being the dominant characteristic of the power spectrum. These horizontal and vertical streaks feature a length scale roughly set by \mathbf{q}_1 . This goes hand-in-hand with the fact that the real-space map now features vertical and horizontal stripe-like patterns, instead of diagonal stripes at lower energies. When QPI is the mechanism for the appearance of these stripes, it is expected that the orientation of these patterns will change depending on the energy. This is again different from the case of static stripe order, where stripe patterns remain fixed even when the energy is varied [87].

3.4.3 *Quantifying the Range of the Potential*

In hindsight it is clear why the signals for QPI are different for smooth and pointlike disorder. The observed power spectra are sensitive to the length scales associated with the disorder potential, since the distribution of the weight in the Fourier maps is set by the characteristic wavevectors of the scattering potential. This is seen in Fig. 3.18, which shows the Fourier transform of the scattering potential $V(\mathbf{r})$ for various screening lengths L . In plotting these we varied V_{sm} in Eq. 3.17 so that $V(\mathbf{r} = \mathbf{0})$ is the same for different L . It can be seen that for all the values of L that we consider, the Fourier transform of $V(\mathbf{r})$ features a very steep dropoff with increasing momentum. The dropoff is most prominent for bigger values of L , but is also seen for small ones as well. The Fourier amplitudes at large momenta are larger for small L , but they still decrease markedly as $|\mathbf{k}|$ is increased. As this implies that the matrix elements of the scattering potential for large momenta are small, such large-momenta scattering processes will be far less prominent. This explains why, in the power spectra for smooth impurities, \mathbf{q}_7 is the only octet-momentum peak visible for low and intermediate energies—as seen in Fig. 3.4, \mathbf{q}_7 is the smallest peak for a wide energy range, and its magnitude falls within the range where the Fourier transform of the smooth potential is finite. It is interesting to note that as one moves toward higher energies, \mathbf{q}_1 becomes small enough for its magnitude to fall within the aforementioned range of allowed scattering momenta, and its signals are indeed faintly visible in the power spectrum. It is however nowhere near as visible at higher energies as \mathbf{q}_7 is at lower energies, a fact that can be attributed to coherence factors that suppress scattering processes between states where the gap has the same sign. All this is to be contrasted with pointlike disorder, whose Fourier transform is a constant which depends only on the im-

purity strength and for which kinematical considerations are the main determinant of the allowed scattering processes.

By measuring carefully not only the dispersions but also the spectral weights of the peaks in the power spectrum, it should be possible in principle to get a quantitative estimate of the typical range of the disorder potential. To the best of our knowledge, this has not been attempted yet. Here we will attempt to quantify in a simple manner the dependence of the power spectrum on the screening length of the Coloumb potential. We introduce a number w that will quantify how much spectral weight is associated with large-momentum scattering processes:

$$w(L, E) = \frac{\sum_{\mathbf{q} \in A} |\rho(\mathbf{q}, L, E)| - |\rho(\mathbf{q} = \mathbf{0}, L, E)|}{\sum_{\mathbf{q} \in \text{BZ}} |\rho(\mathbf{q}, L, E)| - |\rho(\mathbf{q} = \mathbf{0}, L, E)|}. \quad (3.18)$$

$\rho(\mathbf{q}, L, E)$ is the Fourier map associated with a single smooth scatterer with screening length L in the center of the field of view, taken at energy E . As before we will also vary V_{sm} in Eq. 3.17 so that $V(\mathbf{r} = \mathbf{0})$ is independent of L . We set $z = 2$. A in this instance is defined as the subset of the Brillouin zone centered about the Γ point where $-a\pi \leq q_x \leq a\pi$ and $-a\pi \leq q_y \leq a\pi$, and $a < 1$. We will set $a = 0.4$ in our numerical calculations.

The point of introducing w is that it is simply the ratio of the integrated power spectrum within A (without the $\mathbf{q} = \mathbf{0}$ contribution) to the integrated power spectrum within the first Brillouin zone (again without the $\mathbf{q} = \mathbf{0}$ part). If most of the weight in the power spectrum is associated with small-momentum scattering processes, w should be close to 1, whereas if more spectral weight is associated with large-momentum processes, such as in the case of a pointlike scatterer, w should be small.




Figure 3.19: Plot of w (defined in Eq. 3.18) versus the screening length L for $E = 0.100$ and $E = 0.250$. Here we consider a single smooth impurity located in the center of the sample. As discussed in the text, the value of V_{sm} is chosen so that $V(\mathbf{r} = \mathbf{0})$ is the same for all values of L . The two data points at $L = 0$ correspond to the values of w obtained for a weak pointlike scatterer ($V = 0.5$) and a unitary pointlike scatterer ($V = 10$) at energy $E = 0.100$. The smallness of these values indicates that large-momentum processes are a prominent part of the power spectra for these pointlike scatterers.

Fig. 3.19 shows plots of w versus L for energies $E = 0.100$ and $E = 0.250$. It can be seen that when L is large (*i.e.*, $L > 2$), w is large and saturates to a fixed value with increasing L . This means that in this regime, the vast majority of the spectral weight is associated with small-momentum processes. On the other hand, when L is small, w becomes small as well, implying that the power spectrum hosts more contributions from large-momentum processes which show up outside A in the power spectrum. We can see that it is only with very small values of L that we start to see behavior resembling that of the pointlike scatterer, in which both small- and large-momentum processes figure prominently in the power spectrum.

Although a detailed study of weight distributions in experimentally-obtained Fourier maps has not yet been undertaken, it appears that the strength of large-momentum scattering, at least as evidenced from STS experiments, is actually quite large. These experimental results suggest that disorder is close to the point-scatterer limit. Given what

3.5 SPATIALLY RANDOM ON-SITE ENERGIES




Figure 3.20: Real-space LDOS maps for a d -wave superconductor with random on-site energies, normally distributed with $M_V = 0.01$. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

is commonly believed about the nature of intrinsic disorder in the cuprates, this is surprising, if not entirely unreasonable. The reason behind the prominence of the large-momentum peaks in QPI spectra can be considered alongside the problem of the sharpness of the octet-model peaks as two of the primary mysteries of the results of QPI.

3.5 SPATIALLY RANDOM ON-SITE ENERGIES

To complete our survey of the effects of various kinds of disorder on STS results, we now turn to yet one more well-known form of disorder: a random and uncorrelated distribution of on-site energies throughout the sample. This is the form of disorder that underlies Anderson localization in metals. Because we do not have isolated impurities in this case, with the on-site energies varying from one site to another and numerous multiple-scattering processes occurring as a result, the T -matrix method cannot be easily applied to this problem to obtain the LDOS. In contrast, the numerical method we use here allows us to obtain LDOS maps directly and efficiently.




Figure 3.21: Fourier-transformed maps for a system with random on-site energies, normally distributed with $M_V = 0.01$. Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. The color scaling varies linearly with energy.

To be more specific, on each site we have a *random* perturbation V_R in addition to the spatially uniform mean-field chemical potential μ . In other words the on-site potential at site \mathbf{r} is given by the sum $\mu + V_R(\mathbf{r})$. For simplicity we will take $V_R(\mathbf{r})$ to be drawn from a Gaussian distribution with the following properties:

$$\langle V_R(\mathbf{r}) \rangle = 0, \quad (3.19)$$

$$\langle V_R(\mathbf{r}) V_R(\mathbf{r}') \rangle = M_V^2 \delta_{\mathbf{r}\mathbf{r}'}. \quad (3.20)$$

Here the angular brackets denote averaging over disorder realizations. The width of the distribution is parametrized by the standard deviation M_V ; we will use this to characterize the strength of the disorder potential.

Why do we pick this form of disorder? From our previous discussion of pointlike and smooth potential disorder, it is clear that in order to reproduce both the real-space and Fourier-transformed results from STS measurements, one must have both real-space maps that simultaneously have LDOS modulations and feature no obvious signs of impurity cores; and power spectra that show peaks arising from internodal and intranodal scattering. Here the random-potential model could sidestep the difficulties faced by our previous hypothesized scenarios. First, if we pick our distribution to be sufficiently narrow (and hence weak), there is the possibility that we could have real-space modulations without having visible impurity cores that arise from isolated potential perturbations, as was the case in the pointlike case we discussed earlier. Second, this form of disorder, similar to the pointlike scatterer, is short-ranged. This would then not have the suppression of internodal scattering that is a feature of smooth potential disorder with finite correlation length. As a result it could potentially feature both small- and large-wavevector peaks in the power spectrum. A similar form of random on-site disorder was considered by Atkinson *et al* [14].

To check whether these expectations are ultimately borne out, we numerically obtain real-space and Fourier-transformed maps for one realization of random on-site disorder. We make disorder weak by setting the width of the distribution to be narrow. The results are plotted in Figs. 3.20 and 3.21.

The real-space maps feature as before modulations whose structure can be discerned, but not to a similar extent as the pointlike- or smooth-scatterer cases. In this particular scenario one cannot tell whether an impurity is present or not—the signatures we have come to expect from the isolated pointlike impurity are not present here at all. Instead what we have are modulations, primarily in the diagonal direc-

tions, with a crisscrossing pattern slightly similar to that found in the smooth-disorder case. Unlike in the smooth scatterer case, however, the stripe-like patterns are far more subdued. The maps obtained here look very similar to those taken from STS experiments.

This is shown even more so by the Fourier-transformed maps. We see that both small- and large-momentum scattering processes contribute to the observed QPI, as evidenced by peaks at small and large diagonal wavevectors. Interestingly, the power spectrum is very similar to that of the multiple-weak-impurity case. In particular, \mathbf{q}_2 , \mathbf{q}_3 , \mathbf{q}_6 , and \mathbf{q}_7 are strongly present, whereas signals of the three remaining \mathbf{q} -vectors are quite weak. This can be attributed to the fact that, in the superconducting state, coherence factors enter into the scattering amplitude [182, 133]. For scattering off of a weak potential, it turns out that the matrix element between two states with momenta \mathbf{k}_1 and \mathbf{k}_2 contains a factor

$$u_{\mathbf{k}_1}u_{\mathbf{k}_2} - v_{\mathbf{k}_1}v_{\mathbf{k}_2}, \quad (3.21)$$

where

$$u_{\mathbf{k}} = \text{sgn}(\Delta_{\mathbf{k}}) \sqrt{\frac{1}{2} \left(1 + \frac{\epsilon_{\mathbf{k}}}{E_{\mathbf{k}}} \right)}, \quad (3.22)$$

$$v_{\mathbf{k}} = \sqrt{1 - u_{\mathbf{k}}^2}. \quad (3.23)$$

This implies that if $\Delta_{\mathbf{k}_1}$ and $\Delta_{\mathbf{k}_2}$ have the same sign, the $\mathbf{k}_1 \rightarrow \mathbf{k}_2$ process will be suppressed. Conversely, whenever $\Delta_{\mathbf{k}_1}$ and $\Delta_{\mathbf{k}_2}$ have the opposite sign, that process will not be suppressed. This explains why the \mathbf{q}_2 , \mathbf{q}_3 , \mathbf{q}_6 , and \mathbf{q}_7 wavevectors—which connect states at which the values of the order parameter have opposite sign—are not suppressed, while the remaining ones are.

The possibility that one can have the Fourier-space signatures of QPI while having some qualitative similarities between the theoretical and

3.6 SPATIALLY RANDOM SUPERCONDUCTING GAP




Figure 3.22: Real-space LDOS maps for a d -wave superconductor with random pairing amplitudes, normally distributed with $M_{\Delta} = 0.01$. The field of view is 100×100 , and the energies shown are $E = \pm 0.100$ and $E = \pm 0.250$.

experimental real-space images suggests that this form of disorder—weak, narrowly distributed random potential disorder—can be responsible for the physics observed in the STS measurements. Having said this, the peaks in the power spectrum resulting from this form of disorder exhibit the same form of fuzziness as in the weak-impurity scenario. Also, the relative suppression of certain octet-model peaks suggests that even with this form of disorder, the same questions that affect the weak-impurity case affect the random site-energy model as well.

3.6 SPATIALLY RANDOM SUPERCONDUCTING GAP

STS measurements have demonstrated that the superconducting order parameter is in fact inhomogeneous [95, 45, 111]. It is then worthwhile to ask whether *gap disorder* could also be responsible for QPI. In this section we will consider the case of a d -wave superconductor with disorder only in the gap; we keep all other parameters (hoppings and




Figure 3.23: Fourier-transformed maps for a system with random pairing amplitudes, normally distributed with $M_\Delta = 0.01$. Shown are energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$, along with arrows showing where the octet wavevectors are expected to be found. Note that only three peaks— \mathbf{q}_1 , \mathbf{q}_4 , and \mathbf{q}_5 —are visible. The color scaling varies linearly with energy.

chemical potential) at their mean-field values. This ensures that we can identify the defining characteristics of QPI from pure gap disorder.

We will assume the simplest model for disorder in the gap that preserves the purely d -wave nature of the superconductor. Here only the nearest-neighbor pairing terms are disordered. The pairing amplitude between nearest-neighbor sites \mathbf{r} and \mathbf{r}' is of the form $\Delta_{0,\mathbf{r}\mathbf{r}'} + \Delta_{R,\mathbf{r}\mathbf{r}'}$, where $\Delta_{0,\mathbf{r}\mathbf{r}'}$ is the mean-field pairing amplitude and $\Delta_{R,\mathbf{r}\mathbf{r}'}$ is a *random* variable taken from some distribution. Like the random-potential case earlier, we assume that $\Delta_{R,\mathbf{r}\mathbf{r}'}$ is normally distributed, with zero mean and a standard deviation M_Δ , and, importantly, we will assume that

the value of $\Delta_{R,\mathbf{r}\mathbf{r}'}$ at one link $(\mathbf{r}, \mathbf{r}')$ is independent of $\Delta_{R,\mathbf{s}\mathbf{s}'}$ at any other link $(\mathbf{s}, \mathbf{s}')$. More precisely,

$$\langle \Delta_{R,\mathbf{r}\mathbf{r}'} \rangle = 0, \quad (3.24)$$

$$\langle \Delta_{R,\mathbf{r}\mathbf{r}'} \Delta_{R,\mathbf{s}\mathbf{s}'} \rangle = M_{\Delta}^2 (\delta_{\mathbf{r}\mathbf{s}} \delta_{\mathbf{r}'\mathbf{s}'} + \delta_{\mathbf{r}\mathbf{s}'} \delta_{\mathbf{r}'\mathbf{s}}), \quad (3.25)$$

for any two nearest-neighbor links $(\mathbf{r}, \mathbf{r}')$ and $(\mathbf{s}, \mathbf{s}')$. This form of gap disorder is short-ranged and as such should give rise to large-momentum scattering. It should be noted that gap maps from STS measurements do show that the gap variations in space obey a bell-curve-like distribution [45], which justifies to some extent this choice of distribution.

Plots for this form of gap disorder are shown in Figs. 3.22 and 3.23. We take M_{Δ} to have a value comparable to that of M_V discussed earlier, so both perturbations are of similar size. The real-space maps exhibit LDOS modulations that are sharper and more noticeable than in the random-potential case. There are stripe-like patterns with streaks in the vertical, horizontal, and diagonal directions. The patterns get far more pronounced with increasing energy. The maps for this case show a marked resemblance to that arising from smooth scatterers, but show considerably more structure in that modulations for more directions are present here than in the smooth-scatterer scenario. There is no signature akin to the single pointlike impurity of a localized center of the LDOS modulations. In this sense the results from pure gap disorder match closely real-space maps from experiment.

Like the random-potential case, wavevector peaks corresponding to large-momenta scattering processes are present. It is worth noting that, unlike the unitary pointlike scatterer and random-potential cases, only three peaks appear in the power spectrum: \mathbf{q}_1 , \mathbf{q}_4 , and \mathbf{q}_5 . This is because of the fact that scattering due to *weak* gap disorder involves only

processes connecting states at which the order parameter has the same sign [125, 135, 176]. Out of the seven wavevectors, only the three aforementioned ones correspond to such scattering processes. Curiously, these three momenta are precisely the same ones that are suppressed in the random-potential case. At larger energies, these peaks become more prominent, paralleling the progression in the real-space picture, where the modulations become more and more apparent with increasing energy.

In attributing QPI partially to gap disorder, however, we stress some caution. Our model of disorder involves an order parameter that varies over a length scale of one lattice constant. However, experimentally obtained gap maps show that this is generally not the case. These gap maps feature domains. The average value of the gap from one domain to another can change drastically, but the gap varies slowly within a domain [95]. While the steep change in the gap as one moves from one domain to another is captured well by our simple model, the near-constant nature of the gap within one domain is not. Thus the precise interplay between the smoothness of the gap within a domain and the sharp shifts in the gap from one domain to another cannot be seen from our simple model. Smooth gap disorder would have a similar effect as smooth potential disorder in suppressing large-momentum scattering, and thus a realistic model would very likely feature power spectra dominated by small-momentum process. That said, the problem of modeling the gap inhomogeneities accurately, incorporating both the inter-domain sharpness and intra-domain smoothness of the gap, is an interesting problem for future work.

3.7 DISCUSSION AND CONCLUSION

By utilizing the powerful real-space Green's function method introduced earlier, it is possible to study, in a systematic manner and with large fields of view, the consequences of various forms of distributed disorder on the physics of quasiparticle scattering interference. We made use of the standard method of modeling the low-energy electronic excitations deep in the d -wave superconducting state. In addition we also looked at the effects of proposed nontrivial tunneling processes that are potentially of relevance to STS experiments in the cuprates.

Much of the established intuition regarding the physics of QPI is based on results for a single impurity. When one considers a random distribution of such impurities, however, one deals with the problem of wave interference in a finite-sized random medium. The intuitive expectation is that speckle patterns are formed in the Fourier maps, and this is precisely what we find.

In the case of a low concentration of impurities with short-ranged potentials, the main difference from the single-impurity case is that the already-diffuse Fourier maps associated with a single impurity turn into speckle patterns that follow closely the weight distribution of the former. This underlies the intuition that the observed QPI can mostly be understood solely on the basis on single-impurity theory. Even in the case of distributed random disorder, in which there is no clear, well-defined sense of an *isolated* impurity, this correspondence with the single-impurity results can be observed clearly. This can be best seen in the case of Gaussian on-site potential disorder, whose power spectra resemble those of the weak single- and multiple-impurity cases.

The real-space patterns exhibit the characteristic energy-dependent stripe-like interference patterns which are also seen in the raw exper-

imental data. However, upon examining our results more quantitatively, we detect problems that suggest that the present way of interpreting STS experiments may have serious deficiencies. The problem is that the experimental QPI peaks are characterized by a sharpness in momentum space that cannot be reproduced with standard methods of modeling STS experiments, including ours. In the experimental maps, the seven sharp peaks can be discerned and in fact be tracked over a large range of energies. These seven peaks are found to disperse in accordance with the predictions of the octet model. The best-case theoretical scenario is the case of a low concentration of weak pointlike impurities, but even here matching our numerically obtained Fourier maps to those obtained from experiment becomes a stretch. The case of a single weak impurity in fact already demonstrates this problem: in addition to peaks, one sees continuous streaks arising from scattering between points on CCEs, whose spectral weight is only enhanced at momenta at the special “tip-to-tip” processes. In other words, the peaks in the cases we consider are not observed to be as prominent as in experiment—they happen to be the points which possess the largest spectral weight along the streaks corresponding to scattering between CCEs. When one goes beyond this single-scatterer paradigm and considers other, more general forms of distributed disorder, this sharpness is further reduced.

This is an exceptional circumstance. One usually expects that idealized models like ours will produce outcomes that are *sharper* than experimental data. The incorporation of the most general forms of disorder, which we implement in this work, should have the effect of adding fuzziness in the Fourier-space picture. In our models we ignore complicating factors such as frequency-dependent self-energies that could alter the picture for larger energies. Given the relative lack

of complications present in our models, this inability to reproduce the sharpness of experimental data is puzzling.

The outcomes of our simulations for many weak pointlike scatterers are perhaps the closest approach to experiment. However, taking these as an explanation is problematic since no impurity cores are seen in experiment. The case of many unitary pointlike scatterers is even more rife with problems because in this case impurity cores are much more visible and the strong scattering processes preclude the formation of prominent peaks in the power spectra, showing instead very fuzzy streaks corresponding to inter-CCE scattering. As a next case, there are very good reasons to believe that the intrinsic disorder in cuprates is of a smooth kind. The CuO_2 planes themselves are quite clean, lacking disorder from doping, while the chemical sources of disorder are located in the insulating buffer layers located some distance away from the superconducting perovskite planes. Our simulations of smooth disorder show that the large-momentum peaks are suppressed, owing to the fact that in the Fourier decomposition of the screened Coulomb potential the large-wavevector components have very small amplitudes. Our results seem to suggest that in order to reproduce the overall weight distribution seen in the experimental Fourier maps, one needs local, pointlike potentials. This is quite puzzling given what is now known about the chemical composition of the cuprates.

Disorder in the form of randomly distributed on-site energies is another scenario that gives rise to real- and Fourier-space maps that are very similar to those found in experiment. These are found to result in modulations in the LDOS without the presence of visible impurities, and power spectra for this form of disorder show peaks that originate from large-momentum scattering processes. The caveat with this form of disorder however is that, like the many-weak-impurity scenario, not all of the peaks are visible in the Fourier maps. We finally note that

as compelling an explanation as this is for the patterns seen in experimental data, it is difficult to argue from microscopic considerations why this form of disorder should exist—unlike pointlike and smooth disorder, whose possible origin in the cuprates can at least be justified on the level of chemistry.

We also examined in some detail the influence of gap disorder in the LDOS maps, using a simple model of gap inhomogeneities. It is well-established that in the cuprates, especially the underdoped ones, the gap magnitude is quite inhomogeneous, varying by a large amount in space. Our calculations show that this form of disorder scatters the Bogoliubov quasiparticles efficiently, generating a distinctive power spectrum and visible real-space patterns. These results suggest that gap disorder could potentially generate QPI as well. However, it is also known from experiment that this form of disorder is characterized by a short-distance cutoff scale on the order of the coherence length ~ 3 nm [95, 45, 75]. Gap disorder is therefore a smooth form of disorder, and its effect should thus be similar to that of smooth potential disorder.

What is the origin of this trouble? One possibility is that the physics underlying QPI in the cuprates is completely different from the standard explanation, which is centered on the quantum-mechanical scattering of Bogoliubov quasiparticles against quenched disorder. One could contemplate exotic possibilities involving the formation of *real* bound states at the special momenta of the octet model—the most obvious way to obtain sharp quantization in momentum space. However, we think that this is far-fetched. Direct, independent evidence for the presence of coherent Bogoliubov quasiparticles with a *d*-wave dispersion exists from angle-resolved photoemission spectroscopy. Moreover, the octet model is qualitatively highly successful in relating the dispersions from QPI to measured dispersions from ARPES. A concrete

possibility that builds on the scattering picture of QPI is that nematic quantum-critical fluctuations strongly enhance the amplitudes of the peaks [86, 85]. We suspect that the culprit is the tunneling process itself. On a quantitative level this is sensitive to the details of the microscopic electronic structure. Recent first-principles work demonstrates this vividly [93]. Kreisel *et al.* find nontrivial effects arising from microscopic details, such as the enhancement of large-momentum peaks in the power spectrum. Similarly, it may well be necessary to study disorder in a much more microscopic manner in order to capture the way it affects the microscopic intra-unit cell electronic structure [176]. We envisage that it may become possible to extract from such precise modeling of the microscopic tunneling process effective, coarse-grained models which can then be studied in the most general disordered case using the methods we have used in this work. The overarching message is that there is in all likelihood more to the beautiful STS images than meets the eye.

3.A APPENDIX: SINGLE UNITARY POINTLIKE SCATTERER

In this section we will briefly discuss the case of a single *unitary* pointlike scatterer. For reasons discussed in depth in the main text, this is not physically relevant for the experimental data we wish to revisit. That said, these are not unphysical—zinc impurities in BSCCO are an example of non-magnetic unitary scatterers, for instance. While QPI in clean cuprates is most likely caused by far weaker impurities, the properties of unitary pointlike scatterers are sufficiently different from those of weak ones that it is worth spending a few words delineating some of these differences.




Figure 3.24: Real-space LDOS maps for the single unitary pointlike scatterer case. Here an isolated pointlike impurity ($V = 10$) is placed in the middle of the sample. The field of view is 100×100 . Shown are maps corresponding to energies $E = \pm 0.100$ and $E = \pm 0.250$. Inset: a close-up view of the impurity.

In Figs. 3.24 and 3.25 we plot real-space and Fourier-transformed maps for a single strong impurity embedded in the middle of the sample. We take $V = 10$, ensuring that the impurity is a unitary scatterer. The real-space maps are qualitatively similar to those of the weak-impurity case. In both cases the impurity cores can easily be discerned. The main difference between the unitary- and weak-scatterer cases is that the LDOS at the unitary-impurity site is almost completely suppressed. Recall that in the weak-impurity case, the LDOS at the impurity site is finite.

The second noteworthy feature of the unitary scatterer is apparent in the Fourier-transformed maps. Because the potential is so strong, scattering between any two points lying on CCEs is allowed, resulting in very prominent streaks in the power spectrum. Many of the peaks from the octet model can be seen, similar to the case of the weak scatterer. However, the peaks that are most prominent here differ from those seen in the weak-scatterer case. Observe that when energies become high, \mathbf{q}_2 , \mathbf{q}_6 , and \mathbf{q}_7 become less visible. Streaks near the corners of the first Brillouin zone corresponding to internodal scattering re-

3.A APPENDIX: SINGLE UNITARY POINTLIKE SCATTERER




Figure 3.25: Fourier-transformed maps for the single unitary pointlike scatterer case, with $V = 10$. Power spectra for both positive and negative bias voltages are shown for energies ranging from $E = \pm 0.050$ to $E = \pm 0.250$. Arrows indicate where the peaks corresponding to the characteristic momenta of the octet model show up in the upper-right quadrant. The color scale is the same for all energies.

mains very prominent, but a peak at \mathbf{q}_3 is not as visible as it is in the weak-impurity case. In contrast, \mathbf{q}_1 , \mathbf{q}_4 , and \mathbf{q}_5 become far more visible and in fact become the most dominant wavevectors in the power spectrum. It is interesting to note that \mathbf{q}_4 and \mathbf{q}_5 are barely visible in the weak-impurity case; this can be attributed to the presence of coherence factors that suppress the amplitudes of these scattering processes [182, 133].

QUASIPARTICLE DENSITY OF STATES,
LOCALIZATION, AND DISTRIBUTED DISORDER IN
THE CUPRATE SUPERCONDUCTORS

4.1 INTRODUCTION

Disorder in the high- T_c superconductors has motivated many key experimental and theoretical advances in the field. Scanning tunneling spectroscopy (STS) has made wide use of the phenomenon of quasiparticle interference, which results from the presence of disorder, to provide a real-space probe of the underlying electronic nature of the cuprates [70, 112, 90, 99, 182, 25, 196, 125, 176, 93, 161]. On the theory side, the d -wave nature of the cuprate superconductors provided the impetus for various theoretical treatments of disorder which led to a number of differing and often contradictory predictions. Early theoretical work utilized a self-consistent treatment of disorder, which was found to result in a finite quasiparticle density of states (DOS) at the Fermi energy [54, 68, 100, 39]. Later work has shown within a similar diagrammatic approach that the DOS, in the most generic case, is suppressed [187, 67]. Other field-theoretical treatments of disorder in d -wave superconductivity found a vanishing DOS at $E = 0$ [118, 156, 155, 10]. The manner in which the DOS vanishes as $E \rightarrow 0$

varies from approach to approach, with exponents found to be either universal or disorder-dependent.

Meanwhile, experiments performed on $\text{YBa}_2\text{Cu}_3\text{O}_{6+\delta}$ consistently show a T -linear term in the specific heat at zero magnetic field, which points to a nonvanishing DOS at $E = 0$ [115, 116, 144]. How this nonzero DOS arises has been the subject of much speculation. According to standard self-consistent T -matrix theory, which assumes that impurities are located within the copper-oxide planes, this contribution is expected. It is interesting to note, however, that this T -linear term in YBCO persists even with very clean samples, prompting a number of exotic explanations, such as loop-current order coexisting with d -wave superconductivity [21, 4, 88, 181], which give rise to a finite DOS without invoking disorder. For $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$, the story is a bit more complicated: it appears that no definitive evidence in favor of or against a zero T -linear coefficient exists, and what is present instead is considerable variation in the measured values of this coefficient. For BSCCO-2212 at low temperatures, it was found that that the coefficient is small but finite and measurable [30, 80]. However, other experiments, performed at higher temperatures, find no discernible evidence in BSCCO-2212 for a coefficient on the same order as found in YBCO [167]. The results for the BSCCO family suggest that the cleaner the sample is, the smaller the T -linear coefficient becomes, with a large degree of variation present.

Given such a wide array of evidence suggesting that high-temperature superconductors do display a finite zero-energy quasiparticle DOS and the lack of any confirmation of alternative explanations, it is worth revisiting the effect of disorder, especially when incorporating inhomogeneities in the cuprates that do not fall under the random-site-energy or multiple-point-impurity categories. Previous numerical work has extensively focused on pointlike impurities and random on-site ener-

gies. In particular, Atkinson *et al.* found that for realistic models (*i.e.*, without a particle-hole symmetric band) with these two forms of disorder, the quasiparticle DOS becomes suppressed near $E = 0$ [14]. They point out that a *constant* DOS, as seen in experiment, cannot arise from either of these two disorder models.

In any case, what is known about the cuprates makes it difficult to argue that pointlike disorder is a possible origin of the finite DOS at the Fermi energy. The consensus regarding the CuO_2 planes is that they are generally clean. Pointlike disorder necessarily takes the form of dopants *within* the CuO_2 plane. Such substitutions will give rise to strong pointlike potentials. The most dramatic case of this is zinc-doped $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$, in which a small number of zinc atoms take the place of copper ones; STS studies of Zn-doped BSCCO show that the zinc impurities show behavior consistent with that of unitary scatterers [129]. In contrast, STS studies of clean cuprates do not show such strong local impurities, and the conductance maps obtained from such materials are more consistent with far weaker forms of disorder [70, 111, 150]. More reasonable is the expectation that impurities lie in the buffer layers adjacent to the CuO_2 planes [42, 124, 126]. As they are located in an insulating layer some distance from the CuO_2 plane, they act as a source of an electrostatic potential which, in contrast to local pointlike potentials, is *smooth*. These smooth potentials lead to small-momentum scattering processes. It is then worth examining the imprint of such smooth forms of disorder on the DOS.

In this chapter, we obtain the quasiparticle DOS of a two-dimensional d -wave superconductor subject to various kinds of disorder: pointlike disorder, random on-site disorder, and smooth disorder. We utilize an exact real-space numerical method that allows for the evaluation of the local density of states of a disordered system with very large system sizes (a typical calculation involves 100,000 sites). The same geome-

try of the system also enables the *direct* calculation of the localization length, which is a quantity that is difficult to extract from the exact diagonalization of small systems, given the large length scales over which localization occurs. An important feature of this work is its use of realistic band-structure and pairing parameters. As our method faces no difficulties with large system sizes, we do not need to resort to making the *d*-wave gap artificially large in order to sidestep finite-size effects in related methods like exact diagonalization, and we can thus make the parameters of our lattice *d*-wave superconductor as close as possible to the real-world properties of the cuprates. Before proceeding, however, we warn the reader that our numerical results apply strictly to purely two-dimensional systems, as any interlayer coupling affects the ensuing DOS at the Fermi energy [67].

For pointlike and random-site-energy models, we find that weak disorder—whether in the form of a low concentration of strong scatterers or a narrow distribution of on-site energies—leads to a vanishing DOS at the Fermi energy. It is only when unrealistic levels of disorder are reached that a finite DOS is generated, and even then there is an observed suppression at $E = 0$. We observe that the manner in which the *d*-wave gap “fills” differs depending on whether one has random-potential or unitary-scatterer disorder. With smooth disorder, however, a finite DOS at the Fermi energy is generated at fairly realistic concentrations (around 10-20%) and, strikingly, the overall structure of the *d*-wave DOS is preserved for all energies even at high dopings.

We also perform an exact calculation of the localization length λ and its dependence on the strength of disorder for the three different kinds of disorder we consider. We find that states near the Fermi energy are strongly localized for all three models—even for weak disorder—and that at intermediate and high energies *within* the *d*-wave gap the localization length is generally found to be very large for low disorder. It is

worth noting that even with a high concentration of smooth scatterers, the localization length at intermediate and high energies is still very large and comparable to that seen in much lower levels of disorder in the random-potential and unitary-scatterer case, indicating that localization effects due to smooth disorder are far weaker than in the case of pointlike disorder. Unitary scatterers in turn have a weaker effect on the localization length than random-potential disorder does.

Finally, we comment on the nature of disorder in the cuprates based on what is known from specific heat experiments, scanning tunneling spectroscopy, and numerical simulations. We caution the reader that a major limitation of our study is that the gap is not computed self-consistently, so we cannot ascertain with any definiteness whether the effects of disorder that we detail here are preempted by the destruction of d -wave superconductivity once some level of disorder is reached. Incorporating full self-consistency in the real-space numerical method we use is technically difficult, especially when the system size is large. This difficulty is a part of a tradeoff we make in order to access large system sizes. That said, exact-diagonalization studies on d -wave superconductors with unitary scatterers, using small system sizes, find that the superfluid density of the uniform-gap case and that of the self-consistent-gap case behave very similarly to each other, except when the concentration is sufficiently large [48]. T_c in turn was found to be much *less* suppressed in the self-consistent case than in the uniform-gap case. It was found that while in the uniform-gap case $p \approx 8.0\%$ almost completely suppresses T_c , in the self-consistent case such suppression occurs at nearly twice that level of disorder. This means that the uniform-gap picture in fact *overstates* the impact of disorder on the suppression of T_c and the superfluid density. This is augmented by the fact that, in other exact diagonalization studies, self-consistency does not fundamentally alter the structure of the DOS of the random-

potential and unitary-scatterer cases [13, 14, 194]. For certain parameter regimes it appears that the DOS for self-consistent and non-self-consistent order parameters are identical. In other regimes, the DOS is smoother and features more pronounced suppression near the Fermi energy in the self-consistent case than in the non-self-consistent one, while remaining similar to each other in other energy ranges. All of this suggests that what we find from our uniform-gap systems provides a good baseline for ascertaining the effects of site disorder on the cuprates, and very likely overestimates the pair-breaking effects of disorder. We defer a fully self-consistent treatment of these three kinds of disorder and their pair-breaking effects to a future publication.

4.2 METHODS

We start with a tight-binding Hamiltonian describing electrons hopping on a square lattice with d -wave pairing:

$$H = - \sum_{\langle i,j \rangle} \sum_{\sigma} t_{ij} c_{i\sigma}^{\dagger} c_{j\sigma} + \sum_{\langle i,j \rangle} \Delta_{ij}^* c_{i\uparrow} c_{j\downarrow} + \sum_{\langle i,j \rangle} \Delta_{ij} c_{i\uparrow}^{\dagger} c_{j\downarrow}^{\dagger}. \quad (4.1)$$

Nearest-neighbor and next-nearest-neighbor hoppings are both present, as is d -wave pairing, implemented by choosing the pairing amplitude to have the form $\Delta_{ij} = \pm\Delta_0$, where the positive (negative) value applies to pairs of nearest-neighbor sites along the x - (y -) direction. From the Hamiltonian, the Green's function takes the following expression:

$$G^{-1}(\omega) = \omega\mathbb{1} - H. \quad (4.2)$$

Note that H and G are $2N_x N_y \times 2N_x N_y$ matrices written in Nambu-space form, where N_x and N_y are the number of lattice sites in the x - and y -directions, respectively. From $G(\omega)$, various quantities can be

obtained. We will focus on the quasiparticle density of states and the localization length.

4.2.1 Quasiparticle Density of States

The quasiparticle DOS at energy E is

$$\rho(E) = -\frac{1}{\pi N_x N_y} \text{ImTr}G(E + i0^+). \quad (4.3)$$

Periodic and open boundary conditions are implemented in the y - and x -directions, respectively. To compute G , we first rewrite G^{-1} in the following block tridiagonal form:

$$\mathbf{G}^{-1} = \begin{pmatrix} \mathbf{P}_1 & \mathbf{Q}_1 & & \dots & & & \mathbf{0} \\ \mathbf{Q}_1^\dagger & \mathbf{P}_2 & \mathbf{Q}_2 & & & & \\ & \ddots & \ddots & \ddots & & & \\ \vdots & & \mathbf{Q}_{j-1}^\dagger & \mathbf{P}_j & \mathbf{Q}_j & & \vdots \\ & & & \ddots & \ddots & \ddots & \\ & & & & \mathbf{Q}_{N_x-2}^\dagger & \mathbf{P}_{N_x-1} & \mathbf{Q}_{N_x-1} \\ \mathbf{0} & & \dots & & \mathbf{Q}_{N_x-1}^\dagger & \mathbf{P}_{N_x} & \end{pmatrix}. \quad (4.4)$$

The \mathbf{P}_i blocks are $2N_y \times 2N_y$ submatrices and contain in their diagonal elements the frequency ω and the on-site energies at sites located on the i th slice of the system, where i runs from 1 to N_x , in addition to hopping and pairing amplitudes between sites within the i th slice. The \mathbf{Q}_i blocks—also $2N_y \times 2N_y$ submatrices—meanwhile contain hopping and pairing amplitudes from the i th slice to its nearest-neighbor slices. Note that the Nambu-space structure of the Green's function has been transferred to the \mathbf{P}_i and \mathbf{Q}_i blocks.

Because all we need is the trace of G to obtain the DOS, it suffices to obtain the diagonal blocks of G . For this purpose we use a block-by-block matrix-inversion algorithm that applies to block tridiagonal matrices [52, 69, 143]. We first define auxiliary matrices \mathbf{R}_i and \mathbf{S}_i in the following way:

$$\mathbf{R}_i = \begin{cases} \mathbf{Q}_i(\mathbf{P}_{i+1} - \mathbf{R}_{i+1})^{-1}\mathbf{Q}_i^\dagger & \text{if } 1 \leq i < N_x \\ \mathbf{0} & \text{if } i = N_x \end{cases} \quad (4.5)$$

and

$$\mathbf{S}_i = \begin{cases} \mathbf{0} & \text{if } i = 1 \\ \mathbf{Q}_{i-1}^\dagger(\mathbf{P}_{i-1} - \mathbf{S}_{i-1})^{-1}\mathbf{Q}_{i-1} & \text{if } 1 < i \leq N_x. \end{cases} \quad (4.6)$$

Once \mathbf{R}_i and \mathbf{S}_i have been computed, the i th diagonal block of G can be obtained straightforwardly from the following expression:

$$\mathbf{G}^{ii} = (\mathbf{P}_i - \mathbf{R}_i - \mathbf{S}_i)^{-1}. \quad (4.7)$$

We note that this procedure is exact and relies on no approximations. We set $N_x = 1000$ and $N_y = 100$ in all calculations.

To ensure the applicability of our numerical results to the cuprates, we use a band structure that is consistent with the details known about the normal-state Fermi surface of such materials: $t = 1$, $t' = -0.3$, and $\mu = -0.8$, where t , t' , and μ are the nearest-neighbor hopping, next-nearest-neighbor hopping, and the chemical potential, respectively. We note that our parametrization of the Fermi surface is limited as higher-order hopping amplitudes are not included, but this simple form of the band structure still captures the important general features of the Fermi surface of the cuprates. We choose the pairing amplitude to be $\Delta_0 = 0.08$; this choice gives $v_F/v_\Delta \approx 11$, in good

agreement with experiment [178]. (All energies are expressed in units where $t = 1$.) An inverse quasiparticle lifetime given by $\eta = 0.001$ is used throughout this work. This smears out the Dirac delta function peaks $\delta(E - E_n)$, where E_n is an eigenvalue of H , into a Lorentzian, $\frac{1}{\pi} \frac{\eta}{(E - E_n)^2 + \eta^2}$, whose full width at half maximum is 2η . Because the DOS of a clean d -wave superconductor with this particular band structure is nonzero up to energies $E \approx \pm 6t$, this choice of broadening roughly corresponds to introducing $\mathcal{O}(10^3)$ bins for the entire energy range. As there are 2×10^5 eigenvalues of the Hamiltonian, this provides more than adequate resolution for the examination of the DOS as a function of energy. Note that this value of η is parametrically much smaller than the energy resolution seen in scanning tunneling experiments (which are typically found to be 2 meV) [196]. Such values of the broadening already incorporate the effects of disorder, so in order to tease out the impact of disorder on the DOS we need to pick a much smaller value of η than seen in experiment.

The advantage of this particular method of obtaining the exact DOS, as opposed to similar methods such as the exact diagonalization of the Bogoliubov-de Gennes Hamiltonian, is threefold. First, this method is much faster in obtaining the DOS than exact diagonalization. As the DOS involves taking the trace of the Green's function, only the diagonal elements of G are needed, which are precisely the quantities outputted by the algorithm in use here. Second, this method can be extended to very large system sizes. The computational complexity depends only linearly on N_x , and consequently the size of that dimension can be increased without much trouble. Importantly, the large sizes that are accessible mean that the need to average over different disorder configurations is largely obviated—a single realization of disorder results in 10^5 values of the local density of states to be averaged over—and hence for the most part we will focus only on a single realization

of disorder for each of the cases we will consider. This makes much sense from a modeling viewpoint, especially as in experiment only *one* realization of disorder is present for a measurement. Finally, as finite-size effects are minimal, we are free to set the hopping and pairing parameters to correspond closely to those known from experiment. In exact diagonalization, the smallness of the system sizes typically used means that in order to visualize the spectrum fully one is occasionally faced with the need to make Δ_0 artificially large, so that within-gap physics are seen with the energy resolution available. In the method we use no such workarounds are necessary.

The only disadvantage of this method is that self-consistency is very difficult to implement in an efficient manner. In a fully self-consistent treatment the order parameter is iteratively determined via an integral of the anomalous Green's function over a range of energies. Consequently, in energy space the Green's function needs to be evaluated over a finely spaced array of points over the full bandwidth for the numerical integral to be accurate, and this process has to be repeated for an unspecified number of times until self-consistency is achieved. The full bandwidth is several times larger than the *d*-wave gap; hence the amount of computational effort required to perform this self-consistent calculation for even one realization of disorder becomes very large and uncontrollable. (This has to be contrasted with exact diagonalization, from which one obtains all the eigenvalues and eigenvectors of the Hamiltonian at once. The gap can then be computed in terms of the eigenvectors once one diagonalization has been completed. While this method is restricted to very small geometries, it is nonlocal in energy space, and thus implementing self-consistency is much easier.) As we have noted in the Introduction, evidence from previous numerical studies of lattice *d*-wave superconductors with strongly pair-breaking unitary scatterers suggests that self-consistent and non-self-consistent

results are not drastically different from one another. We will thus take the results from our uniform-gap systems to provide a reasonable account of the effects of disorder on the various quantities of interest to us.

It is also easy to obtain the *local* quasiparticle density of states (LDOS) from G . Because G is written in a real-space basis, the LDOS $\rho(\mathbf{r}, E)$ is simply

$$\rho(\mathbf{r}, E) = -\frac{1}{\pi} \text{Im}(G_{11}(\mathbf{r}, \mathbf{r}, E + i0^+) + G_{22}(\mathbf{r}, \mathbf{r}, E + i0^+)), \quad (4.8)$$

where G_{11} and G_{22} are the particle and hole parts, respectively, of the Nambu-space Green's function. At this point it is worth emphasizing the fact that, from the way we have defined them, these maps are *not* the same as the local density of states maps obtained from STS studies. The conductance maps obtained in STS experiments are proportional to the local *electron* density of states, which are taken solely from the electron part of the Green's function: $\rho_{\text{tunn}}(\mathbf{r}, E) = -\frac{1}{\pi} \text{Im}G_{11}(\mathbf{r}, \mathbf{r}, E + i0^+)$. In contrast, the quasiparticle DOS at energy E , as defined in Eq. 4.8, includes contributions from both the electron and hole Green's function. We will frequently show these maps to visualize the extent to which disorder affects the degree of inhomogeneity in the quasiparticle wavefunctions at a particular energy E .

We also calculate, for completeness, the quasiparticle DOS of a *clean* d -wave superconductor in order to provide a baseline from which one can examine the impact of disorder. Unlike the disordered case, we perform this calculation in momentum space. We use the formula

$$\rho(E) = \sum_{\mathbf{k} \in \text{BZ}} \delta(E - E_{\mathbf{k}}), \quad (4.9)$$

where $E_{\mathbf{k}}$ are the eigenvalues of the clean Hamiltonian, given for positive energies by

$$E_{\mathbf{k}} = \sqrt{\epsilon_{\mathbf{k}}^2 + \Delta_{\mathbf{k}}^2}. \quad (4.10)$$

Here $\epsilon_{\mathbf{k}} = -2t(\cos k_x + \cos k_y) - 4t' \cos k_x \cos k_y - \mu$ and $\Delta_{\mathbf{k}} = 2\Delta_0(\cos k_x - \cos k_y)$ are the normal-state dispersion and the gap function in momentum space, respectively. Only positive energies need to be considered because of particle-hole symmetry. For consistency with the real-space calculations of the disordered cases, we also broaden the delta functions that enter Eq. 4.9 into a Lorentzian with broadening $\eta = 0.001$. In our momentum-space calculations we discretize the first Brillouin zone into a grid with 4000×4000 points. This choice results in a smooth DOS as a function of E which is free from finite-size effects.

4.2.2 Specific Heat

The quasiparticle contribution to the specific heat C is easily derived from the density of states by means of the following equation [68],

$$C = 2 \times \frac{\partial}{\partial T} \int_0^\infty dE \rho(E) E \frac{1}{e^{E/k_B T} + 1}, \quad (4.11)$$

where the factor of two arises from the two spin species present. We are interested in C in the low-temperature regime, so we can neglect the dependence of $\rho(E)$ on T , and because $T \ll 4\Delta_0$ (the d -wave gap edge, which itself is much bigger than T_c) we can impose a cutoff $E_c \approx 4\Delta_0$ so that only energies within the d -wave gap are integrated over. As such, Eq. 4.11 becomes

$$C = 2 \times \frac{1}{k_B T^2} \int_0^{E_c} dE \rho(E) E^2 \frac{e^{E/k_B T}}{(e^{E/k_B T} + 1)^2}. \quad (4.12)$$

It can further be shown that the contribution of $\rho(E = 0)$ to the specific heat is

$$C_0 = \gamma_0 T = \frac{1}{3} \pi^2 \rho(E = 0) k_B^2 T. \quad (4.13)$$

When C_0/T is plotted versus T , the plot is flat, and the y -intercept of this plot is equal to γ_0 . In our numerical results we will typically set $k_B = 1$ and measure the temperature T in units of the hopping energy t ($t \approx 0.150$ eV ≈ 1700 K).

Note that the scaling of C with T is dependent on how ρ scales with E . At low energies the DOS of a clean d -wave superconductor is a linear function of E ; thus the quasiparticles of a clean d -wave superconductor contribute a T^2 -dependent term to C . When this coexists with a finite quasiparticle DOS at $E = 0$, the most general scaling of C due to the d -wave quasiparticles is

$$C \approx \gamma_0 T + \alpha T^2, \quad (4.14)$$

and a C/T -versus- T plot would have a slope equal to α and a y -intercept equal to γ_0 . In the most general disordered case we should not expect this form of scaling to arise, as disorder can lead to a non-linear dependence of ρ on E . However, a finite value of γ_0 is a feature that unambiguously suggests the presence of a finite DOS at the Fermi energy.

4.2.3 Localization Length

The geometry of our system is particularly amenable to exact calculations of the localization length λ , owing to the fact that N_x can be made very large relative to N_y , allowing us to measure the localization length even when it is much bigger than the transverse dimension.

This calculation is all but impossible using exact diagonalization, as that method is restricted to fairly small system sizes whose linear dimension is much smaller than typical localization lengths.

We will use the following definition of λ [106, 23, 92, 186]:

$$\lambda^{-1} = -\frac{1}{2(N_x - 1)} \ln \frac{\sum_{ij\sigma\sigma'} |G_{ij\sigma\sigma'}^{N_x 1}(E)|^2}{\sum_{ij\sigma\sigma'} |G_{ij\sigma\sigma'}^{11}(E)|^2}. \quad (4.15)$$

The $\frac{\sum_{ij\sigma\sigma'} |G_{ij\sigma\sigma'}^{N_x 1}(E)|^2}{\sum_{ij\sigma\sigma'} |G_{ij\sigma\sigma'}^{11}(E)|^2}$ factor measures the transmission probability from the left end of the system (the 1st slice) to the right end (the N_x th slice); the denominator in the aforementioned factor is for normalization. The sums are performed over all sites and spin indices within the relevant block. The off-diagonal block $G^{N_x 1}(E)$ can be recursively computed from the diagonal block $G^{11}(E)$ by an algorithm that applies to block tridiagonal matrices [52, 69, 143]. Using the \mathbf{P}_i , \mathbf{Q}_i , \mathbf{R}_i , \mathbf{S}_i , and \mathbf{G}^{ii} matrices obtained earlier, any off-diagonal blocks of G can be computed using this formula:

$$\mathbf{G}^{ij} = \begin{cases} -(\mathbf{P}_i - \mathbf{R}_i)^{-1} \mathbf{Q}_{i-1}^\dagger \mathbf{G}^{i-1,j} & \text{if } i > j, \\ -(\mathbf{P}_i - \mathbf{S}_i)^{-1} \mathbf{Q}_i \mathbf{G}^{i+1,j} & \text{if } i < j. \end{cases} \quad (4.16)$$

We calculate the localization length only for fixed values of N_x and N_y . We do not extract the actual localization length via finite-size analysis. We thus provide the necessary caveat that the values of λ that we cite here are meaningful only in comparison with systems with *identical* system sizes. That is, a direct comparison is possible between λ 's computed with the same N_x and N_y but for different disorder types and strengths, but *not so* when these system-size parameters are altered relative to one another.

4.3 MODELS OF DISORDER

In this chapter we will focus on three distinct models of disorder. Many of these forms of disorder have been discussed in the older literature on the subject, and in particular some of them can be treated, on some level, analytically in either the Born approximation or the T -matrix approximation. Here we will make use of the ability to simulate systems with very large system sizes to cover regimes where the approximations that enable analytical treatments of disorder fail. Below we will enumerate these models of disorder, their properties, and the degree to which these describe the actual disorder present in the cuprates.

4.3.1 *Random-Potential Disorder*

The first model is random and *spatially uncorrelated* on-site energies. We assume that the potential at each lattice site consists of two parts: the uniform chemical potential and a normally distributed random component V with zero mean and variance σ^2 :

$$\langle V(\mathbf{r}) \rangle = 0, \quad (4.17)$$

$$\langle V(\mathbf{r}_1)V(\mathbf{r}_2) \rangle = \sigma^2\delta_{\mathbf{r}_1\mathbf{r}_2}. \quad (4.18)$$

From the perspective of diagrammatic perturbation theory, this is a particularly tractable model of disorder: given the above conditions,

the Fourier transform of the two-point *averaged* correlation function of the disorder potential is a constant in momentum space:

$$\begin{aligned}
 W(\mathbf{k}) &= \sum_{\mathbf{r}} \langle V(\mathbf{r}) V(\mathbf{0}) \rangle e^{-i\mathbf{k}\cdot\mathbf{r}} \\
 &= \sum_{\mathbf{r}} \sigma^2 \delta_{\mathbf{r}\mathbf{0}} e^{-i\mathbf{k}\cdot\mathbf{r}} \\
 &= \sigma^2.
 \end{aligned} \tag{4.19}$$

This property of the model allows one to analytically obtain the self-energy easily using the Born approximation in the limit that σ is small [100]. Physically this model can be obtained from the multiple point-impurity model when one takes the strength of these impurities to be very weak and the spacing between impurities very small.

A related version of this disorder potential was studied numerically by Atkinson *et al.*; however they utilized box disorder instead of Gaussian distributions [14]. We on the other hand will focus exclusively on normally-distributed on-site energies. This form of disorder is physically realistic, as recent work has shown that narrowly-distributed Gaussian disorder of this sort could give rise to quasiparticle scattering interference (QPI) patterns in *d*-wave superconductors that are in reasonably good agreement with those seen in experiments on BSCCO [161].

4.3.2 Multiple Unitary Scatterers

The second model we will discuss is another paradigmatic form of disorder in the cuprates: unitary pointlike scatterers situated *within* the copper-oxide plane. Unitary scatterers in *d*-wave superconductors have been extensively studied experimentally and theoretically. Zinc dopants within the CuO_2 planes of BSCCO are the most well-known

studied form of unitary scatterers in the cuprates, and in fact their resonances have been directly imaged in STS experiments [129]. Unitary scatterers also arise in the cuprates in the form of vacancies within the CuO_2 plane. Like the Gaussian random-disorder case discussed earlier, unitary scatterers, which induce scattering phase shifts equal to $\delta_0 = \pi/2$, are quite tractable to model in practice: the T -matrix for a single pointlike impurity is momentum-independent, allowing one to obtain the full Green's function, including the impurity and its effects, in an exact manner. This can then be extended to the many-impurity case in the dilute limit (*i.e.*, at low concentrations p) in the form of a multiple-scattering T -matrix [68]. (Note that if one takes the strength of the impurities to be small, the phase shift is $\delta_0 \approx 0$, and the corresponding T -matrix problem becomes identical to the Born-scattering limit of the Gaussian random-potential case discussed previously [68, 100].)

We will eschew the T -matrix approach and instead obtain the full Green's function and the DOS exactly using the methods described in Section 4.2. This will allow us to examine cases where the concentration p is large enough that the system enters the strong-disorder regime. We will vary p to cover small, intermediate, and large concentrations; the strength of the impurity is fixed at $V_u = 10$, and we will make this potential attractive, to mimic the effect of zinc impurities, which are attractive potential scatterers [109, 93]. These impurities are distributed randomly over the entire system, with each lattice site having a p chance of hosting a unitary impurity and a $1 - p$ probability of not having one. Our choice of $V_u = 10$ gives a resonance energy at around $E \approx -0.06$ —the negative-bias peak in the bare electron LDOS at the sites adjacent to an isolated impurity is far more prominent than the positive-bias one—which is near, but not at, the Fermi energy. (To perform a sanity check, we checked the case of an isolated impurity

with $V_u = 100$, which yielded a resonance energy of $E \approx -0.045$. Increasing the impurity potential tenfold indeed pushed the resonance closer to the Fermi energy, but only by a small amount. In fact, if we do a single-impurity (*i.e.*, non-self-consistent) T -matrix calculation [195], assuming unitary scatterers with $V_u \rightarrow \infty$ and using the same band-structure and pairing details as in our exact numerical calculations, we find that the resonance is at $E \approx -0.04$. For generic band structures and arbitrary but strong V_u the resonance due to a strong, attractive scatterer is located close to, but not at, the Fermi energy, although for the purposes of this chapter its precise location is not very important.) Note that the effect of unitary scatterers on the DOS of d -wave superconductors has been studied by Atkinson *et al.* [13, 14], but we will go beyond their work by varying p such that both dilute and strong-disorder limits are covered, and by delving deep into the statistics of the DOS at the Fermi energy in considerable detail.

4.3.3 Smooth Disorder

The third and final form of disorder that we will discuss is off-plane disorder. As we have noted earlier, for the cuprates, disorder due to doping is generally due to dopants that are located some distance away from the CuO_2 planes. Doping in the cuprates is accomplished using oxygen atoms, and these oxygens are in general not found within the conducting planes. For BSCCO, the BiO planes host the excess oxygens arising from doping. In the case of YBCO, the doped oxygens are found in the one-dimensional CuO chains some distance away from the CuO_2 planes. YBCO is a particularly interesting case to consider because the amount of doping, and hence disorder, can be controlled rather precisely: very clean samples have been synthesized. Thermal

conductivity experiments on clean YBCO find that transport does not resemble either Born or unitary scattering (*i.e.*, the previous two models at low levels of disorder) [66]. Thus it is an interesting theoretical puzzle as to why precisely a finite DOS at the chemical potential is consistently found in specific heat studies of YBCO, even with clean samples.

We will attempt to revisit the effects of off-plane disorder on the quasiparticle DOS of a d -wave superconductor. Off-plane dopants will produce a screened Coulomb potential which affects the electrons on the CuO_2 plane in the form of a *smooth* disorder potential [130, 42, 22]. In the absence of a more microscopic model of disorder, we will take the disorder potential from one off-plane dopant located on the a - b plane at \mathbf{r}_n to have the following reasonably general form:

$$V_n(\mathbf{r}) = V_0 \frac{e^{-\frac{s(\mathbf{r}, \mathbf{r}_n)}{L}}}{s(\mathbf{r}, \mathbf{r}_n)}. \quad (4.20)$$

For brevity we have defined $s(\mathbf{r}, \mathbf{r}_n)$ as

$$s(\mathbf{r}, \mathbf{r}_n) = \sqrt{(\mathbf{r} - \mathbf{r}_n)^2 + l_z^2}, \quad (4.21)$$

and L is the screening length of the Coulomb potential, l_z is the distance along the c -axis from the dopant to the CuO_2 plane, and V_0 quantifies the “strength” of the potential. For our calculations we take $L = 4$, $l_z = 2$, and $V_0 = 0.5$. Because we do not exactly know the details of this disorder potential, we will assume two different scenarios for how this form of disorder is spatially distributed. For the first scenario, we will take the general disorder potential to have the *same*

sign, such that the net potential, expressed as a function of the doping concentration p , takes the following form:

$$V_s(\mathbf{r}) = \sum_{n=1}^{pN_x N_y} V_n(\mathbf{r}). \quad (4.22)$$

The second scenario assumes that there is an equal number of positive- and negative-strength potentials,

$$V_z(\mathbf{r}) = \sum_{n=1}^{pN_x N_y} (-1)^{a(n)} V_n(\mathbf{r}), \quad (4.23)$$

where $a(n)$ is a random integer. This leads to a potential whose spatial average is zero, and whose average over disorder configurations (*i.e.*, positions of the dopants, with the number of dopants held fixed) is also zero:

$$\langle V_z(\mathbf{r}) \rangle = 0. \quad (4.24)$$

The second scenario relies on a finely-tuned equality between the number of positive- and negative-strength dopants, and as such we do not claim that it necessarily corresponds to a realistic disorder potential. Nevertheless, from a theoretical standpoint V_z is a particularly interesting form of disorder because, like the Gaussian random-potential disorder case discussed earlier, its spatial and configuration average is zero. However $V_z(\mathbf{r})$ differs from the Gaussian case because it is *not* spatially uncorrelated: its disorder-averaged two-point correlator is not a delta function. Rather, this correlator decays much more slowly than a delta function. The length scales associated with this disorder potential drastically affect the allowed scattering processes. Recall that a d -wave superconductor has four nodes where gapless Bogoliubov quasiparticles exist at $E = 0$, which then morph into banana-

shaped contours of constant energy (CCEs) once energy is increased from zero. When one has elastic scattering off of pointlike impurities, there is no restriction on scattering processes aside from phase-space considerations: scattering has to occur between states lying on CCEs [182, 25, 93, 161]. With smooth disorder, however, the matrix elements of the potential vanish very quickly as momentum is increased, leading to a suppression of large-momentum scattering processes [126]. For this form of disorder, the dominant scattering processes occur only within one node, and to a first approximation scattering between states on different nodes can be neglected. This has been studied from the perspective of quasiparticle scattering interference, and smooth disorder potentials have been found to result in the marked suppression of large-momentum peaks in the Fourier-transformed LDOS [125, 161].

The distinction between pointlike disorder (*e.g.*, random normally-distributed on-site potentials and multiple unitary impurities) and smooth disorder is rarely discussed on a theoretical level. Prominent exceptions are the pioneering and extensive work by Nunner *et al.* on Coulomb-potential disorder [124, 126, 125], by Durst and Lee on extended linear scatterers [40], and field-theoretical work motivated by the possibility that scattering in the cuprate superconductors is primarily forward (*i.e.*, small-momenta) in nature [118]. It has been argued that, from the standpoint of effective field theory, the microscopics of the disorder determine the symmetry class of the effective theory of the disordered system, and consequently pointlike and smooth disorder belong to different universality classes [10]. While this does make sense from this particular viewpoint, from a more microscopic perspective such as ours such a distinction is not as clear-cut: one can, at least in principle, continuously tune the length scales of the disorder potential to come close to the pointlike limit, so it is difficult to argue that the lattice tight-binding Hamiltonian exhibits such a sharp dis-

inction between two different universality classes. There is also a difficulty in extending these field-theoretical results to the intermediate- and strong-disorder regimes, as these take as a starting point the presence of *weak* disorder. Nevertheless, as we shall see with our numerics, smooth disorder does lead to effects that differ dramatically from either random Gaussian disorder or multiple-impurity models.

The main variable we use to manipulate the amount of disorder in the superconductor is the concentration p of off-plane dopants. To be more specific, p here is the number of off-plane dopants per copper site *at* the CuO_2 plane. From what is known about LSCO, BSCCO, and YBCO, p is generally a large fraction which is usually of the order of $p \approx 0.1-0.2$. The precise doping level of YBCO is a complicated quantity to determine because it is not at all obvious how many of the oxygen dopants go to the chains and to the planes; we will not incorporate these subtleties in our calculations, but we do note that microwave conductivity measurements on YBCO are generally found to be consistent with a concentration of defects on the CuO chains given by $p \approx 0.1$ [22]. We will cover this regime of doping, as this is the most physically relevant one, although we will cover low and high concentrations as well. It is not clear *a priori* whether a density of $p \approx 0.1-0.2$ corresponds to weak or strong disorder, so we will scan through p to see precisely what regimes are covered by these impurity concentrations.

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW

We now discuss our numerical results for the quasiparticle density of states. We first focus on random-potential disorder. Fig. 4.1 shows the quasiparticle DOS as a function of energy for various values of σ .

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW




Figure 4.1: Plots of the quasiparticle DOS as a function of energy E for the Gaussian random-potential model, for various values of σ .

There are a number of interesting features in these plots that are worth mentioning. We focus first on the DOS near $E = 0$. For small values of σ (*i.e.*, $\sigma = 0.125$ and $\sigma = 0.25$), the DOS vanishes markedly at $E = 0$. For these cases the DOS scales roughly linearly with E near $E = 0$. The weakest disorder distribution we consider ($\sigma = 0.125$) has a DOS curve that is concave upward between $E = 0$ and the coherence peaks. This changes for $\sigma = 0.25$, for which the DOS is almost perfectly linear from zero energy up to the coherence peaks, and from $\sigma = 0.35$ upwards the DOS curves are all *concave downward*. At $\sigma = 0.35$ and $\sigma = 0.50$, a finite DOS at $E = 0$ is generated, but despite this offset the DOS still scales approximately linearly with E . For higher values of σ , the DOS at the Fermi energy is still finite, but there is a very visible dip around $E = 0$ relative to nearby energies. In the strong-disorder regime, the DOS scales linearly with E only within a small neighborhood of $E = 0$, then becomes dramatically concave downward as energies increase.




Figure 4.2: Snapshots of the real-space quasiparticle density of states for random Gaussian disorder with increasing standard deviation σ (top to bottom) and energy E (left to right), extracted from the middlemost 80×80 subset of the full system. The leftmost column shows plots of the DOS as a function of energy for a particular σ , along with plots of the clean case for comparison. The same disorder realizations as in Fig. 4.1 are used here. The color scale is the same for all plots.

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW




Figure 4.3: Plots of the quasiparticle DOS as a function of energy E for the multiple unitary-scatterer model, for various impurity concentrations.

At $E \approx 0.3$, one can see the coherence peaks becoming more rounded and decreasing in height with increasing σ . With relatively weak disorder, the peaks retain their prominence, but as disorder becomes stronger these peaks flatten. In fact, for the strongest disorder cases we consider ($\sigma = 1.41$ and $\sigma = 2.00$) the DOS near (but not at) $E = 0$ barely differs from the DOS at $E \approx 0.3$. For energies between $E = 0$ and $E \approx 0.3$, the slope of the DOS decreases with increasing σ . The overall effect of increasing disorder of this kind is to shift spectral weight away from the coherence peaks towards a broad range of low and intermediate energies, consequently filling in the d -wave gap.

Qualitatively there are three distinct regimes that are encountered as random on-site disorder is increased. At low values of σ , the superconductor is only weakly disordered: the DOS vanishes at $E = 0$ and coherence peaks are prominent. At intermediate values of σ , a finite




Figure 4.4: Snapshots of the real-space quasiparticle density of states for an ensemble of unitary pointlike scatterers ($V_U = 10$) with increasing impurity concentration p (top to bottom) and energy E (left to right), extracted from the middlemost 80×80 subset of the full system. The leftmost column shows plots of the DOS as a function of energy for a particular p , along with plots of the clean case for comparison. The same disorder realizations as in Fig. 4.3 are used here. The color scale is the same for all plots.




Figure 4.5: Plots of the quasiparticle DOS as a function of energy E for the multiple smooth-scatterer model with *positive* net potential, for various impurity concentrations.

value of the DOS forms at the Fermi energy, but the DOS still varies linearly with E over a broad energy range, and traces of the coherence peaks (now rounded and diminished in height) still remain. Finally, when σ is large, we enter the strong-disorder regime, where the DOS is linear only within a small neighborhood of $E = 0$ and saturates very quickly to a constant value (albeit with considerable random fluctuations about that value). The DOS is suppressed at $E = 0$ relative to the value to which it eventually saturates, and in fact tends toward zero once more as disorder is increased. In this regime almost no trace of the structure of the DOS of the clean d -wave superconductor remains.

To closely examine the origins of both the generation of a finite DOS at $E = 0$ and the smoothing of the coherence peaks, we extract real-space maps of the *quasiparticle* local DOS (LDOS) for various disorder strengths and energies. We take these samples from the middle 80×80 section of the full system. These maps are shown in Fig. 4.2.

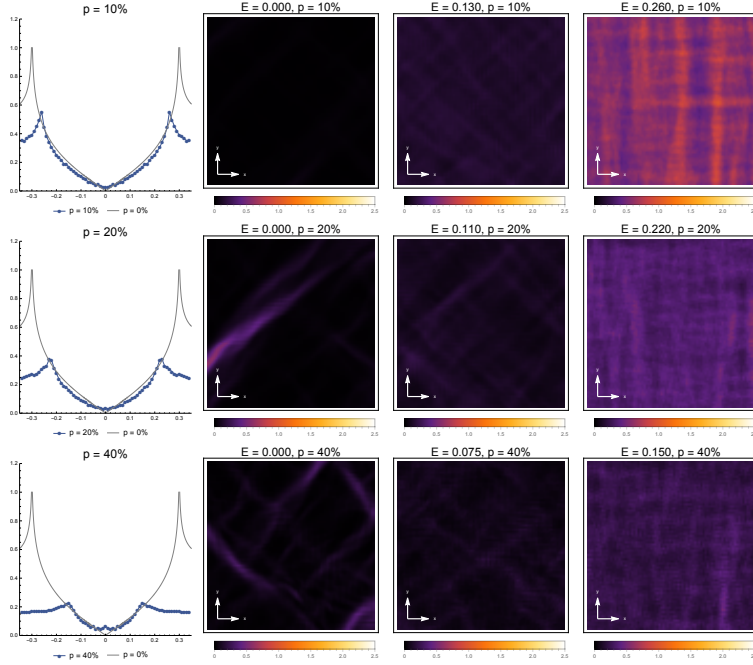


Figure 4.6: Snapshots of the real-space quasiparticle density of states for smooth disorder (with *positive* net potential) with increasing impurity concentration p (top to bottom) and energy E (left to right), extracted from the middlemost 80×80 subset of the full system. The energy at the rightmost column corresponds to the location at which the coherence peaks can be found, while the energy at the middle column is half the coherence-peak energy. The leftmost column shows plots of the DOS as a function of energy for a particular p , along with plots of the clean case for comparison. The same disorder realizations as in Fig. 4.5 are used here. The color scale is the same for all plots.

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW




Figure 4.7: Histogram of the values of the disorder potential for smooth disorder with positive net potential for three values of p . The width of each bin is 0.01. Notice that the mean of the disorder potential is nonzero, leading to a shift in the average chemical potential of the overall system.

At $E = 0$, the weak-disorder ($\sigma = 0.25$) LDOS is almost zero and is spatially featureless. When disorder is increased, regions where the LDOS is nonvanishing form even at $E = 0$. At moderate levels of disorder ($\sigma = 0.50$) these regions tend to be isolated, surrounded by a sea of vanishing DOS. These are sufficient however to produce a finite DOS when averaged over the entire system. When disorder is tuned to be strong ($\sigma = 1.00$), the LDOS map at $E = 0$ displays considerable randomness: patches where the LDOS vanishes coexist with regions where the DOS is visibly nonzero, thereby resulting in a nonzero average DOS.

As energies are increased the $\sigma = 0.25$ maps start exhibiting modulations in the LDOS that arise from quasiparticle interference in the presence of weak disorder. As disorder is increased, this structure becomes less and less visible: the $\sigma = 1.00$ maps at $E = 0.150$ and $E = 0.300$ show randomness that is not much different than the maps obtained at $E = 0$. The strong-disorder maps show at higher energies similar structures as the zero-energy case, with regions where the LDOS is




Figure 4.8: Plots of the quasiparticle DOS as a function of energy E for the multiple smooth-scatterer model with *zero* net potential, for various impurity concentrations.

heavily suppressed existing alongside areas with nonzero DOS. The presence of these patches where the LDOS is almost zero at large σ is responsible for the overall suppression of the averaged DOS relative to less disordered cases.

We repeat this analysis for the unitary-scatterer disorder model. For this form of disorder we show the quasiparticle DOS as a function of energy E in Fig. 4.3. When a small number of impurities are present (*e.g.*, $p = 0.125\%$), the DOS is barely altered from the clean case: the DOS tends toward zero at $E = 0$, increases linearly for a broad energy range, and displays sharp coherence peaks at $E \approx 0.300$. The same behavior holds for higher concentration of levels such as $p = 0.25\%$ and $p = 0.50\%$. We can see that the coherence peaks become slightly lower for these cases.

A major feature of these plots for a broad range of p is the rounding off of the DOS at an energy scale that appears to be dependent on the

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW




Figure 4.9: Snapshots of the real-space quasiparticle density of states for smooth disorder (with *zero* net potential) with increasing impurity concentration p (top to bottom) and energy E (left to right), extracted from the middlemost 80×80 subset of the full system. The leftmost column shows plots of the DOS as a function of energy for a particular p , along with plots of the clean case for comparison. The same disorder realizations as in Fig. 4.8 are used here. The color scale is the same for all plots.




Figure 4.10: Histogram of the values of the disorder potential for smooth disorder with zero net potential for three values of p . The width of each bin is 0.01. The mean of the disorder potential is zero, and the average chemical potential of the system as a whole is not shifted.

concentration. Near $E = 0$, the DOS scales linearly. As p is increased, the d -wave gap fills in a particular manner: more spectral weight accumulates at a characteristic energy scale, so that instead of a linear DOS as in the clean case, one sees the DOS encountering a “hump” that becomes more pronounced when p is increased. With increasing p the DOS surrounding $E = 0$ starts accumulating larger values of DOS, all while the coherence peaks become shorter and flatten, showing a transfer of spectral weight from the coherence peaks towards the region around the Fermi energy. It is interesting to note that the way the gap is filled is different for the case of unitary scatterers than for random on-site disorder: for small p , spectral weight is moved from the coherence peaks towards the neighborhood of the Fermi energy, with a width roughly set by the impurity concentration, whereas for random Gaussian disorder the spectral weight is transferred to a far broader range of energies, with strong deviations from the clean case occurring even at energies away from $E = 0$. For higher values of p , the DOS resembles the large- σ random-disorder cases discussed ear-

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW




Figure 4.11: Plot of the density of states at and near $E = 0$ for the multiple-smooth-scatterer case for various impurity concentrations p . For the $p = 20\%$ and $p = 40\%$ the resonance is seen to have a width of approximately 0.006.

lier. One feature that is consistently present—even at high values of p , with coherence peaks completely flattened and the DOS near the Fermi energy finite—is a visible dip *at* $E = 0$.

Real-space maps of the LDOS for a d -wave superconductor subject to a variety of unitary-impurity concentrations are shown in Fig. 4.4. At $p = 1.0\%$, the $E = 0$ LDOS map is largely almost zero, save for small areas that show large, nonzero values of the LDOS. A closer examination shows that these arise from interference effects from the presence of a few impurities bunched up together within a small area, arranged together such that a resonance forms. These resonances are very rare—in the 80×80 map we take, only one particular group of closely-spaced impurities generates such nonzero LDOS values at $E = 0$, whereas groups of a few impurities near one another do appear quite frequently. Despite their relative rarity, the presence of such regions with large average LDOS is enough to produce a small but nonzero average DOS for the entire sample. When the concentration is increased, we see behavior in the $E = 0$ maps that is strongly reminiscent of that seen in the maps from the Gaussian random disorder case. At $p = 4.0\%$,




Figure 4.12: Plot of the density of states at $E = 0$ for the multiple-smooth-scatterer case as a function of screening length L . For a given p , the positions of the smooth scatterers are fixed, with only the screening length and the amplitude of the disorder potential adjusted as discussed in the text. At fixed p the zero-energy DOS increases monotonically with L . In addition, at fixed L the DOS at $E = 0$ increases with increasing p .

regions where the LDOS is nonzero appear more frequently, but they are isolated and are largely surrounded by areas where the LDOS is suppressed. The $p = 16.0\%$ case shows a remarkably large number of lattice sites with large values of the LDOS. Clearly in this case the large impurity concentration means that there is a large probability that an impurity is placed in close proximity to another impurity, resulting in a nonzero LDOS.

At higher energies the $p = 1.0\%$ and $p = 4.0\%$ cases show modulations that are due to quasiparticle scattering interference (QPI) from multiple impurities. In particular the $p = 1.0\%$ map at $E = 0.300$ shows strikingly prominent modulations in the LDOS due to the presence of disorder; the $p = 4.0\%$ map at the same energy also shows visible modulations, but the larger number of impurities results in an average DOS that is lower than the $p = 1.0\%$ case. The $p = 16.0\%$ case, on the other hand, shows almost no visible traces of patterns arising

4.4 QUASIPARTICLE DENSITY OF STATES: AN OVERVIEW




Figure 4.13: Smooth disorder potentials used in Fig. 4.12, with concentration $p = 10\%$, evaluated at various screening lengths L . Shown are $L = 1$, $L = 2$, $L = 4$, and $L = 8$. The color scale is the same for all plots. Notice that as L is increased the disorder potential becomes smoother and more spatially correlated.

from QPI. Instead what one sees is a very inhomogeneous map featuring both sites with very strong suppression of the LDOS and sites at which the LDOS is large. For this particular concentration, the degree of inhomogeneity does not change markedly upon increasing E .

The suppression of the DOS at $E = 0$ for both random-potential and unitary-scatterer disorder has been discussed at length by Senthil and Fisher with field-theoretic methods [155] and by Yashenkin *et al.* using diagrammatic techniques incorporating weak-localization corrections to the T-matrix results [187, 67]. This suppression—found to be logarithmic in both approaches—can be understood as being due to the inclusion of diffusive modes that, in the absence of symmetries other than spin rotation invariance, lead to an overall suppression of the DOS. Yashenkin *et al.* also find that the addition of artificial nesting symmetries (*e.g.*, a particle-hole-symmetric normal-state band structure in the presence of unitary scatterers) can lead to additional diffusive modes that *enhance* the DOS at the Fermi energy. It is interesting to note that even in strong-disorder regimes where these approximations do not hold—diagrammatic and field-theoretical treatments both implicitly rely on a relatively narrow distribution of disorder for them to

be sensible—this logarithmic suppression at the Fermi energy is still very much evident for both random-potential and unitary-scatterer disorder.

We finally discuss the case of smooth disorder. We first focus on the case where the dopants have the same sign of the impurity strength—*i.e.*, the full potential is given by Eq. 4.22. Fig. 4.5 shows the quasiparticle DOS for a d -wave superconductor with such disorder, for various doping concentrations p . The behavior of the DOS near $E = 0$ has a number of interesting features when p is increased. First, at low p , the DOS is close to zero. As p is increased, the DOS gradually acquires a finite value, and at higher concentrations ($p = 20\%$ and $p = 40\%$) the DOS has a small bump at $E = 0$ relative to the value of the clean DOS. The neighborhood of the Fermi energy shows a gradual rounding of the DOS from a sharp V-shape in the clean and mildly disordered cases to a smooth U-shape for higher impurity concentrations. For all p , coherence peaks are present and quite prominent, but these shorten and move towards the Fermi energy as p is increased. This can be attributed to the fact that for this particular form of disorder, the mean of the disorder potential is nonzero, and the chemical potential is shifted away—only slightly for lower p , and considerably more strongly for larger and larger p , as seen in Fig. 4.7. It is interesting to note that despite the fact that this form of potential seemingly represents a strong modification to the d -wave superconductor, the effect is mainly to transfer spectral weight from the coherence peaks to the Fermi energy, with a corresponding rounding of the DOS, without impacting the DOS that much in the intermediate-energy regimes. There is also no visible suppression at $E = 0$, as was the case in the pointlike disorder models we discussed earlier. It seems that the overall effect of this particular form of disorder, at least as the quasiparticle DOS is concerned, is qualitatively *much* weaker than the random Gaussian

on-site energy and the multiple unitary-scatterer models at roughly similar disorder widths or impurity concentrations.

Real-space plots are shown in Fig. 4.6. The plots at $E = 0$ show how a nonzero DOS is generated in the neighborhood of the Fermi energy. At $p = 10\%$, the effect is only mild, as the LDOS is almost spatially uniform. With increasing concentration visible patterns start to show up in the LDOS maps. These patterns are interesting because they correspond to only a small portion of the entire system, but do generate, upon averaging over space, an overall nonzero DOS centered around $E = 0$. Unlike similar maps for the pointlike disorder cases, the patterns—which manifest themselves as streaks of nonzero DOS amid a featureless, almost-zero background—display a smoothness that is not present in the highly disordered pointlike cases. While displaying patchiness, it exhibits spatial variations that are much more ragged than in the smooth case. Meanwhile the maps taken at higher energies show crisscrossing patterns which arise naturally from quasiparticle interference due to scattering off of a highly random smooth disorder potential. Unlike the maps showing pointlike disorder, the modulations here are much smoother, owing to the fact that these arise from small-momenta scattering processes.

We next turn to the case where there is an equal number of positive- and negative-strength dopants—*i.e.*, the disorder potential shown in Eq. 4.23. This will prove to be a much more interesting case than the smooth-disorder scenario we had just discussed. We show plots of the DOS for this disorder potential in Fig. 4.8. A number of remarkable features are present in these plots which we will now discuss in detail. We focus first on the region around $E = 0$. At low p , the DOS vanishes, but at $p = 10\%$ the DOS acquires a value that is appreciably larger than that of the clean or low-doping cases. At this doping the DOS at $E = 0$ has a slight upward hump, and the DOS surrounding the Fermi

energy has a U-shape and is considerably rounded off compared to the shape of the clean DOS. At higher dopings, a very prominent spike in the DOS at $E = 0$ start to form: this spike is localized at $E = 0$, and falls off quickly towards the base of a “valley.” It can be seen that the area around the Fermi energy hosts a considerable amount of spectral weight relative to the clean case as p is increased.

These effects near the Fermi energy are far more pronounced because elsewhere there are no significant deviations from the clean DOS. Even for very large dopings (*e.g.*, $p = 40\%$), the DOS at intermediate and high energies are almost unchanged from that of the clean case. The main significant change at these energy ranges happens at the coherence peaks ($E \approx 0.3$), which become shorter and more rounded with increasing disorder. However the rounding and shortening are nowhere near as pronounced or as strong as those in the random-potential or unitary-scatterer cases. Recall that in these other cases, the coherence peaks are destroyed at some level of disorder ($\sigma \approx 0.5$ for random potential disorder, and $p \approx 8\%$ for unitary scatterers). However, even at $p = 40\%$ doping, smooth disorder preserves coherence peaks. More emphatically, the global structure of the d -wave DOS is preserved even for very large dopings.

This is remarkable given how randomly distributed the disorder potential is. This can be seen in histograms of the disorder potential values for this particular form of smooth disorder, which we show in Fig. 4.10. One can see that they are almost normally distributed, with widths not far off from the weaker incarnations of the random-potential case we discussed earlier. The difference of course lies in the presence of spatial correlations in the smooth disorder potential, which are completely absent for pointlike disorder. Evidently, unlike random-potential or unitary-scatterer disorder, which show dramatic spectral-weight transfers from the coherence peaks to a broad range

of energies, for this particular form of smooth disorder only moderate spectral weight transfer occurs, with the bulk accumulating near the Fermi energy and almost none in intermediate-energy regimes.

The $E = 0$ maps in Fig. 4.9 show how a spike in the average DOS is generated. At low p , few if any streaks are visible, and these faint streaks occur against a background where the LDOS is heavily suppressed. As p increases, more of these streaks are visible, and in the $p = 40\%$ case these streaks are strong enough that averaging over the LDOS yields a finite value. The $E = 0.150$ maps show, as in the other smooth-disorder case we studied, diagonal crisscrossing patterns that can be attributed to quasiparticle scattering interference. Note that the modulations in real space are slowly varying, which as before can be attributed to the fact that, in this disorder scenario, nearly all scattering is forward. The fact that mostly diagonal streaks can be seen is due to the fact that scattering occurs heavily within one node only, and the only \mathbf{q} -vector corresponding to such intranodal scattering is \mathbf{q}_7 , which is diagonal and small. At the coherence-peak energies ($E = 0.300$), the diagonal streaks are now mainly replaced by modulations in the vertical and horizontal directions—a reflection of the fact that these LDOS maps are still heavily determined by quasiparticle scattering interference. At this energy regime the vertical/horizontal momentum \mathbf{q}_1 becomes most dominant, leading to the prominent modulations in the horizontal and vertical directions. The maps at higher energies show a remarkable degree of similarity with each other, despite vastly different amounts of doping, indicating that the transfer of spectral weight away from these energies is largely muted. This is very different from what we have seen for random-potential or unitary-scatterer disorder.

The origin of the sharply enhanced DOS at $E = 0$ is unknown, but we will try to characterize this effect as fully as possible numerically.

First of all, the resonances are sharply located at $E = 0$, and are very narrow. Fig. 4.11 shows a close-up view of the DOS within a small window of the Fermi energy. We find that the resonances, which are most visible at $p = 20\%$ and $p = 40\%$, have a width of $\Delta E \approx 0.006$ centered about $E = 0$, and that these subsequently plateau into a flat profile a short distance away from the Fermi energy. From our numerical results it appears that these zero-energy resonances are uncorrelated with the underlying smooth disorder potential. It is an intrinsically many-impurity effect, since results from single-impurity simulations do not show a sharp spike in the local DOS at zero energy. It also depends rather sensitively on the length scales associated with the smooth disorder potential. In Fig. 4.12 we plot the DOS at $E = 0$ versus the screening length L for three different impurity concentrations p , keeping the positions of the impurities at a given p fixed. In these plots we change V_0 as L is varied in Eq. 4.20 so that $V(\mathbf{r} = \mathbf{0})$ remains the same for all values of L we consider. This choice ensures that the resulting smooth disorder potentials feature the same degree of spatial variations, even as L is varied. As we have seen in the $L = 4$ case heavily discussed earlier, at fixed L the $E = 0$ DOS depends on p , with the DOS increasing as p is increased. More remarkably, however, we can see that at fixed p , the zero-energy DOS increases monotonically as L is increased. This is interesting because at face value the smooth-disorder potentials at various L appear to be very similar to each other. This is seen in Fig. 4.13, which shows the different smooth disorder potentials used at fixed $p = 10\%$. These are similar in appearance, but evidently lead to considerable differences in the values of $\rho(E = 0)$. This suggests that the range of the potential plays an important role in the emergence of these resonances at zero energy.

We note that a mechanism for the *enhancement* of the DOS at $E = 0$ was discussed by Yashenkin *et al.*, who point out that diffusion modes

due to additional symmetries could lead to an increase in the DOS at $E = 0$ [187, 67]. It is not clear at all if this mechanism has any relation with the real-space streaks which generate the spike at the Fermi energy in our numerics. It was argued that symmetries such as particle-hole symmetry in the normal state lead to this enhancement; however, the normal-state band structure we use does not have any special symmetries, so this cannot explain this phenomenon. It should be noted too that Yashenkin *et al.*'s analysis relies on pointlike scatterers treated within a self-consistent T -matrix approximation, which does not describe the smooth disorder potentials which generate the enhanced DOS at $E = 0$. It is thus an interesting, if possibly very difficult, problem to apply the analysis of Yashenkin *et al.* to smooth impurity potentials. Treating smooth disorder analytically is a formidable challenge, unlike random-potential and unitary-scatterer disorder, and tractability is generally possible only in the nodal approximation, at which the Born or T -matrix approximations can be used. We will thus leave an explanation of these strong zero-energy enhancements of the DOS due to smooth disorder as an open problem.

4.5 CORRELATION BETWEEN THE LDOS AND THE DISORDER POTENTIAL

As discussed earlier, the behavior of the real-space LDOS varies as the amount of disorder is increased, with low-disorder cases exhibiting more visible modulations in the LDOS that are due to QPI. At high energies these modulations follow closely the details of the disorder potential. As disorder is increased, these modulations become less prominent. We can get some insight into how “strong” the disorder in the system is by computing the coefficient of correlation $R(E)$ between




Figure 4.14: Plot of the correlation coefficient R between the local density of states in the middlemost 80×80 patch of the system and the disorder potential in that region for different types of disorder, for varying disorder strength, as a function of energy. For all three plots the correlation coefficient is negative—that is, there is an overall *anticorrelation* between the LDOS and the disorder potential.

the local density of states at energy E and the disorder potential. $R(E)$ is defined in the following manner:

$$R(E) = \frac{\sum_{ij}(V(i,j) - \bar{V})(\rho(i,j,E) - \overline{\rho(E)})}{\sqrt{(\sum_{ij}(V(i,j) - \bar{V})^2)(\sum_{ij}(\rho(i,j,E) - \overline{\rho(E)})^2)}}. \quad (4.25)$$

Here $V(i,j)$ is the disorder potential at site (i,j) , $\rho(i,j,E)$ is the quasi-particle DOS at site (i,j) and energy E , and \bar{V} and $\overline{\rho(E)}$ are the average values of the disorder potential and the DOS, respectively, over the area where we perform the calculation. We compute R between the middlemost 80×80 LDOS patch of the system at energy E and the disorder potential in that same patch of the system. Plots of $R(E)$ are shown in Fig. 4.14. This is motivated by a similar analysis performed by McElroy *et al.* on experimentally-obtained LDOS data from BSCCO; they find that there is moderate *anticorrelation* between the locations of the dopant defects and LDOS minima [111]. Our analysis differs from theirs in that we know the details of the disorder potential directly, and the cross-correlation is between the potential and the LDOS, not between the impurity location and the LDOS.

In the case of random-potential disorder, what we find is that the LDOS is only moderately anticorrelated with the disorder potential, even for weak disorder. When $\sigma = 0.25$, R decreases from a small value ($R \approx -0.2$) until it saturates at $R \approx -0.5$ at $E \approx 0.25$, indicating that the high-energy LDOS displays more similarity with the underlying disorder potential than the low-energy LDOS. As σ is increased, the LDOS and the disorder potential become even less anticorrelated. $R(E)$ at $\sigma = 0.5$ shows only a moderate degree of dependence on energy, and at $\sigma = 1.00$ $R(E)$ is almost energy-independent and has a small value, indicating that the two variables are only weakly anticorrelated.

For multiple unitary scatterers the situation becomes markedly different. The $R(E)$ obtained for the $p = 1.0\%$ case exhibits a very visible dependence on energy. At low energies the LDOS is very weakly anticorrelated with the disorder potential, but this anticorrelation increases sharply as energy is increased, a sign that higher-energy LDOS maps match the features of the disorder potential more than the lower-energy maps do; for instance, $R \approx -0.6$ at $E \approx 0.25$. This trend is even noticeable once E is increased past the d -wave gap edge, where it can be seen that R continues to be more and more anticorrelated with increasing E . This behavior can be seen to a good extent in the $p = 4.0\%$ case, for which R shows a similar degree of energy-dependence in the intermediate- and high-energy ranges as in the $p = 1.0\%$ case. The $p = 16.0\%$ case is interesting, as in that case R is much less energy-dependent than in the cases involving lower concentrations, similar to the strong-disorder ($\sigma = 1.00$) case of the random-potential model, but the overall coefficient indicates that *stronger* anticorrelation is present between the two variables. This can be explained by the fact that unitary pointlike scatterers suppress the LDOS *at* the impurity sites, which contributes to the overall anticorrelation between the LDOS and the disorder potential.

The smooth-disorder cases feature behavior that is starkly different from the random-potential or unitary-scatterer models. For one, we obtain strongly energy-dependent $R(E)$ at all concentrations we consider (10%, 20%, and 40%). In addition, the behavior of R does not appear to vary as p is altered. At low energies, there is almost no anticorrelation between the LDOS and the disorder potential, but the anticorrelation sharply increases as E is increased. R reaches very large values at high energies—for instance, $R \approx -0.7$ at $E \approx 0.3$ —and in these regimes the LDOS maps bear a remarkable resemblance to plots of the smooth disorder potential, with regions where the LDOS is suppressed coinciding

with patches at which the disorder potential is positive and vice versa. Another interesting aspect is that strong fluctuations in R exist, independently of p . This is unlike random-potential or unitary-scatterer disorder, for which we saw that the fluctuations in R are minimal. It is important to note from these plots that the resonances in the DOS at $E = 0$ are almost completely *uncorrelated* with the disorder potential—the origin of these resonance streaks at zero energy appears not to originate from local features of the disorder potential.

4.6 PROPERTIES OF THE DENSITY OF STATES NEAR $E = 0$

As a considerable number of properties of the cuprate superconductors rely on the physics of the low-energy quasiparticles near the Fermi energy, we will examine more closely the behavior of the DOS near $E = 0$ as disorder is increased. We have seen that, in the random-potential and unitary-scatterer models of disorder, when the amount of disorder is increased, $\rho(E = 0)$ acquires a finite value, then drops once more towards zero after a certain disorder strength is reached. To see if this behavior is robust, we show in Fig. 4.15 plots of the mean and standard deviation of $\rho(E = 0)$ as the amount of disorder is increased for each of the four models of disorder we use, with five realizations used per value of the disorder strength parameter. All in all, a total of 500,000 LDOS values for each value of the disorder strength parameter are used to generate this plot. A similar if considerably more detailed analysis of LDOS distributions on the Anderson model was performed by Schubert *et al.* in order to obtain criteria for Anderson localization using finite-size scaling [152]. We will not repeat their finite-size analysis here. It should be noted that, under certain conditions, information about the distribution of the LDOS at $E = 0$ can be extracted by ob-




Figure 4.15: Plots of the mean and standard deviation of the quasiparticle *local* DOS at $E = 0$ for different types of disorder. Five realizations are utilized for each value of the disorder strength parameter for each type of disorder; an average over 5×10^5 values of the local DOS is taken.

4.6 PROPERTIES OF THE DENSITY OF STATES NEAR $E = 0$

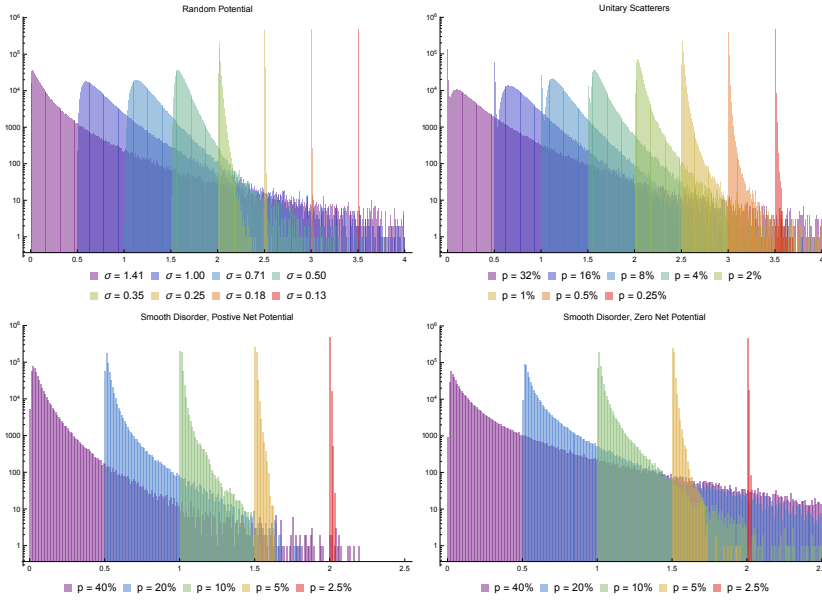


Figure 4.16: Histogram of the *local* DOS at $E = 0$ for different types of disorder. A logarithmic scale is used for the y -axis. Five realizations are used per value of disorder strength parameter for each type of disorder. 5×10^5 values of the LDOS for each value of disorder parameter are shown here. To show the variation in the behavior of the distributions for various disorder strengths, the histogram for a particular value of disorder strength is offset from the preceding one. The bin width is 0.01.

taining the ^{17}O Knight shift values from nuclear magnetic resonance experiments [128, 192, 193]. In particular, Zhou *et al.* find an asymmetric distribution of Knight shifts in YBCO with charge order, which suggests that the LDOS at the Fermi energy is distributed similarly, and argue that a likely explanation of this is quasiparticle scattering off of defects [192].

Let us discuss first the random-potential model. In the weak-disorder regime, the mean and standard deviation of the DOS are both close to zero and exhibit almost no dependence on σ . Starting at approximately $\sigma = 0.35$ the mean DOS becomes finite, increasing as σ is increased, and, more interestingly, the standard deviation of $\rho(E = 0)$ depends strongly on the value of σ . This trend continues until $\sigma = 1.00$: as disorder is increased past that point, the mean DOS starts to decrease, while the standard deviation continues to increase until $\sigma = 1.4$ is reached. In these strong-disorder regimes, the way that $\rho(E = 0) \rightarrow 0$ is of a fundamentally different nature than the way the weak-disorder DOS tends toward zero: the distribution of the strong-disorder DOS, while heavily weighted towards zero, exhibits very large spatial variations. The weak-disorder case on the other hand is almost fully concentrated at zero, with almost negligible variations in space.

Surprisingly similar behavior can be seen in the unitary-scatterer model. One can see that in the low-concentration regime (*i.e.*, up to $p \approx 0.5\%$), both the mean and the standard deviation of the LDOS are almost zero. Then at around $p = 1.0\%$ both the mean and standard deviation display a strong dependence on p , with both increasing as the impurity concentrations are increased. This behavior stops at around $p = 16.0\%$, at which point the mean LDOS reaches the largest value (out of the values of p we consider), and the mean starts to decrease once p is increased. The standard deviation continues to increase past $p = 16.0\%$ up until $p = 32.0\%$, signaling that despite the decrease in

the mean LDOS, the spatial variations remain considerable. It is interesting to note that both the mean and standard deviation of the DOS at $E = 0$ in this case depend on p very similarly to the way the same two quantities depend on σ in the Gaussian random-potential case discussed before, despite the considerable differences present between the two disorder scenarios.

Despite the huge difference in the effects seen in the quasiparticle DOS and local DOS maps between smooth and pointlike disorder, the DOS at $E = 0$ for the smooth-disorder case does display a similar dependence on the disorder strength as for pointlike disorder. For the positive-net-potential case, low doping concentrations show a mean LDOS close to zero, with a corresponding small standard deviation indicating small spatial variations in the LDOS. Both the mean and standard deviation exhibit a dependence on p up to the (quite unphysical) doping $p = 40\%$. At that point the mean LDOS becomes a maximum, but the standard deviation continues to increase past that point. The zero-net-potential case meanwhile shows much more spatial variation than the positive-net-potential case. Low dopings show small mean and standard deviations, and as p is increased these two quantities depend strongly on p . Interestingly, at $p = 10\%$ the standard deviation starts to depend more strongly on p ; consequently, at intermediate and high dopings the LDOS at $E = 0$ has a considerable amount of spatial variation. The mean LDOS also has a strong dependence on p .

The extent to which the LDOS at the Fermi energy varies over space can be visualized neatly by taking histograms of these LDOS values for various values of the disorder strength parameter. These histograms are shown in Fig. 4.16. To facilitate comparisons between LDOS distributions corresponding to different disorder strengths, we use the same bin width for each histogram. For random-potential disorder, it can be seen that the weak-disorder cases feature very narrowly dis-

tributed LDOS values. When disorder is increased until $\sigma = 1.00$ the distributions start to broaden, and as a consequence the peaks of the distributions shift rightwards, becoming lower, with the mean moving away from zero. For values of $\sigma > 1.00$ the distribution starts to narrow, with much of the distribution being concentrated near zero, but there remains a large amount of spatial variation. Because of the large weight at and near zero, $\rho(E = 0)$ is suppressed in these cases, but the distribution is much more variable than in the weak-disorder case. We note in passing that throughout the range of disorder strengths we consider, $\rho(E = 0)$ is consistently distributed log-normally, which is remarkable given how dramatically different the overall statistics of these distributions are as disorder is varied.

Moving on to unitary-scatterer disorder, at small p the distribution is centered mainly around $\rho(E = 0) \approx 0$, but with a small number of LDOS values with larger values arising from the random interference effects discussed earlier. These effects become more and more numerous as p is increased, leading to broader and flatter distributions at intermediate impurity concentrations. The behavior of the LDOS distributions in the multiple unitary-scatterer case parallels very closely that of the Gaussian random-potential disorder, with distributions for both cases widening and then subsequently narrowing once more as p or σ is increased. The main difference here is that the distribution of the LDOS for unitary scatterers is *bimodal* for moderate and large values of p : as the LDOS is suppressed almost completely *at* impurity sites, these represent a considerable number of LDOS values that are zero, and these peaks in the distributions are present independently of the variations arising from the very presence of these impurities. When one takes these impurity-site LDOS values out of consideration, the LDOS distribution is log-normal, similar to the case of random Gaussian disorder (which, unlike the unitary-scatterer model, does not

exhibit a special subset of lattice sites at which the LDOS is maximally suppressed).

In the case of smooth disorder with positive net potential, one can see that the generation of a finite DOS is achieved by an increase in the spatial variation, resulting in the broadening of the distribution. Similarly, when we consider smooth disorder with zero net potential, as p is increased, the LDOS distributions at $E = 0$ become very broad. While this effect is also seen in the other pointlike forms of disorder we looked at earlier, here the broadening is more pronounced, and much more so compared to the positive-net-potential case. We also do not hit the strong-disorder regime where these LDOS distributions start to narrow while still exhibiting strong spatial variations, which we encountered in the random-potential and unitary-scatterer disorder models.

We end this section by noting that our results for weak disorder match closely with what field-theoretic treatments of disorder find, which is that the DOS at $E = 0$ vanishes [118, 155, 10]. A crucial assumption made in the construction of these field theories is that the distribution of the disorder is narrow [10]. Indeed, we find that weak disorder of whatever form leads to a very small DOS at the Fermi energy. What our numerical results suggest however is that the DOS is not vanishing only up to some threshold value of disorder which invalidates the construction of these field-theoretic models. Instead what we find is that the DOS at $E = 0$ varies smoothly as the amount of disorder is increased, suggesting that crossovers, rather than sharp transitions, occur as one moves from weak to intermediate disorder and from intermediate to strong disorder.




Figure 4.17: Plots of C/T as a function of temperature T for different types of disorder. Gray dotted lines indicate fits of the C/T curves to the form $C/T = \gamma_0 + \alpha T$, the scaling expected from d -wave quasiparticles with a nonzero DOS at $E = 0$. The numerically-obtained C/T exhibits visible deviations from this scaling.

4.7 LOW-TEMPERATURE SPECIFIC HEAT

The next quantity we will consider is the low-temperature specific heat. We will examine the contributions of the d -wave quasiparticles to the specific heat, neglecting the effect of phonons which arise at higher temperatures. As mentioned earlier, a clean d -wave superconductor has a DOS which vanishes at $E = 0$ linearly, and this gives rise to a T^2 -dependent term in the specific heat C . Interestingly, in specific heat experiments, it is found that this T^2 -dependent term is difficult to disentangle from the signal [144]. Instead the most prominent contributions to the specific heat are the phonon contribution (scaling as T^3) and the contribution due to a finite density of states at zero energy, which scales as T , similar to a normal metal. We thus begin our discussion of specific heat with the necessary warning that it is difficult

to match the dependence on temperature of C from our numerical calculations with that found in specific heat experiments. What can be unambiguously compared between simulation and experiment, however, is the magnitude of γ_0 , the coefficient of the linear-in- T term in C which is proportional to the DOS at $E = 0$.

Shown in Fig. 4.17 are plots of C/T versus T for various types of disorder. We first discuss random-potential disorder. It can be seen that when σ is small, the specific heat scales as $C \propto \gamma_0 T + \alpha T^2$, with γ_0 very small, reflecting the fact that the DOS at the Fermi energy at weak random-potential disorder is suppressed. The behavior of γ_0 closely follows that of the DOS at $E = 0$, as a large jump in γ_0 is found at $\sigma \approx 0.35$. Even at moderately strong disorder, the specific heat is still found to scale as $C \propto \gamma_0 T + \alpha T^2$, at least up to $T \approx 0.03$ (approximately 50 K). When disorder is strong enough, the scaling finally starts to deviate considerably from that found in the weak-disorder cases. For instance, when $\sigma \approx 1.00$, C/T becomes concave downward. The large value of C/T as $T \rightarrow 0$ seen in that case is a reflection of the very large DOS at $E = 0$.

For the case of multiple unitary scatterers, the specific heat results are by and large similar to the random-potential case. Low concentrations of unitary scatterers show a very small value of γ_0 , and with large values of γ_0 reached only until $p \approx 2.0\%$ is reached. It bears noting that at low temperatures the specific heat roughly scales as $C \propto \gamma_0 T + \alpha T^2$ at low and moderate concentrations of unitary scatterers. The unitary-scatterer cases however feature mild kinks in the C/T -versus- T plots at low temperatures which are not present in the random-potential cases. These kinks arise from the particular form of the DOS profiles in the unitary-scatterer cases, which show both a rounding of the DOS at energy scales set by the scattering rate, and ultimately its suppression at $E = 0$. The kink in the C/T profile becomes

more prominent with increasing p , and in the strong-disorder regime the plot becomes, as in the random-potential case, concave downward.

Finally, smooth-potential disorder gives rise to specific heat behavior that is rather demonstrably different from that arising from random-potential or unitary-scatterer disorder. Low concentrations of smooth scatterers (*e.g.*, $p \approx 2.5\%$ or $p \approx 5.0\%$) show $C \propto \gamma_0 T + \alpha T^2$ scaling of the specific heat, with correspondingly small values of γ_0 , reflecting the relatively small DOS at the Fermi energy due to these levels of smooth disorder. However, the unusual behavior of the DOS at $E = 0$ at higher concentrations p manifests itself in a strange kink in the plot of C/T versus T , showing strong deviations from the scaling one would expect from both d -wave dispersion and a finite DOS at the Fermi energy. The large value of C/T as $T \rightarrow 0$ results naturally from the enhancement of the DOS at $E = 0$, and as T is increased C/T dips, then rises linearly once more past a certain temperature. It is worth noting that the deviations from the expected scaling are fairly localized within a small region near $T = 0$, with the specific heat returning to quadratic scaling $C \propto \alpha T^2$ once temperature is raised past some threshold value.

Given the aforementioned difficulty of measuring precisely α from experiment, we cannot say much about how consistent with experiment our numerically-obtained scaling for C is. However, what we obtain for γ_0 can be compared with that found from experiment with definiteness. We will return to a comparison with results from specific heat experiments at the conclusion of this chapter.

4.8 QUASIPARTICLE LOCALIZATION

The final quantity of interest to us is the localization length λ . Unlike the DOS and the specific heat, the localization length is not an experi-




Figure 4.18: Plots of the localization length λ (in units where the lattice constant $a = 1$) for different types of disorder, taken at three different energies E . The x -axis shows the disorder strength parameter, given by σ for random-potential disorder (leftmost plot) and the impurity concentration p for unitary-scatterer and smooth disorder (middle and rightmost plots).




Figure 4.19: Plot of the localization length λ as a function of energy E for different types of disorder.

mental observable; no experiment exists which measures the quantity described by Eq. 4.15. However it is a very important quantity in that it gives information as to how localized the states at a particular energy are. It is a rather difficult quantity to measure in finite-size simulations of lattice systems because more often than not λ is much bigger than the system size. The numerical method we use however circumvents this difficulty by allowing one dimension of the system to be much longer than the other. Thus we can use one definition of the localization length which involves the transmission probability between two ends of an elongated two-dimensional system [106, 23, 92, 186]. This enables us to *directly* and exactly calculate the localization length for the full disordered system. As a first exercise we calculate the localization length λ using Eq. 4.15 on the same set of disorder configurations as used in Figs. 4.15 and 4.16. In Fig. 4.18 we show λ for three differ-

ent values of E : $E = 0$ (corresponding to states at and near the Fermi energy), $E = 0.15$ (for states far from either the Fermi energy or the coherence peaks, but still within the d -wave gap), and $E = 0.3$ (states at and near the coherence peaks). We also plot in Fig. 4.19 the localization length as a function of energy for various forms of disorder using the same configurations used in Figs. 4.1, 4.3, 4.5, and 4.8.

Let us discuss localization in the random-potential model first. We begin with the states at and near the Fermi energy. The dependence of λ at $E = 0$ on σ appears to be unusual: it is approximately constant from $\sigma = 0.13$ to $\sigma = 0.25$, then hits a peak at around $\sigma = 0.35$ before decreasing with increasing disorder. The localization lengths for these states are small at weak disorder ($\lambda \approx 170$), and the strong-disorder λ is even smaller— $\lambda \approx 50$ at $\sigma = 1.00$, smaller in fact than the transverse dimension of the system.

The localization lengths at intermediate and high energies show more consistent behavior than the low-energy case. These decrease monotonically as disorder is increased. It is worth noting that while these states are quasi-extended at low disorder, with a larger localization length than for the $E = 0$ states, there is a range of σ where these higher-energy states have a smaller λ than states near the Fermi energy, which coincides at the range where $\lambda(E = 0)$ peaks. We will later show that the contrast in behavior seen here between the $E = 0$ case and that for higher energies is also seen in other forms of disorder. However a remarkable fact about random-potential disorder is that, of the various types of disorder we consider, this has the most dramatic impact on the behavior of the localization length. For one, it can be seen from the results that $\lambda(E = 0.15) > \lambda(E = 0.3)$ for all values of σ we use, implying that the intermediate-energy states are less localized than the higher-energy ones—a feature that is not seen in other forms of disorder we consider. Also, the closeness of the values of λ at

different E for all σ is much less pronounced in the unitary- or smooth-scatterer cases. These cases exhibit a more visible and rigid separation of λ as a function of energy for a wide range of disorder strengths—*i.e.*, $\lambda(E = 0) < \lambda(E = 0.15) < \lambda(E = 0.3)$ for these cases, which the random-potential case clearly does not show. There is a disorder strength— $\sigma \approx 1.00$ —at which the localization lengths for the three different energies are approximately the same number; this corresponds to the onset of the strong-disorder regime.

We can see these effects more clearly when the localization length is plotted versus energy. Notice that for all disorder strengths we consider, the states near the Fermi energy are strongly localized, and their localization lengths at $E = 0$ are close in value to one another even as the amount of disorder is varied. For weak disorder ($\sigma = 0.125$ and $\sigma = 0.25$) the localization length rises from a small value at $E = 0$ into a prominent peak at some small energy ($E \approx 0.02$ for $\sigma = 0.125$ and $E \approx 0.01$ for $\sigma = 0.25$), after which it decreases as energy is increased. It bears noting that the localization lengths at intermediate and high energies at these disorder levels are still quite large, at around 200–600 lattice constants. At $E \approx 0.3$ (the coherence-peak energy), the localization length for the $\sigma = 0.125$ case starts to increase; this effect is not visible when disorder is stronger. When disorder is increased, the localization length stops exhibiting these energy-dependent features: when $\sigma = 0.50$, λ is almost energy-independent, and this is even more the case for $\sigma = 1.00$, indicating that the states are strongly localized at all energies.

We next discuss unitary-scatterer disorder. Focusing first on the $E = 0$ case, we see that it exhibits the same unusual dependence on p as the random-potential case at the same energy does on σ . At low impurity concentrations $\lambda(E = 0)$ increases slightly with increasing p , reaching a peak at $p = 1.0\%$ before decreasing monotonically as a function of

p . This is in stark contrast with the behavior of λ at higher energies, which monotonically decrease with increasing p for *all* p we consider. It is worth comparing these plots to the ones derived for the random-potential case. Here we can see that, in the unitary-scatterer model, the low-disorder cases at intermediate and high energies have a far larger localization length than in the random-potential model. The impact of unitary scatterers is less pronounced than Gaussian random-potential disorder at low disorder, but with stronger disorder the behavior of the localization length for this case starts to become similar to that of the random-potential case. At higher impurity concentrations, the values of λ for different E approach each other as p is increased, with $\lambda(E = 0) \approx \lambda(E = 0.15) \approx \lambda(E = 0.3)$ at $p = 16.0\%$, which corresponds to the strong-disorder regime of this particular form of disorder.

The localization length for the unitary-scatterer model exhibits a very different dependence on energy from the Gaussian random-potential case, at least for small amounts of disorder. Near the Fermi energy, the states are strongly localized, and as with the previous disorder model we discussed the localization lengths at $E = 0$ are close in value to each other. At low concentrations, the localization length increases from $E = 0$ up to some energy, then after that point it increases once more with increasing energy, but at a decreased rate. This is seen in the $p = 0.25\%$ and $p = 1.0\%$ cases. Evidently, past a certain threshold energy the states become far less localized, with very large localization lengths at intermediate and high energies (around 300-700 lattice constants), and states at higher energies are less localized than those at intermediate energies—in stark contrast to the Gaussian random-potential case. When p is increased, however, these energy-dependent features become far less noticeable, as can be seen when $p = 4.0\%$, indicating that when disorder is large enough, the effects of localization become visible at all energies, and not just at small energies. At these

large- p regimes the behavior of the localization length with increasing energy becomes very similar to that seen in the strong-disorder random-potential case in that little, if any, dependence on energy can be discerned, and in that all states are localized, even at high energies.

We finally consider the localization due to smooth disorder. Here we will consider smooth disorder potentials whose spatial average is zero—*i.e.*, disorder potentials described by Eq. 4.23. Here the smoothness of the disorder potential makes itself particularly manifest. First, for the states near the Fermi energy, it can be seen that $\lambda(E = 0)$ does not exhibit a sharp peak at some disorder strength, unlike what is seen for random-potential or unitary-scatterer disorder. Instead its profile is flat at low p , and it then smoothly decreases as p increases. It is interesting to note that $\lambda(E = 0)$ manages to be fairly large even at high impurity concentrations. Notably, when one has unphysically high p (*e.g.*, $p = 40\%$ or $p = 80\%$), the localization length at the Fermi energy is still $\lambda(E = 0) \approx 80$ -100. For comparison's sake, that point is reached with random-potential disorder at $\sigma \approx 0.50$ and with unitary-scatterer disorder at $p \approx 8.0\%$ —levels of disorder which are strong enough to destroy coherence peaks. From just the consideration of states near the chemical potential, the impact of smooth disorder on λ is much less pronounced than either of these other cases.

The absence of any strong impact on the localization lengths is even visible at higher energies. Here it can be seen that the localization lengths for $E = 0.15$ and $E = 0.3$ are very large— $\lambda \approx 500$ for low p . Even at $p = 20.0\%$ we find that $\lambda \approx 300$. Such large values of λ are seen only at low levels of disorder for the random-potential model ($\sigma \approx 0.18$) and the unitary-scatterer model ($p \approx 1.0\%$). Even at very high smooth-impurity concentrations such as $p = 40\%$ and $p = 80\%$, we find that $\lambda \approx 150$ -200; these localization lengths correspond to low disorder levels in the random-potential and unitary-scatterer models

of disorder. Another notable observation is the fact that $\lambda(E = 0.15)$ and $\lambda(E = 0.3)$ are quite close to each other for all p . We *never* reach the onset of the strong-disorder regime that we observe in the other pointlike models of disorder—that is, the disorder strength at which $\lambda(E = 0) \approx \lambda(E = 0.15) \approx \lambda(E = 0.3)$. We find that at the absurdly unphysical $p = 80\%$ concentration $\lambda(E = 0.15) \approx \lambda(E = 0.3)$, but $\lambda(E = 0)$ remains parametrically much smaller than either. This is a clear sign that, even with very large off-plane impurity concentrations, the impact of this form of disorder on the localization of states at all energies is much more muted than in random-potential or unitary-scatterer disorder—*especially* at intermediate or high energies.

The plots of λ versus E for the smooth-disorder case exhibit a number of differences from the other two forms of disorder we have considered. First, the states near the Fermi energy are strongly localized, but as the energy is increased the localization length increases sharply for all p we consider until some value of E is encountered, at which point the localization length exhibits a far less pronounced dependence on E . At low concentrations (*e.g.*, $p = 5\%$), the localization length by and large increases as energy is increased, but with considerable random fluctuations. When the concentration is increased, the localization length grows more slowly with E . It is interesting to note that the localization length trends upward past $E \approx 0.3$, the energy where coherence peaks are found, indicating that states at energies higher than the coherence-peak positions are quite extended in space. These behaviors are different from those seen in unitary-scatterer or random-potential disorder, although there are similarities—at low energies λ for smooth disorder behaves similarly as in unitary-scatterer disorder, while at higher energies there is a noticeable increase in λ starting at $E \approx 0.3$, similar to what is seen in weak random-potential disorder. Even at very large values of p the behavior of the localization length is

still similar to that at low concentrations; at $p = 40\%$ λ is still visibly energy-dependent, suggesting once more that even in these regimes disorder of this form has a far weaker effect than the other two types of disorder we have considered. It is instructive to compare smooth disorder at $p = 20\%$ to, say, unitary-scatterer disorder at $p = 1.0\%$ or random-potential disorder at $\sigma = 0.25$ —the localization lengths for these three cases occupy a similar range to each other.

Our numerical results for the localization length are in good qualitative agreement with the analytical results obtained by Lee, who performed self-consistent calculations for weak Gaussian random-potential and dilute unitary-scatterer disorder in the d -wave superconducting state [100]. Some caveats need to be mentioned, however, as our numerics exhibit more detail and structure about the localization properties of these disorder models. Lee argued that the states near the Fermi energy are localized, although the extent to which these states are localized away from $E = 0$ (instead of being quasi-extended) was found to depend on whether the scattering is in the Born limit or the unitary limit. In the Born-scattering limit of Gaussian random disorder it was found that localization is negligible away from $E = 0$, whereas for unitary scatterers localization can be observed at energies $E < \Gamma_0$, where Γ_0 is the scattering rate in the superconducting state as $E \rightarrow 0$. In our numerical results we find that the states within the vicinity of $E = 0$ are special in being much more localized than their neighbors in energy space for *all* weak-disorder models we consider. We find that the dip in the localization length at $E = 0$ for the unitary-scatterer case is narrower than Lee's calculations suggest—that is, the energy range over which the quasiparticles are sharply localized is considerably narrower than Lee's estimate of the scattering rate Γ_0 . Away from $E = 0$ the behavior of the localization length is in much more quantitative agreement with Lee's predictions: $\lambda(\omega) \approx v_F/\Gamma(\omega) \approx 1/\omega$ for

random-potential disorder in the Born limit and $\lambda(\omega) \approx \omega$ for unitary scatterers, which are behaviors similar to what we can observe in the weak-disorder cases we discussed earlier. Our numerical results are also in good agreement with earlier numerical work on random-potential and unitary-scatterer models of disorder [186, 47, 194].

The behavior of the localization length as a function of E at weak disorder resembles that predicted by Senthil and Fisher from field theory [155]. Their inclusion of diffusive modes—as elucidated in the complementary diagrammatic approach by Yashenkin *et al.* [187, 67]—implies that additional behavior due to these modes, not captured by self-consistent diagrammatic theory, should account for the differences between these approaches [10]. Senthil and Fisher argue that, at least in the case of unitary scatterers in the dilute limit, there are three regimes: the ballistic regime, the diffusive regime (at $E \approx \Gamma_0$), and finally the localized regime near $E = 0$. The distinction between the ballistic and diffusive regimes cannot be clearly delineated from our numerics, but the crossover from the ballistic/diffusive regimes to the sharply localized regime can be seen very clearly in the weak-disorder cases we consider. Also, we find, in agreement with Senthil and Fisher’s results, that the localization length as $E \rightarrow 0$ in fact approaches a finite constant—in striking contrast to the predictions by Nersesyan *et al.*, who find a diverging localization length as $E \rightarrow 0$ [118]. Our calculations find that this constant localization length at the Fermi energy is independent of the disorder strength in the weak-disorder regime, and stands in contrast to the behavior of the localization length at higher energies, which is found to be dependent on the disorder strength.

4.9 DISCUSSION AND CONCLUSION

We have revisited the effects of disorder in high-temperature superconductors using exact real-space methods which allow large system sizes to be studied, and have ensured that the parameters we have used in our models hew closely to what is known about the cuprates from experiment. We have focused primarily on the density of states and the localization length, two quantities that are of central importance in the study of disordered systems, and made use of various models of site-energy disorder—random Gaussian potentials, multiple unitary scatterers, and off-plane dopants—which are found to result in vastly different behavior depending on which particular model is used.

Our main motivation for looking at the density of states of disordered d -wave superconductors once more is the observation—seen consistently in experiments as disparate as specific heat measurements, ARPES, and STS—that there appears to be a nonzero density of states in the cuprate superconductors, even those for which the samples can be made very clean, such as YBCO. The persistent appearance of such a signal has prompted a number of explanations that do not invoke disorder, and at the very least suggests the possibility that physics beyond the usual paradigm of d -wave superconductivity has to explain this. We reconsider the possibility that disorder is responsible for this nonzero density of states, and find that disorder of a form rarely considered in the older literature on the subject can in fact be a plausible explanation for this phenomenon.

The idea that the cuprates host different variants of disorder is not strange or even new, as STS experiments can directly visualize the disorder present in these materials and find that throughout the phase diagram of BSCCO, the signatures of disorder are present—whether in the form of quasiparticle interference in the superconducting state,

or the real-space inhomogeneities in the DOS and pairing gaps in the pseudogap regime. However, the very chemistry of the cuprates naturally precludes the possibility that disorder is present within the CuO_2 planes. The most natural form of disorder, at least from a chemical standpoint, appears to be dopants located some distance from the conducting planes. Doped cuprates host a nonzero number of oxygens at off-plane sites, and they exert an effect on the physics within the CuO_2 planes by means of a screened Coulomb potential that modifies the chemical potential at sites located within the conducting planes. The longer-ranged nature of these potentials makes them trickier to model than unitary scatterers or random-potential disorder, but the numerical methods presented here allow the effects of these forms of disorder to be simulated with great efficiency. We have also been able to obtain the localization length, a quantity that, thanks to its large size, is unable to be extracted from exact diagonalization studies of small systems, and closely examine its behavior as a function of disorder strength and energy for different models of disorder used.

Examining first random-potential disorder, we find that its effect on the DOS is to flatten the coherence peaks at the edge of the d -wave gap, and that the dominant spectral-weight transfer processes appear to be from the coherence peaks to intermediate energies, with not much spectral weight transferred to the region near the Fermi energy. A large finite DOS at $E = 0$ is not generated until fairly strong levels of disorder are reached. We consistently see that the DOS at the Fermi energy is suppressed relative to that at nearby energies; that the DOS profile at that region is V-shaped, in stark contrast to what is seen in STS experiments; and that coherence peaks are considerably flattened, even when disorder is weak. For this form of disorder the localization length exhibits an interesting dependence on energy and disorder strength: states near $E = 0$ are localized, but the localization length

sharply increases moving away from the Fermi energy, until it starts decreasing monotonically as energy is increased.

Multiple unitary scatterers are found to exhibit spectral-weight transfers from the coherence peaks to a particular energy scale, resulting in the presence of a hump-like feature in the DOS at small energies, with otherwise small deviations from the clean case at small impurity concentrations. The DOS consistently exhibits suppression at $E = 0$, and manages to acquire a large finite value only when fairly large concentrations are reached. As the concentration is increased the d -wave gap gets filled and the coherence peaks become more and more flattened. The behavior of the localization length for this form of disorder is drastically different from the random-potential case, especially at low levels of disorder. The localization length is small for states near the Fermi energy, then increases sharply until some energy is reached, and subsequently increases once more, but at a far slower rate.

Off-plane scatterers are the most interesting case, insofar as even a large concentration of such dopants turns out not to destroy the d -wave profile of the DOS—spectral weight transfers are minimal at best—while generating a finite DOS at $E = 0$ at levels of disorder that are not far off from what is seen in experiment. For the parameters we have used in our disorder potential, concentrations around 10-20% result a small but visibly finite DOS at the Fermi energy and a U-shaped DOS profile for small energies, which are consistent with experiment. At higher concentrations, an unusual resonance forms at $E = 0$; this appears to be an intrinsically many-impurity effect, as there is no obvious correlation between the disorder potential and the resulting resonant DOS. The localization length is found to be much bigger than that seen in the previous two disorder models used. While the states near $E = 0$ have a short localization length, away from that region the localization lengths are very large, even when the concentrations are

sizable—for comparison’s sake we have found that a concentration of 20% off-plane scatters has roughly the same effect on the localization length for a broad energy range as an ensemble of unitary scatterers with concentration 1.0%, or random-potential disorder with $\sigma = 0.25$. These results all point to the fact that smooth scatterers have far less of an effect on the DOS and the localization properties of a d -wave superconductor than the other two disorder models, even when the amount of smooth disorder is large.

It is worth asking whether we can make any definitive conclusions regarding the nature of disorder in the cuprates from our results. Disorder makes itself felt in a panoply of effects seen in various experiments, but isolating its effect with any definiteness is difficult given the vast array of strongly correlated phenomena present in the cuprates. We have focused mostly on single-particle properties in the form of the DOS, and it bears noting that many of the effects due to disorder we have described could be due to other effects as well. Disorder broadens the DOS, but so do interactions (in the form of self-energies) and finite temperatures. We work in the $T \rightarrow 0$ limit, so the latter alternative is ruled out, but even then we cannot rule out the possibility that nontrivial physics beyond the mean-field model of a d -wave superconductor we work with can explain the bulk of what is seen in experiment. The best we could do in the meantime is to look at the extent to which disorder—and disorder alone—reproduces experiment.

How does one square the presence of a finite DOS at $E = 0$ with the amount of disorder present in the cuprates, assuming that disorder alone is responsible for the broadening? Zero-field specific heat measurements on YBCO find a residual T -linear term in the specific heat whose coefficient is $\gamma \approx 2 \text{ mJ}\cdot\text{mol}^{-1}\cdot\text{K}^{-2}$ [115, 116, 144]. Using Eq. 4.13, we find that $\rho(E = 0) \approx 0.1$. Interestingly, angle-resolved photoemission spectroscopy provides a similar value for the residual

DOS at $E = 0$. The widths of energy distribution curves taken from ARPES experiments on clean BSCCO suggest that the scattering rate in the superconducting state is around $\Gamma \approx 15$ meV near zero binding energy [82, 170, 91]. Using the formula

$$\rho(E = 0) = \sum_{\mathbf{k}} \frac{\Gamma}{\epsilon_{\mathbf{k}}^2 + \Delta_{\mathbf{k}}^2 + \Gamma^2}, \quad (4.26)$$

this too leads to $\rho(E = 0) \approx 0.1$. These provide constraints in the amount of disorder in the cuprates, assuming that this finite value of the DOS is due purely to disorder.

Unitary scatterers can be safely ruled out. STS experiments show few, if any, signals of unitary scatterers in real-space conductance maps of clean BSCCO. They do not show the resonances one sees in zinc-doped BSCCO. The presence of vacancies however could be one source of unitary-scatterer disorder in the cuprates. How numerous would they have to be to produce a finite density of states consistent with experiment? From our numerics it appears that $p = 2.0\%$ and $p = 4.0\%$ are the closest matches to this, but these concentrations of unitary scatterers appear to be too high to describe clean BSCCO. In fact, these are too large to describe even zinc-doped BSCCO—the STS experiments on these doped materials use a zinc-dopant concentration of $p = 0.6\%$ [129], and conductance maps from these studies show very prominent resonances that are not present in clean BSCCO.

Weak random-potential disorder can also be ruled out as a primary source of the finite DOS ultimately for two reasons. First, by the argument we used above for unitary-scatterer disorder, the level of Gaussian disorder needed to reproduce $\rho(E = 0) \approx 0.1$ is around $\sigma = 0.50$. At this level of disorder, the coherence peaks are completely flattened and smeared out. This is in contrast to what is seen in STS experiments, which consistently find a spatially-averaged LDOS with prominent co-

herence peaks in the superconducting state of BSCCO. Second, at this level disorder is strong enough that the usual telltale signatures of QPI are no longer present. As discussed before, this form of disorder is consistent with QPI when σ is very small [161]. When disorder of this sort is weak, peaks in the power spectrum of the LDOS corresponding to what the octet model predicts are visible and prominent, and the real-space maps show crisscrossing patterns consistent with experiment. However this is destroyed when disorder is increased, and STS studies of BSCCO show that disorder is never strong enough to prevent the formation of modulations governed by QPI—disorder has to be weak enough that QPI is preserved. The strong levels of disorder that would produce a finite DOS at $E = 0$ consistent with the large self-energies found in ARPES would on the other hand not result in QPI. This suggests that QPI due to weak random-potential disorder occurs *on top* of other effects that are primarily responsible for the finite DOS at $E = 0$.

This leaves us with smooth disorder due to off-plane dopants. Many aspects of these dopants remain mysterious, and important properties—the screening length, the strength of the potential, and even the exact placement of these dopants—are not known with any degree of accuracy. Nevertheless, in our treatment of these dopants we have attempted to be consistent with a number of crucial facts. First, the dopant concentration is generally large, and second, the dopants are located some distance away from the CuO_2 planes, which leads to small-angle scattering. We find that the effects of smooth disorder on the DOS are much more muted than in the other two disorder cases, with minimal impact on the heights of the coherence peaks and only small spectral-weight transfers to the region near the Fermi energy. This is seen too in our calculations of the localization length in the presence of this form of disorder, which is found to remain quite large for a wide energy range even for large impurity concentrations p . We

find that $\rho(E = 0)$ acquires a value within the range $[0.05, 0.10]$ for a rather wide range of p —this would correspond to $p \approx 10\text{-}20\%$, depending on which disorder scenario one has. The more realistic scenario, in which the impurity strengths of the scatterers all have the same sign, features considerably more suppression of the DOS at the Fermi energy than the case where the spatially-averaged disorder potential is zero. The zero-average scenario has a number of very interesting features at large concentrations ($p \approx 20\%$, for instance), such as resonances at $E = 0$ whose origins appear to be unrelated to the exact details of the disorder potential. While these prominent resonances are not seen in experiment, lower impurity concentrations show much more muted LDOS patterns at $E = 0$, which, while yielding a nonzero DOS at the Fermi energy when averaged, are far less observable than at higher concentrations, and the value of the DOS appears to be fairly consistent with experiment.

Having said this, studies of quasiparticle scattering interference in BSCCO do consistently demonstrate that small- *and* large-momenta scattering processes occur in BSCCO, which is something that *purely* smooth disorder cannot take into account on its own. Purely smooth disorder such as what we discussed in this section has been shown to give rise to Fourier-transformed maps where large-momenta peaks are missing [125, 161]. Because so much of the chemistry of the cuprates is consistent with off-plane disorder, and because strong, pointlike potentials are rarely encountered in BSCCO, it is a bit of a mystery why the observed QPI exhibits large-momenta peaks. It is of course entirely possible that these effects occur in tandem with each other—smooth potentials cause the finite DOS, while relatively weak pointlike disorder causes QPI—but a full resolution still awaits, and possibly requires a much more microscopic modeling of the tunneling process [93].

We additionally caution the reader that our work has focused on strictly two-dimensional d -wave superconductors, and as such we have neglected the effects of coupling to the third dimension. The suppression of the DOS in the presence of in-plane pointlike disorder has been shown in field-theoretical work to occur strictly in 2D, and the logarithmic divergences responsible for this effect are cut off when interlayer coupling is included [118]. The observed dips we see in the in-plane disorder cases would be lost the more three-dimensional the system becomes, and this leaves open the possibility that, in the presence of interlayer coupling, this finite DOS could be due in part to the presence of pointlike forms of disorder. We thus stress that our results do not by any means suggest that smooth disorder is the be-all and end-all cause of the finite DOS at the Fermi energy. However, as noted earlier, YBCO is noted to have clean CuO_2 planes, so any influence of in-plane disorder on the DOS is likely to be very weak, regardless of the presence of interlayer coupling.

The possibility that the finite DOS at the Fermi energy in the superconducting state of the cuprates is due to disorder—smooth disorder, in particular—does not leave other explanations wanting, however, and one should not rule these out completely. It is possible that disorder is present alongside other, more exotic effects involving strong interactions (quantum criticality, for instance). In such a scenario there would be even more broadening involved. When the self-energies incorporating both disorder and interactions contain a nontrivial dependence on frequency or temperature, numerous interesting effects could conceivably occur. It would be interesting to see if alternative explanations invoking, say, quantum criticality or coexisting order result in the preservation of crucial aspects of the d -wave state, as the smooth-disorder scenario does.

On a completely different note, our results suggest a number of avenues for future work. First, the incorporation of full self-consistency is one possibility, albeit a very technically challenging one, at least from the point of view of our methods. While self-consistency *may* not be completely necessary—it might very well be that the superconductivity in the cuprates is decidedly non-BCS-like—it would be very interesting to see how smooth disorder affects the superconducting order parameter. The non-self-consistent results in this chapter suggest that smooth disorder has a far more muted effect on the single-particle properties of the d -wave superconductor than unitary-scatterer or random-potential disorder, so it is reasonable to guess that a fully self-consistent treatment would result in the preservation of d -wave superconductivity up to very high off-plane impurity concentrations, and consequently a large T_c even when the superconductor is disordered. A second possibility is to revisit the exact calculation of the superfluid stiffness, T_c , and optical conductivity in the superconducting state [149, 148] in the presence of off-plane disorder, and to examine if superconductivity is ever destroyed by smooth disorder. Our results suggest that even something as relatively anodyne as disorder—especially a relatively overlooked form of disorder like off-plane dopants—can produce surprisingly rich physics that accounts for many observed experimental properties of the cuprate high-temperature superconductors.

SELF-ENERGIES AND QUASIPARTICLE SCATTERING INTERFERENCE

5.1 INTRODUCTION

The copper-oxide superconductors are well-known to be strongly correlated materials. Many phenomena exhibited by the cuprates evade explanations based on weakly interacting quasiparticles. Perhaps the most notorious example of this is the “strange metal,” the normal state of these materials near optimal doping. This shows behavior that, as probed by transport, is very different from that seen in conventional metals, which are described well by Fermi-liquid theory [84]. Another similarly perplexing phase of these cuprates is the pseudogap, in which the density of states is prominently suppressed near the Fermi energy, exhibiting numerous exotic phenomena such as various ordered phases, gap inhomogeneities, and “Fermi arcs”—disconnected segments in momentum space hosting gapless excitations—as seen in experiments such as scanning tunneling spectroscopy (STS) and angle-resolved photoemission spectroscopy (ARPES) [164, 121]. Even the superconducting state, which is comparatively well-understood among the various phases of these materials, is highly unusual: it has d -wave pairing, leading to gapless, Dirac-like nodal quasiparticles

[157, 36]. It appears to be much more stable against disorder than d -wave mean-field BCS theory predicts, while unusual interaction effects, as probed by ARPES, are seen to emerge as the temperature approaches T_c [141, 140, 138, 139], which in turn is much higher than in conventional superconductors. The $T = 0$ states at low and high doping are rather firmly established as an antiferromagnetic Mott insulator and a conventional Fermi liquid, respectively, but the intermediate-doping states remain to be fully understood. A full microscopic theory of these materials consistent with all of these phenomena has yet to be developed.

Much understanding can nevertheless be gained by adopting a phenomenological approach towards modeling these various phases of the cuprates. Starting from a weakly-interacting picture, interaction or disorder effects can be included by putting in the appropriate self-energy, which “dresses” up the mean-field description one starts out with. For instance, many of the unusual transport properties of the strange metal, such as linear-in- T resistivity, can be captured by the marginal Fermi liquid self-energy first introduced by Varma *et al.* [173]. While this self-energy enters the Fermi-liquid propagator in what appears to be an innocuous manner, it results in the complete absence of quasiparticles at $T = 0$: the quasiparticle weight of a marginal Fermi liquid vanishes at zero temperature. This MFL self-energy has been shown to account for much of the transport anomalies seen in the cuprates, although its microscopic origins remain largely unknown. Similarly, much insight can be derived by treating the d -wave superconducting state as a mean-field, albeit unconventional, BCS superconductor. This starting point is largely justified by experiment. In general, ARPES finds that the Bogoliubov quasiparticles inside the superconducting state are well-defined excitations [82, 46, 163, 102, 176], while STS similarly finds behavior consistent with *coherent* quasiparticles scattering

off of disorder, leading to quasiparticle scattering interference (QPI) [70, 112, 61, 90, 99, 50]. From a purely phenomenological standpoint, the d -wave superconducting state can be reasonably studied starting with this mean-field description, with self-energies included to model phenomena that deviate strikingly from the mean-field expectation.

Recently, a number of ARPES experiments on both normal and superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ (Bi-2212) across a wide doping range have provided a more complete picture of the various phenomena in these materials, with the self-energy playing a crucial role in both phases. In the superconducting state, it is observed that the superconducting gap is not the sole factor determining T_c —contrary to expectations from BCS theory. Instead, the quasiparticle scattering rate exhibits a pronounced uptick near T_c . It appears that T_c is set by the scale at which the gap and the scattering rate cross over into each other, and the temperature at which the gap closes is *larger* than T_c [141, 140, 138, 139]. Meanwhile, in the normal state, experiments affirm the validity of the marginal Fermi liquid description at optimal doping, but in addition find that the ARPES data are well-described by a self-energy that interpolates smoothly between a Fermi-liquid one at extreme overdoping and a marginal-Fermi-liquid one at optimal doping. Such a doping-dependent self-energy has been central to the proposed “power-law liquid” phenomenology first proposed by Reber *et al.* [142].

Our goal in this chapter is to provide a detailed theoretical exploration of the effects of these self-energies, in both the normal and superconducting phases, on the real-space local density of states as probed by experiments. We will focus on QPI, which has not been looked at in related high-temperature STS studies on Bi-2212. Very few experimental studies on the temperature-dependent behavior in the superconducting state have been performed thus far [53, 132, 131],

and the effects of self-energies on STS spectra have been largely unexplored save for a small number of studies [132, 9]. Given this situation, we will provide a template demonstrating how STS results might look like, providing a guide for future experiments. QPI can be used as a real-space method of probing the momentum-space structure of the excitations: one takes the power spectrum of the differential conductance maps, and the most prominent wavevectors appearing can be used to map out the underlying band structure of these materials [70, 112, 182, 25, 61, 90, 99, 50]. In addition, STS experiments can, in principle, demonstrate whether the excitation spectra are coherent or not. The presence of sharp peaks in the power spectrum of the differential conductance maps taken from the d -wave superconducting state at low temperatures is a clear-cut demonstration of the existence of the Bogoliubov quasiparticles as sharp, phase-coherent excitations [189]. This fact is corroborated by evidence from ARPES suggesting that the excitations in the superconducting state at optimal doping are long-lived, unlike those in the normal state at the same doping [82, 46, 163, 102, 176]. These peaks in the power spectrum behave exactly as the heuristic “octet model” suggests. If these Bogoliubov quasiparticles are no longer long-lived, there is no reason to suspect that these peaks will continue to be present. These will be broadened and, if the scattering rate is large enough, will be rendered diffuse enough that these no longer exist as well-defined peaks. Throughout this chapter we will examine closely in several case studies the effect of the quasiparticle scattering rate on the power spectrum of the LDOS in the superconducting and normal state.

We first study the superconducting state as temperature is varied, and consider three different possible scenarios and how these can be seen in ARPES and STS. The first, which we call the “gap-closing” scenario, is well-known from BCS theory. Here the gap shrinks contin-

uously as temperature is increased until T_c is reached, at which point it vanishes. The scattering rate is constant as a function of temperature. The second scenario is “gap-filling/closing” and is argued to be seen in ARPES experiments. Here the gap shrinks with increasing T , vanishing at some temperature T_p , but, importantly, $T_c \neq T_p$. In addition, additional spectral weight fills in at low energies as T is increased. This can be accounted for by a temperature-dependent imaginary part of the self-energy which takes on a value comparable to that of the gap at temperatures near T_c . The third scenario is “gap-filling,” wherein the superconducting gap is temperature-independent, while the scattering rate is strongly temperature-dependent, as in the second scenario. We observe the gradual disappearance of the octet-model peaks as the scattering rate becomes very large. In the two scenarios where the gap closes, we observe that the octet-model peaks can be seen to disperse when the energy is fixed and temperature is varied, but that these peaks lose coherence if the scattering rate becomes very large.

As for the normal state, three scenarios are also considered. The first is the ordinary Fermi liquid, the second is the marginal Fermi liquid, and the third is a realistic marginal Fermi liquid whose spectral function exhibits considerable momentum-space anisotropy, with the nodal regions being much more coherent than the antinodal ones. We see that the power spectrum of the LDOS in the first two cases appears superficially similar to each other—the main feature for both is a set of caustics which correspond to the scattering wavevectors between points along the Fermi surface. The difference between the two sets of spectra is quite subtle: the caustics in the marginal Fermi liquid power spectrum are much more broadened than those in the ordinary Fermi liquid power spectrum, owing to the smaller self-energies present in the ordinary Fermi liquid compared to those in the marginal Fermi

liquid. Finally, the momentum-dependence of the self-energy of the anisotropic marginal Fermi liquid results in a highly anisotropic LDOS power spectrum as well—scattering between incoherent portions of the Fermi surface results in very broadened segments of the caustics in the power spectrum, while the wavevectors corresponding to scattering between coherent quasiparticles give rise to sharp caustic segments.

We note that STS as a probe is particularly vulnerable to finite-temperature smearing, which can obscure the features described in our numerics. We thus augment our single-impurity results without thermal smearing with macroscopically disordered and thermally smeared simulations to provide guides to experimentalists. It is in principle possible to deconvolute the thermally smeared STS data to obtain differential conductances that feature only *intrinsic* broadening; this has been performed in a number of STS studies [132]. However it is nevertheless worthwhile to examine the extent to which the features described in the single-impurity, thermally unsmeared case survive when multiple impurities and thermal smearing are included. We find that the thermally smeared case obscures many of the features seen in the superconducting state, with the octet model peaks disappearing even when the thermally unsmeared simulations suggest they are present. In the normal state cases we study, however, the general features of the thermally unsmeared results—the caustics—survive even with thermal smearing included.

5.2 SELF-ENERGIES AND BROADENING

When considered as phenomenological inputs and in the limit of weak disorder, self-energies do not fundamentally alter any of the fundamental physics of quasiparticle scattering interference in both the nor-

mal and the superconducting state. Their main nontrivial effect is to broaden the density of states relative to the clean, non-interacting limit. In what follows we will illustrate these effects in the normal and superconducting states.

Consider a normal metallic system described by a Hamiltonian H without any self-energies. The density of states ρ at energy E is

$$\rho(E) = \sum_n \delta(E - \epsilon_n), \quad (5.1)$$

where ϵ_n is an eigenvalue of H and n is some quantum number. (In a translationally-invariant system, this quantum number could be the momentum \mathbf{k} , for example, and the sum amounts to integrating over \mathbf{k} .) If one has a finite-sized system, the spectrum of H is discrete, and the DOS consists of spikes at energies equal to ϵ_n . Now we include the effect of self-energies. In this system, the self-energy is defined as

$$\Sigma^N(n, \omega) = G_0^{-1}(n, \omega) - G^{-1}(n, \omega), \quad (5.2)$$

where G is the Green's function for the full system (with interactions, disorder, or both) and G_0 is the noninteracting/clean Green's function, written in the basis which diagonalizes H (*i.e.*, the set of eigenstates $|n\rangle$) [107]. The self-energy is assumed to incorporate all the effects of interactions or disorder, so the Green's function for the full system has the same symmetries as that of the non-interacting/clean one. The retarded full Green's function can be written as

$$G(n, \omega) = \frac{1}{\omega - \epsilon_n - \Sigma^N(n, \omega)}. \quad (5.3)$$

We can define a spectral function $A(n, \omega) = -\frac{1}{\pi} \text{Im}G(n, \omega)$. It has the following form:

$$A(n, \omega) = -\frac{1}{\pi} \frac{\Sigma_2^N(n, \omega)}{[\omega - \Sigma_1^N(n, \omega) - \epsilon_n]^2 + [\Sigma_2^N(n, \omega)]^2}. \quad (5.4)$$

Here Σ_1^N and Σ_2^N are the real and imaginary parts, respectively, of the self-energy. In the limit $\Sigma^N \rightarrow 0$, this reduces to a delta function describing the noninteracting system:

$$\begin{aligned} \lim_{\Sigma \rightarrow 0} A(n, \omega) &= -\lim_{\Sigma^N \rightarrow 0} \frac{1}{\pi} \frac{\Sigma_2^N(n, \omega)}{[\omega - \Sigma_1^N(n, \omega) - \epsilon_n]^2 + [\Sigma_2^N(n, \omega)]^2} \\ &= \delta(\omega - E_n). \end{aligned} \quad (5.5)$$

The DOS for the full system is

$$\rho(E) = \sum_n A(n, \omega \rightarrow E). \quad (5.6)$$

This means that, in the presence of Σ^N , the density of states at an energy E does not consist merely of states which satisfy $\epsilon_n = E$. For one, $\Sigma_1^N(n, \omega)$ shifts the real parts of the poles of the Green's function from $\omega = \epsilon_n$ to $\omega - \Sigma_1^N(n, \omega) = \epsilon_n$. More importantly, the spectrum is broadened and $\rho(E)$ now incorporates nonlocal contributions from states located away from E in energy space. This will be reflected in the local density of states (LDOS) as well: a map of the LDOS taken at energy E will include contributions from states at other energies, weighted by Eq. 5.4.

None of our discussion fundamentally changes when one considers the superconducting state. The full Green's function in Nambu space, including self-energies, is

$$\tilde{G}^{-1} = \begin{pmatrix} \omega - \epsilon_n - \Sigma^N(n, \omega) & -\Sigma^A(n, \omega) \\ -\Sigma^A(n, \omega) & \omega + \epsilon_n + \Sigma^N(n, -\omega)^* \end{pmatrix}. \quad (5.7)$$

ϵ_n is the normal-state energy, and $\Sigma^N(n, \omega)$ and $\Sigma^A(n, \omega)$ are the normal and anomalous self-energies, respectively [29]. Under this definition, in the superconducting state the real part of the anomalous self-energy is equal to the pairing gap. In the cases involving d -wave superconductors that we will discuss, we will focus only on normal-state self-energies, and we will take the anomalous self-energy to be frequency-independent, so that in the translationally-invariant case the gap has the usual noninteracting d -wave form given by $\Sigma^A(\mathbf{k}) = \Delta(\mathbf{k}) = 2\Delta_0(\cos k_x - \cos k_y)$.

We start with a normal-state self-energy of the form $\Sigma^N(\omega) = \Sigma_1^N(\omega) + i\Sigma_2^N(\omega)$. We assume that the self-energy depends only on ω , and that $\Sigma_1^N(\omega) = -\Sigma_1^N(-\omega)$ and $\Sigma_2^N(\omega) = \Sigma_2^N(-\omega)$. It can be shown that the the spectral functions corresponding to the particle and hole parts of the Green's functions, $A_1(n, \omega) = -\frac{1}{\pi}\text{Im}G_{11}(n, \omega)$ and $A_2(n, \omega) = -\frac{1}{\pi}\text{Im}G_{22}(n, \omega)$, are

$$A_1(n, \omega) = -\frac{1}{\pi} \frac{\Sigma_2^N(\omega)}{[\omega - \Sigma_1^N(\omega) - E_n]^2 + [\Sigma_2^N(\omega)]^2} \quad (5.8)$$

and

$$A_2(n, \omega) = -\frac{1}{\pi} \frac{\Sigma_2^N(\omega)}{[\omega - \Sigma_1^N(\omega) + E_n]^2 + [\Sigma_2^N(\omega)]^2}, \quad (5.9)$$

where $E_n = \sqrt{\epsilon_n + \Delta_n}$ are the energies of the Bogoliubov quasiparticles. Without self-energies these spectral functions consist of delta

functions at energies E_n . As in the normal case, the spectral functions are broadened by $\Sigma_2^N(\omega)$, and the presence of $\Sigma_1^N(\omega)$ shifts the real parts of the poles by $\Sigma_1^N(\omega)$. The full spectral function $A(n, \omega)$ is

$$A(n, \omega) = u_n^2 A_1(n, \omega) + v_n^2 A_2(n, \omega), \quad (5.10)$$

where u_n^2 and v_n^2 are coherence factors, given by $u_n^2 = \frac{1}{2}(1 + \frac{\epsilon_n}{E_n})$ and $v_n^2 = \frac{1}{2}(1 - \frac{\epsilon_n}{E_n})$ [24]. Consequently the density of states at energy E takes the following form:

$$\rho(E) = \sum_n [u_n^2 A_1(n, \omega \rightarrow E) + v_n^2 A_2(n, \omega \rightarrow E)]. \quad (5.11)$$

5.3 METHODS

Here we will briefly sketch the methods we utilize in the chapter. Both real- and momentum-space methods are used to ensure that our numerical results can be compared well with STS and ARPES. We first focus on real-space methods. To obtain quantities such as the local density of states, we start with the Bogoliubov-de Gennes Hamiltonian, written in a site basis:

$$H = - \sum_{ij\sigma} t_{ij} c_{i\sigma}^\dagger c_{j\sigma} + \sum_{ij} (\Delta_{ij}^* c_{i\uparrow} c_{j\downarrow} + \text{h.c.}). \quad (5.12)$$

We will parametrize the normal-state Fermi surface with a minimal single-band model capturing most of the salient features of the normal state of optimally-doped BSCCO. We set the nearest-neighbor and next-nearest neighbor hopping amplitudes to be $t = 1$ and $t' = -0.3$, respectively. The chemical potential μ is tuned to ensure that the hole doping concentration is $p \approx 16\%$. In the superconducting state the

pairing amplitude is of a d -wave nature; this is ensured by taking $\Delta_{ij} = \Delta_0$ and $\Delta_{ij} = -\Delta_0$ whenever i and j are nearest-neighbor sites in the x - and y -directions, respectively.

All real-space information about the spectrum of Eq. 5.12 can be extracted from the Green's function G . The bare Green's function G_0 —without disorder or interactions—can be written in terms of the lattice Hamiltonian H as follows:

$$G_0^{-1}(\omega) = \omega\mathbb{1} - H. \quad (5.13)$$

As we have defined them, G_0 and H are $2N_x N_y \times 2N_x N_y$ matrices living in Nambu space, as in Eq. 5.7. We will incorporate disorder or interactions into this mean-field formalism by means of a self-energy $\Sigma(\omega)$, which is another $2N_x N_y \times 2N_x N_y$ matrix with a similar Nambu-space substructure as G_0 . Most generally, $\Sigma(\omega) = \Sigma^N(\omega) + \Sigma^A(\omega)$ in the d -wave state; however we will assume that d -wave pairing has already been incorporated into the bare Green's function, so only the normal part of the self-energy enters into consideration. The full Green's function becomes

$$G^{-1}(\omega) = G_0^{-1}(\omega) - \Sigma(\omega), \quad (5.14)$$

and before proceeding we need to input first the needed form of $\Sigma(\omega)$. Note that in principle, $\Sigma(\omega)$ could be momentum-dependent; this can be incorporated into a lattice description by putting the appropriate off-diagonal couplings into Eq. 5.14.

By judiciously choosing the indexing of the sites, G^{-1} can be written in block tridiagonal form. We then invert G^{-1} using an efficient algorithm for block tridiagonal matrices [52, 69, 38, 136, 184, 103, 143, 104, 94]. The details of this method have been worked out in detail in prior

work, so we will not repeat them here [161, 162]. The advantage of this method is that it is extremely fast compared to exact diagonalization; allows general forms of disorder to be included, unlike the T -matrix method, which is exact only for pointlike impurities; and enables self-energies to be included explicitly in the Green's function, allowing the study of the unusual effects of self-energies on measurable real-space quantities. We take $N_x = 1000$ and $N_y = 120$.

The local density of states $\rho(\mathbf{r}, E)$ can be obtained from G using the following equation:

$$\rho(\mathbf{r}, E) = -\frac{1}{\pi} \text{Im} G_{11}(\mathbf{r}, \omega \rightarrow E). \quad (5.15)$$

To study quasiparticle scattering interference, we first introduce a single weak ($V = 0.5$) pointlike scatterer in the middle of the sample. We obtain the LDOS map of the central 100×100 region from Eq. 5.15 and take the *absolute value* of its Fourier transform to obtain the QPI power spectrum $P(\mathbf{q}, \omega)$. The general case of a macroscopically disordered sample can be modeled by randomly distributing a number of these weak scatterers across the sample. To provide a guide for experimentalists, we also include results in which thermal broadening is present. It is known that the differential conductance as measured by STS at temperature T is broadened by a factor $\Gamma_t \approx 3.5k_B T$ —this is simply the width of the first derivative of the Fermi-Dirac distribution function which enters into the expression for the density of states at temperature T [132]. This can be incorporated into our model by adding this temperature-dependent thermal smearing factor—note here that Γ_t is the full width at half maximum—to the intrinsic broadening due to disorder and interactions.

Another major quantity we are interested in is the spectral function $A(\mathbf{k}, \omega)$. Assuming that we have only normal self-energies entering

the Green's function, the spectral function can be computed directly from the dispersion of the Bogoliubov quasiparticles using Eqs. 5.8, 5.9 and 5.10, with $n \rightarrow \mathbf{k}$. Here, $E_{\mathbf{k}} = \sqrt{\epsilon_{\mathbf{k}} + \Delta_{\mathbf{k}}}$, and $\epsilon_{\mathbf{k}} = -2t(\cos k_x + \cos k_y) - 4t' \cos k_x \cos k_y - \mu$ and $\Delta_{\mathbf{k}} = 2\Delta_0(\cos k_x - \cos k_y)$. To numerically calculate this, the first Brillouin zone is divided into a discrete 1000×1000 grid, which is large enough to render finite-size effects insignificant.

5.4 SELF-ENERGIES IN THE SUPERCONDUCTING STATE

In this section we will focus our attention on the various effects of self-energies in the superconducting state which can be seen in STS and ARPES. The main phenomenon of interest is “gap filling,” which is seen across a wide range of dopings via ARPES and STS [141, 140, 138, 139, 132]. We will examine the phenomenological consequences of a nontrivial temperature-dependence of the scattering rate Γ on the observed spectral function, $A(\mathbf{k}, \omega)$, and the power spectrum of the LDOS, $P(\mathbf{q}, \omega)$, both for the single-impurity case without thermal smearing (the idealized case) and the case with an dilute array of weak impurities with thermal smearing (to simulate actual tunneling data).

Recall that STS experiments on the superconducting cuprates show weak and energy-dependent modulations in the LDOS due to QPI. QPI results from scattering off of weak impurities, which generate Friedel oscillations around impurities. Because of the unusual, banana-like shape of the contours of constant energy (CCEs) of d -wave superconductors, the most dominant scattering processes are those from states on one tip of a “banana” to those on another, and these dominant wavevectors appear as peaks in the power spectrum of the LDOS—this in a nutshell is the so-called “octet model.” Indeed, the peaks seen in




Figure 5.1: Plot of the gap and the quasiparticle scattering rate as a function of temperature in the gap-closing (left), gap-filling and -closing (middle), and gap-filling (right) scenarios. The behavior seen in the middle plot—corresponding to the gap-filling/closing scenario—is seen in ARPES measurements by Reber *et al.* on optimally-doped BSCCO. The markers label the values of the gap and scattering rate at selected temperatures which are used in plots throughout this section.

experimentally-obtained power spectra behave entirely in accordance with the predictions of this simple model of d -wave Bogoliubov quasiparticles scattering off of weak impurities. That said, the vast majority of STS experiments on the superconducting state of the cuprates have been performed at temperatures well below T_c , where the quasiparticle scattering rate Γ is fairly small and is only weakly dependent on temperature. However, various experiments have shown that Γ is not temperature-independent, as one would expect from elastic scattering off of impurities—it instead exhibits a very pronounced dependence on T as temperatures are increased. In fact, recent ARPES results suggest that $\Gamma(T)$ is roughly of the same size as the superconducting gap $\Delta_0(T)$ itself as $T \rightarrow T_c$ [138]. Furthermore, the same ARPES results show that $\Delta_0(T)$ does not go to zero at T_c , as one would expect from BCS theory. Instead, d -wave pairing correlations are seen to exist beyond T_c , and persist up to a higher temperature scale which appears to decrease as doping is increased. We show in Fig. 5.1 plots of the superconducting gap and the quasiparticle scattering rate as a function of temperature for three different scenarios: the BCS scenario, in

which the gap closes as T is increased, becoming zero at T_c ; the gap-filling and -closing scenario, seen in ARPES experiments by Reber *et al.* on optimally-doped BSCCO ($T_c \approx 90$ K), in which the gap shrinks and the quasiparticle scattering rate increases as T is increased, but the gap remains finite past T_c [138]; and the gap-filling scenario, in which the gap remains roughly temperature-independent while the scattering rate increases at T near T_c . We will carry out the exercise of obtaining results measurable by STS experiments in the cuprates as temperature is increased, assuming consistency with ARPES results. It is an interesting experiential question to see if the peaks suggested by the “octet model” still appear when the quasiparticle scattering rate is very large, as appears to be the case when $T \approx T_c$.

The temperature-dependence of the superconducting gap and the scattering rate can be parametrized simply as follows. As argued by Reber *et al.*, the experimentally-measured gap amplitude at optimal doping can be fit to the following BCS-like functional form,

$$\Delta_0(T) = \Delta_0(0) \times \tanh \left(\alpha \sqrt{\frac{T_p}{T} - 1} \right), \quad (5.16)$$

where T_p is the temperature at which the gap fully closes, $\Delta_0(0)$ is the value of the gap at $T = 0$, and α is a dimensionless number of order unity [139]. In our numerics we will take $\Delta_0(0) = 0.096$, $T_p = 100$ K, and $\alpha = 2.32$. We remind the reader that $\Delta_0(T)$ enters the momentum-space gap function as $\Delta(\mathbf{k}, T) = 2\Delta_0(T) \times (\cos k_x - \cos k_y)$. As for the scattering rate, we use the form obtained by Chubukov *et al.* [28], which we write in the following manner:

$$\Gamma(T) = \Gamma_0 + \Omega \sinh \left(\frac{T_b}{T} \right). \quad (5.17)$$

Here T_b is some very large temperature scale included phenomenologically in order to provide a good fit with the experimental results and Γ_0 is the elastic scattering rate. Reasonably good fits can be obtained by using $\Gamma_0 = 0.015$, $\Omega = 2350$, and $T_b = 1100$ K. We will neglect any momentum-dependence of the scattering rate. These two functional forms in tandem with each other explain very well the phenomenology of the closing and the filling of the gap as seen in experiments.

The BCS case features only the closing of the gap, and only elastic scattering is present; as such we will take $\Gamma(T) = 0.015$ in that case. To allow us to compare the results of the first case with the BCS case, we take the same functional form for the BCS case as in Eq. 5.16. This ensures that the values of the superconducting gap are the same at each temperature, and that the effects of the self-energies in the first case can be isolated very clearly and contrasted with the trivial effects seen in the BCS case. It is very important to note that in the case with both the filling and closing of the gap, T_p is *not* equal to T_c , whereas in the BCS case $T_p = T_c$. Finally, for the gap-filling case, we will freeze $\Delta_0(T)$ at its $T = 0$ value, and let the scattering rate vary with temperature as in Eq. 5.17.

To illustrate clearly the differences between BCS and gap-filling phenomenology, we first discuss the BCS case with only the closing of the gap and show in Fig. 5.2 plots of $A(\mathbf{k}, \omega \rightarrow E = 0.100)$ and $P(\mathbf{q}, \omega \rightarrow E = 0.100)$ for various temperatures. The main changes one can observe with increasing temperature at fixed frequency are due to way the CCEs—as seen directly in $A(\mathbf{k}, \omega)$ —are altered by the decreasing size of Δ_0 as T is increased. At $T = 10$ K, the superconducting gap is large, implying that at the low frequencies ($E = 0.100 \approx 15$ meV) at which these plots were taken the banana-shaped contours only cover a small part of the underlying normal-state Fermi surface. As Δ_0 shrinks with increasing temperature, more and more of the un-

5.4 SELF-ENERGIES IN THE SUPERCONDUCTING STATE




Figure 5.2: Gap-closing phenomenology at various temperatures. T_c here is 100 K. Left to right: The spectral function $A(\mathbf{k}, \omega)$; the Fourier transform of the LDOS $P(\mathbf{q}, \omega)$; linecuts of $P(\mathbf{q}, \omega)$ in the nodal and antinodal directions; $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing; and linecuts of $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing. Arrows indicate the locations of the peaks predicted by the octet model. All plots are taken at $E = 0.100$.

derlying Fermi surface becomes covered by the “bananas.” However, because Γ is constant as a function of temperature, the spectral functions taken at various temperatures remain similarly sharp—the CCEs maintain their shape without much visible smearing. These imply that for the power spectrum of the LDOS, the peaks corresponding to the “octet model” remain very much visible. Because no change in intrinsic broadening occurs as temperature is increased, the octet-model peaks retain their sharpness throughout the temperature ranges we consider, and even the caustics corresponding to scattering between the off-tip segments of the “bananas” are still visible and do not get broadened. The only change that occurs as temperature is changed is in the positions of the characteristic peaks of the power spectrum. Because Δ_0 decreases in size as T increases at fixed frequency, the CCEs all increase in size with increasing T , and consequently the seven octet-model peaks disperse as T is changed at fixed frequency. For instance, \mathbf{q}_7 —the smallest diagonal scattering wavevector, corresponding to tip-to-tip scattering within one “banana”—is seen to increase in magnitude as T is increased. When the gap finally fully closes, the QPI power spectrum consists of sharp, well-defined caustics characteristic of a normal metal. With realistic disorder (*i.e.*, a 0.5% concentration of weak pointlike scatterers) and finite-temperature smearing, the expected (unconvoluted) LDOS power spectra is seen to feature the loss of the octet-model peaks as temperature is increased. In particular, only at 10 K does the disordered and thermally-smearred power spectrum show these peaks. However, a sharp transition in the features of the power spectrum once the gap fully closes is still visible even at the high temperatures at which these occur—there is a change from a highly anisotropic power spectrum in the superconducting state, with pronounced spectral weight near the corners and suppressed antinodal scattering wavevectors, to the caustics seen in the zero-gap case.

If deconvolution is carefully applied to the real-space differential conductance data, only the intrinsic (that is, non-thermal) broadening will affect the LDOS and the octet-model peaks should be recovered.

Dramatically different behavior is seen once a strongly temperature-dependent but momentum- and frequency-independent quasiparticle scattering rate is included, as is the case in the gap-filling/closing scenario, the results for which we plot in Fig. 5.3. We used the same superconducting gap for each temperature as in the BCS case, so all differences between the two sets of plots can be attributed to the presence of a T -dependent Γ . At low temperatures both $A(\mathbf{k}, \omega \rightarrow E = 0.100)$ and $P(\mathbf{q}, \omega \rightarrow E = 0.100)$ are identical, as in that particular regime there is little difference between the two scenarios. However, when $T \approx T_c$, Γ is no longer parametrically smaller than Δ_0 but is instead almost of similar size, and thus the effects of the intrinsic broadening are no longer trivial. Consider first the behavior of $A(\mathbf{k}, \omega)$. At 85 K, the CCEs are still well-defined, albeit broadened considerably compared to the BCS case, with more spectral weight found in the streaks emanating from the ends of the contours which follow the underlying Fermi surface. At $T = T_c = 90$ K, even more broadening is present, and yet more spectral weight moves towards the streaks. At 95 K, $\Gamma \approx \Delta_0$, and consequently the spectral function resembles neither that of a d -wave superconductor nor that of a normal metal. Instead, it shows a quasiparticle excitation spectrum which resembles Fermi arcs. That is, there is considerable spectral weight present near the nodes, and one sees less spectral weight as one moves along the underlying Fermi surface towards the antinodes. Once the gap fully closes, what is seen is the expected isotropic normal-state spectrum in which the spectral weight along the CCE is uniform. In our plots where the gap is fully zero, we have assumed the value of the scattering rate to be equal to that given by the marginal Fermi liquid self-energy at $T = 100$ K.




Figure 5.3: Gap-filling and -closing phenomenology at various temperatures. T_c here is 90 K. Left to right: The spectral function $A(\mathbf{k}, \omega)$; the Fourier transform of the LDOS $P(\mathbf{q}, \omega)$; linecuts of $P(\mathbf{q}, \omega)$ in the nodal and antinodal directions; $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing; and linecuts of $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing. Arrows indicate the locations of the peaks predicted by the octet model. All plots are taken at $E = 0.100$.

The strong temperature-dependence of Γ has an even more pronounced effect on the single-impurity $P(\mathbf{q}, \omega)$ without thermal smearing. At 85 K, the patterns are the same as in the BCS case, with the difference that the octet-model peaks that were visible and sharp in the BCS case are now muted—the intensities of the peaks are quite weak in the gap-filling scenario. At 90 K even more smearing of the QPI patterns becomes apparent, and some peaks, such as \mathbf{q}_7 , almost completely disappear. The points corresponding to certain other octet-model peaks such as \mathbf{q}_2 and \mathbf{q}_6 , while still discernible, are so broadened as to be almost undefined at this point, and only streaks corresponding to diagonal internodal scattering remain as the prominent signal. At 95 K and beyond, all octet-model peaks cease to be well-defined signals. Instead what remains are caustics which track scattering along the underlying normal-state Fermi surface, but with variable weights depending on the location of the initial and final states on the Fermi surface, resulting in a nonuniform distribution of spectral weight along the caustics. This is considerably different from the QPI power spectrum of a normal metal with a momentum-independent scattering rate, wherein the magnitude of P along the caustics is uniform. Finally, at the point where the gap has fully closed, we see a return to a metallic QPI power spectrum, with uniform weight along all the caustics, but which is considerably smeared compared to that seen in the BCS scenario. The addition of thermal broadening and distributed disorder however results in power spectra which are very similar to those of the BCS case. This makes it difficult to distinguish the gap-filling/closing scenario from the BCS one from unconvoluted data, and one needs to perform a deconvolution of the data to recover the power spectrum with only intrinsic broadening present.

The change in the behavior of $P(\mathbf{q}, \omega)$ from a small-gap d -wave superconductor with very large Γ to a normal-state metal with zero gap

is quite stark—in contrast to what is seen in $A(\mathbf{k}, \omega)$, where the change appears to occur smoothly. Compared to the QPI power spectrum for the normal metal, the spectrum at energies below the gap for a broadened d -wave superconductor is much more suppressed in the antinodal directions (*i.e.*, the $(0,0) \rightarrow (0, \pm\pi)$ and $(0,0) \rightarrow (\pm\pi, 0)$ directions in \mathbf{q} -space). It also features much more spectral weight near the corners. All of this can be attributed to the fact that, because the gap is finite in this regime, the scattering matrix element is affected by coherence factors. In the presence of a weak perturbation in the chemical potential (as in the case of our numerics), scattering between two states where the d -wave gap has the same sign is suppressed compared to that between states where the gap has opposite signs [182, 133]. With a finite gap, this would explain why the intensities of wavevectors corresponding to scattering in the antinodal directions (which would be between states where the gap has the same sign) are weaker compared to those of internodal scattering wavevectors, resulting in the comparatively strong signal near $(\pm\pi, \pm\pi)$. This coherence-factor effect completely disappears upon the closing of the gap. It should be emphasized that this dramatic change in the spectrum as $\Delta_0 \rightarrow 0$ is still visible even in the unconvoluted thermally-smearred spectrum.

The third scenario we consider is one in which the superconducting gap remains finite and temperature-independent, while the quasiparticle scattering rate increases monotonically as temperature is raised. We plot results for this case in Fig. 5.4. We used the same scattering rate as in the gap-filling/closing case; note that at higher temperatures the scattering rate becomes of the same size as the gap. Unlike in the second case, because the gap remains a constant as T is increased, there is no change in the position and size of the CCEs as seen in $A(\mathbf{k}, \omega \rightarrow E = 0.100)$. What differs is the sharpness of these contours in momentum space. At $T = 85$ K, the contours remain sharp,

5.4 SELF-ENERGIES IN THE SUPERCONDUCTING STATE




Figure 5.4: Gap-filling phenomenology at various temperatures. Left to right: The spectral function $A(\mathbf{k}, \omega)$; the Fourier transform of the LDOS $P(\mathbf{q}, \omega)$; linecuts of $P(\mathbf{q}, \omega)$ in the nodal and antinodal directions; $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing; and linecuts of $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing. Arrows indicate the locations of the peaks predicted by the octet model. All plots are taken at $E = 0.100$.

with only a small amount of spectral weight found beyond the ends of the “banana.” When temperature is increased, the scattering rate increases, and the contours become less sharp, with more and more spectral weight found in the tails. At the highest temperatures, what had once been well-defined banana-shaped contours resemble more and more the underlying Fermi surface, but with anisotropy in the spectral weight along the Fermi surface. While most of the spectral weight remains near the nodes, considerably more weight has shifted towards the tails, which track the Fermi surface and which now extend all the way to the antinodes. However, unlike the scenario in which the gap both closes and fills, here the shape of the contours is largely preserved even with increasing broadening.

Because no change in the gap occurs with increasing temperature, the peaks seen in the single-impurity $P(\mathbf{q}, \omega)$ without thermal smearing do not disperse when frequency is fixed and temperature is varied. The main change that occurs is in the sharpness of the peaks, which is affected by how large the quasiparticle scattering rate is. At $T = 85$ K, the peaks can still be seen, but with more bluriness than at lower temperatures due to the large Γ at this temperature scale. Increasing T from this point onwards results in these peaks becoming progressively more broadened and less visible, turning into blurry patches with nonzero spectral weight. At the highest value of the scattering rate we considered, no isolated peaks are visible. With distributed disorder and thermal smearing, the plots show similar behavior as the thermally unsmearred single-impurity results, insofar as no shifts in the spectral weight as T increases appear in the spectra due to the constancy of the gap, but no peaks can be discerned at these high temperatures, and in experiment one has to deconvolute the dI/dV data to disentangle the intrinsic broadening from finite-temperature effects.

5.4 SELF-ENERGIES IN THE SUPERCONDUCTING STATE




Figure 5.5: Frequency-dependence at $T = 10$ K—a temperature at which all three scenarios are essentially identical—of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Arrows indicate the locations of the peaks predicted by the octet model. Note that the scales used for plotting the LDOS power spectra change with frequency.




Figure 5.6: Frequency-dependence at $T = 90$ K in the gap-closing scenario of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Arrows indicate the locations of the peaks predicted by the octet model. Note that the scales used for plotting the LDOS power spectra change with frequency. In this scenario, this temperature is *less* than T_C .

Further differences between the BCS and the two gap-filling scenarios can be seen by plotting both $A(\mathbf{k}, \omega)$ and $P(\mathbf{q}, \omega)$ for various frequencies. At the lowest temperatures, all three scenarios result in the same behavior, as seen in Fig. 5.5: the small scattering rate results in both sharp features in the spectral function, and well-defined peaks in the LDOS power spectrum whose position in \mathbf{q} -space changes as ω is varied, in agreement with the octet model. Because at low temperatures thermal smearing only has a weak effect, the disordered and thermally-smearred power spectra show octet-model peaks that are clearly discernable.

5.4 SELF-ENERGIES IN THE SUPERCONDUCTING STATE




Figure 5.7: Frequency-dependence at $T = 90$ K in the gap-closing/filling scenario of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Arrows indicate the locations of the peaks predicted by the octet model. Note that the scales used for plotting the LDOS power spectra change with frequency. In this scenario, this temperature is the same as T_c .




Figure 5.8: Frequency-dependence at $T = 90$ K in the gap-filling scenario of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Arrows indicate the locations of the peaks predicted by the octet model. Note that the scales used for plotting the LDOS power spectra change with frequency.

At higher temperatures, how these CCEs and QPI peaks appear in measurements depends on the degree of broadening present. In the BCS scenario, the CCEs in the spectral function remain sharp as frequency and temperature are changed. In comparison, in the two gap-filling scenarios the CCEs feature much more smearing, which in turn affects how prominent the QPI peaks appear in the LDOS power spectrum. Fig. 5.6 shows $A(\mathbf{k}, \omega)$ and $P(\mathbf{q}, \omega)$ taken for the gap-closing scenario at $T = 90$ K, while Fig. 5.7 shows similar quantities with gap-closing/filling assumed, and Fig. 5.8 shows the case with only the filling of the gap. At this temperature it is already apparent that in the gap-closing/filling case the QPI peaks broaden so much that it is difficult to see them clearly. What had been very visible QPI peaks in the gap-closing scenario have turned into barely-discernible patches in the gap-closing/filling scenario, while at higher energies no trace of the QPI peaks remain. Similarly, in the gap-filling scenario, one can see that because the gap is temperature-independent, the LDOS power spectrum resembles that of the low-temperature case, but with so much more smearing that the octet-model peaks become far less discernable. In all three cases, the thermal smearing at $T = 90$ K is so large that the fine features seen in the single-impurity unsmeared data are lost in the smeared data, and the plots appear qualitatively similar to each other.

We note further that in the two scenarios in which the gap closes (shown in Figs. 5.6 and 5.7), the shrinking of the gap with increasing T alters the shape of the CCEs as seen in $A(\mathbf{k}, \omega)$, and consequently the positions of the QPI peaks in $P(\mathbf{q}, \omega)$ change as well. The smallness of the gap ensures that the superconducting coherence peaks, located at $E_c \approx \pm 4\Delta_0$, are shifted closer to the Fermi level. At energies which satisfy $|E| > |E_c|$ the spectral function and QPI power spectrum in the superconducting state are largely similar to those of the normal




Figure 5.9: Plots of the spectral function $A(\mathbf{k}, \omega)$ (left), the power spectrum of the single-impurity LDOS without thermal smearing (middle), and the power spectrum of the multiple-impurity LDOS with thermal smearing (right) at $T = 95$ K in the gap-closing/filling scenario, taken at the Fermi energy ($E = 0$). The spectral function at this regime bears a marked resemblance to the “Fermi arcs” found in the pseudogap regime of the underdoped cuprates.

state at E , except for additional features which arise from the presence of shadow-like streaks in the spectral function, which in turn are an effect of the coherence factors which enter Eq. 5.10. The similarity to the normal-state LDOS power spectrum here is such that even in the BCS case, which has minimal broadening, no traces of the octet-model peaks appear at these high energies.

We end this section by revisiting our earlier observation that the combination of small but nonzero d -wave pairing correlations and a large scattering rate at $T > T_c$ can give rise to Fermi arc-like patterns in the spectral function. It is interesting to note that this can be seen right *at* the Fermi energy itself. In Fig. 5.9 we plot the spectral function and the LDOS power spectrum at the Fermi energy at $T = 95$ K for the gap-filling and -closing scenario. In the absence of broadening, the d -wave superconducting state would result in zero-energy states being localized only at the nodes—the four points on the Fermi surface at which the superconducting gap is zero. With broadening, however, there is now a finite density of zero-energy states in the neighborhood of the nodes. When the scattering rate is small, the effect is minor, and apart from a small arc centered near the nodes the zero-energy

states disappear a short distance away from the nodes. However, once $\Gamma \approx \Delta_0$ the regions about the nodes which support low-energy states become large: the “arc” along the Fermi surface which supports zero-energy states becomes longer and broader, and a comparable lack of spectral weight is found at the antinodes. The QPI power spectrum is quite pronounced even at the Fermi energy, and is completely different from that of a d -wave superconductor or a normal metal. Instead it shows streaks near the corners due to strong internodal scattering, and large low- \mathbf{q} patches showing strong intranodal scattering. We note that this particularly simple set of ingredients (nonzero d -wave pairing past T_c and a large quasiparticle scattering rate) has already been proposed as an explanation for the Fermi arcs found via ARPES in the underdoped cuprates [123, 120, 28, 175]. It is an interesting experimental challenge to see if these Fermi arc-like patterns can be seen by STS in the optimally-doped cuprates above T_c .

5.5 SELF-ENERGIES IN THE NORMAL STATE

We next turn our attention to the effects of self-energies on the spectra in the normal state. We had briefly touched upon aspects of this in the previous section when we considered the ARPES and STS spectra at temperatures in which the gap fully closes. We will more closely examine the consequences when the self-energy in the normal state depends on frequency, temperature, and momentum. Our main focus will be on the marginal Fermi liquid phenomenology in the optimally-doped cuprates, and we will obtain concrete experimental predictions for STS which are indicative of marginal Fermi liquid behavior. We will in turn contrast the results for the marginal Fermi liquid from that of the ordinary Fermi liquid, which is argued to be the normal state of




Figure 5.10: Plots of the self-energies for the Fermi liquid (red line) and marginal Fermi liquid (blue line) at $T = 100$ K. Here $\lambda = 0.5$, $\Gamma_0 = 0$, and $\omega_c = 1$.

overdoped cuprates. Finally, to faithfully represent real-world ARPES data, we add at the end momentum-space anisotropy in the self-energy in order to reproduce the observation that the spectra at the antinodes are considerably more incoherent than those found in the nodal region of the Fermi surface. As with the superconducting cases considered earlier, we will evaluate the LDOS power spectrum both for a single isolated impurity without thermal smearing and for a macroscopically disordered sample with thermal smearing to incorporate effects likely to be seen in STS experiments.

We will assume that the self-energy has the “power-law liquid” form suggested by Reber *et al.* from ARPES data on Bi-2212 across a wide range of dopings. This is simply given by

$$\Sigma''(\omega, T) = \lambda \frac{(\omega^2 + \pi^2 T^2)^\alpha}{\omega_c^{2\alpha-1}} + \Gamma_0, \quad (5.18)$$

where ω is the frequency, T the temperature, Γ_0 a temperature- and frequency-independent impurity scattering rate, ω_c a frequency cut-off, and α a doping-dependent exponent which is argued from ARPES data to be equal to 0.5 at optimal doping and near 1 at extreme over-

doping [142]. This parametrization conveniently captures both the marginal Fermi liquid ($\alpha = 0.5$) at optimal doping [173] and the ordinary Fermi liquid ($\alpha = 1.0$) at the overdoped side of the phase diagram. Plots of the self-energy for both the marginal Fermi liquid and the ordinary Fermi liquid at 100 K are shown in Fig. 5.10. In our numerics the parameters are chosen to hew closely to the phenomenological fits found by Reber *et al.* We will first neglect any momentum-space anisotropy in the self-energy; we will consider these effects later. We will set $\lambda = 0.5$, $\Gamma_0 = 0$, and $\omega_c = 1$ in our computations.

As an instructive case we first discuss the spectra of an ordinary Fermi liquid. Plots of $A(\mathbf{k}, \omega = 0)$ and $P(\mathbf{q}, \omega = 0)$ for this case are shown in Fig. 5.11. Because of the isotropic nature of the self-energy, the spectral weight at Fermi surface is uniform at all temperatures considered. The spectral function here is narrow at the Fermi energy due to the small value of the imaginary part of the self-energy. Consequently the single-impurity LDOS power spectrum has sharp and well-defined features which broaden as temperature is increased. The main feature of $P(\mathbf{q}, \omega = 0)$ are caustics which indicate scattering wavevectors from one part of part of the Fermi surface to another, as expected from a metal. With randomly distributed impurities and thermal smearing, the LDOS spectra still manages to be visible at reasonably high temperatures, even without deconvoluting.

The situation for a marginal Fermi liquid is largely similar. In Fig. 5.12 we have plotted both $A(\mathbf{k}, \omega = 0)$ and $P(\mathbf{q}, \omega = 0)$ for a marginal Fermi liquid ($\alpha = 0.5$) at the Fermi energy for various temperatures. As the self-energy scales goes as $\propto T$ at the Fermi energy, the width of the spectrum at the Fermi surface also increases as T increases. Like the spectral function, the LDOS power spectrum shows progressively more broadening as temperature is increased. When distributed disorder and thermal broadening are both present, the LDOS power spec-




Figure 5.11: Ordinary Fermi liquid phenomenology at various temperatures. Left to right: The spectral function $A(\mathbf{k}, \omega)$; the Fourier transform of the LDOS $P(\mathbf{q}, \omega)$; linecuts of $P(\mathbf{q}, \omega)$ in the nodal and antinodal directions; $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing; and linecuts of $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing. All plots are taken at $E = 0.000$.

5.5 SELF-ENERGIES IN THE NORMAL STATE




Figure 5.12: Marginal Fermi liquid phenomenology at various temperatures. Left to right: The spectral function $A(\mathbf{k}, \omega)$; the Fourier transform of the LDOS $P(\mathbf{q}, \omega)$; linecuts of $P(\mathbf{q}, \omega)$ in the nodal and antinodal directions; $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing; and linecuts of $P(\mathbf{q}, \omega)$ in the presence of multiple weak impurities and finite-temperature smearing. All plots are taken at $E = 0.000$.

tra are broadened and feature speckle, but retain most of the structure present in the single-impurity case—caustics can still be observed at 100 K, but much of the spectrum becomes overwhelmed by noise at higher temperatures, rendering it difficult to extract these patterns at high temperatures without deconvoluting the data.

It has to be noted that at fixed frequency and temperature the results for the ordinary and marginal Fermi liquid cases are not drastically different from each other, except for the amount of broadening present—the marginal Fermi liquid has much more intrinsic broadening than the ordinary Fermi liquid. Thus one key signature that one may look for in ARPES and STS experiments is that, assuming that the overdoped cuprates have a Fermi-liquid normal state, the spectral widths at fixed T and ω become larger as doping is decreased towards optimal doping. This is of course assuming that the normal state of the optimally-doped cuprates is in fact well-described by electrons dressed with a marginal Fermi liquid self-energy. While a marginal Fermi liquid features no quasiparticles at $T = 0$ —unlike an ordinary Fermi liquid—it is clear that this description of the normal state should produce results that resemble those arising from a much more broadened version of the ordinary Fermi liquid at finite temperature. The unusual frequency- and temperature-dependence of the marginal Fermi liquid can also be measured using both ARPES and STS, and one should see a change in the scaling of the broadening of the spectra with temperature and frequency as doping is changed. If one sees these caustics in the STS spectra in the normal state of the optimally-doped cuprates, then the “dressed Fermi liquid” description of the normal state is valid. However, if these are not present, then a much more different theory involving exotic hidden excitations may be found to be necessary.

The frequency-dependence of the spectral function and the LDOS power spectra are plotted in Figs. 5.13 and 5.14 for the ordinary Fermi

5.5 SELF-ENERGIES IN THE NORMAL STATE




Figure 5.13: Frequency-dependence at $T = 100$ K of the spectra of an ordinary Fermi liquid. Shown are plots of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Note that the scales used for plotting the LDOS power spectra are the same for all frequencies.




Figure 5.14: Frequency-dependence at $T = 100$ K of the spectra of a marginal Fermi liquid. Shown are plots of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Note that the scales used for plotting the LDOS power spectra are the same for all frequencies.

5.5 SELF-ENERGIES IN THE NORMAL STATE




Figure 5.15: The widths of the spectral function (left) and the single-impurity LDOS power spectrum (middle) versus the imaginary part of the self-energy for the marginal Fermi liquid and the ordinary Fermi liquid, both with momentum-independent self-energies, at a variety of temperatures and frequencies. These are evaluated from the widths of the momentum-distribution curves along the nodal directions for the spectral function and from the widths of caustics along the antinodal direction for the LDOS power spectrum. The rightmost graphic illustrates how the spectral widths, as defined in the text, are extracted from linecuts of $A(\mathbf{k}, \omega)$ and $P(\mathbf{q}, \omega)$. In this example the self-energy is of marginal-Fermi-liquid form, and $T = 100$ K and $E = 0$.

liquid and the marginal Fermi liquid, respectively, at 100 K. Note that for both these models both the spectral function and the LDOS power spectra broaden as frequency is increased at fixed temperature. The spectra do differ at high energies due to the renormalization of the band structure due to the real part of the self energy, which is different for both cases. It can be seen at negative frequencies the marginal Fermi liquid hits a van Hove singularity at a lower (negative) frequency than the ordinary Fermi liquid does owing to this renormalization. However this effect is quite troublesome to detect in practice, as disentangling this effect requires detailed knowledge of the bare band structure, and it is relatively unimportant compared to the scale set by the imaginary part of the self-energy. As such we will not direct any more focus on this phenomenon in this chapter.

We next examine the precise dependence of the broadening of the spectral function and the LDOS power spectrum on the self-energy; plots of these are shown in Fig. 5.15. Here, the self-energies used cover a wide range of frequencies and temperatures for both the marginal Fermi liquid and the ordinary Fermi liquid. The widths of the momentum-distribution curves along the nodal directions are proportional to the imaginary part of the self-energy. We can see this directly by obtaining the full width at half maximum of these MDCs; these widths scale *linearly* with Σ'' . As for the single-impurity LDOS power spectrum, the widths of the caustics broaden in a different manner from that of the spectral function. Quantifying this broadening is a bit trickier than for the spectral function, because the power spectrum features considerably more structure within the Brillouin zone due to backfolding. We define one measure of this broadening in the following manner. Along the $(0,0) \rightarrow (0,\pi)$ direction, there is a peak which corresponds to scattering between the antinodal portions of the Fermi surface. We define the width of the caustic as the distance between the midpoint between the peak and the minimum along the linecut at the central plateau near $(0,0)$ and the midpoint between the peak and the global minimum of this linecut. While our resolution in \mathbf{q} -space is very limited, it can be seen that the widths of these caustics scale roughly as the *square root* of Σ'' , regardless of whether the self-energy is of marginal Fermi liquid or ordinary Fermi liquid form. Furthermore, the extracted widths of the spectra of the marginal Fermi liquid are parametrically much larger than those of the spectra of the ordinary Fermi liquid.

We show in Fig. 5.16 the widths of the spectral function and the LDOS power spectrum for both the marginal Fermi liquid and the ordinary Fermi liquid as a function of frequency (for positive frequencies) for a variety of temperatures as extracted from our simulations, in addition to one-parameter fits of the form $W_{sc} = A\Sigma''(\omega, T)$ and $W_{ps} =$

5.5 SELF-ENERGIES IN THE NORMAL STATE




Figure 5.16: The widths of the spectral function (top row) and the single-impurity LDOS power spectrum (bottom row) versus frequency for the marginal Fermi liquid (left column) and the ordinary Fermi liquid (right column), evaluated at various temperatures. The fits used are taken from the complete data plotted in Fig. 5.15. The limited resolution available in the LDOS power spectrum results in the relatively jagged behavior of the plots compared to that seen in plots of the spectral function.




Figure 5.17: Energy-distribution curves taken at the nodal and antinodal points on the Fermi surface for an anisotropic marginal Fermi liquid. Here $T = 100$ K and $\beta = 0.2$ (see Eq. 5.20 for the functional form of the self-energy).

$B\sqrt{\Sigma''(\omega, T)}$ for the widths of the spectral function and the LDOS power spectrum, respectively. Here $\Sigma''(\omega, T)$ is of either marginal Fermi liquid or ordinary Fermi liquid form, and the parameters A and B are obtained from the data shown in Fig. 5.15. It should be noted that at the energy ranges we have considered, the widths of the caustics in the LDOS spectra grow more slowly with frequency compared to the widths of the MDCs; this reflects the rough square-root dependence of the caustic widths on the imaginary part of the self-energy.

Finally, we end this section by considering a marginal Fermi liquid with a realistic amount of momentum-space anisotropy in the self-energy. A variety of ARPES measurements on optimally-doped Bi-2212 have shown that the spectral function at the antinodal region of the Brillouin zone is much less coherent than at the near-nodal region [169, 2, 168, 127, 81]. The degree to which the spectral function is incoherent is most visible in energy-distribution curves taken at nodal and antinodal points along the Fermi surface; the nodal EDCs show a more prominent peak at the Fermi energy compared to the nodal ones. This suggests that the full self-energy is anisotropic in momen-

tum space. Abrahams and Varma [2] argue that a good form of the self-energy is given by the following expression:

$$\Sigma''(\mathbf{k}, \omega, T) = \Gamma_a(\mathbf{k}) + \lambda \sqrt{\omega^2 + \pi^2 T^2}. \quad (5.19)$$

In this equation, the first term is the scattering rate due to disorder and is momentum-dependent and temperature- and frequency-independent. The second term contains the marginal Fermi liquid self-energy and is momentum-independent. The anisotropic elastic scattering rate is argued to arise from impurities located away from the copper-oxide planes, which induce only small-momentum scattering. To model this anisotropic scattering rate, we take it to have the following functional form:

$$\Gamma_a(\mathbf{k}) = \beta \left(\frac{2 + \cos 2k_x + \cos 2k_y}{4} \right). \quad (5.20)$$

This form of the scattering rate ensures that it is small near the nodes—it is zero at $(\pm\frac{\pi}{2}, \pm\frac{\pi}{2})$, in fact—and that it has maxima at $(0, \pm\pi)$ and $(\pm\pi, 0)$. Importantly, this form preserves all the symmetries of the square lattice. The choice $\beta = 0.2$ gives rise to EDCs which show large anisotropy between the nodal and antinodal points on the Fermi surface, as seen in Fig. 5.17.

Plots of $A(\mathbf{k}, \omega = 0)$ and $P(\mathbf{q}, \omega = 0)$ for this anisotropic marginal Fermi liquid at a variety of temperatures are shown in Fig. 5.18. Note first that the spectral function at the near-nodal region is fairly sharp, while moving towards the antinodes we see that much more broadening becomes present, with considerable spectral weight being present in the regions between the Fermi surface at the antinodal regions. In the isotropic cases we considered earlier, there is zero spectral weight in these regions, as these parts of the Brillouin zone lie far beyond the




Figure 5.18: The spectra of an anisotropic marginal Fermi liquid at various temperatures. Left to right: The spectral function $A(\mathbf{k}, \omega)$; the Fourier transform of the LDOS $P(\mathbf{q}, \omega)$; linecuts of $P(\mathbf{q}, \omega)$ in the nodal and antinodal directions; $P(\mathbf{q}, \omega)$ with finite-temperature smearing; and $P(\mathbf{q}, \omega)$ in the presence of many weak impurities. All plots are taken at $E = 0.000$. For ease of visualization, the scale used here is smaller than that used in the Fermi liquid and isotropic marginal Fermi liquid plots.




Figure 5.19: Momentum-distribution curves taken along nodal $((0,0) \rightarrow (\pi, \pi))$ and antinodal $((0,0) \rightarrow (0, \pi))$ cuts in the Brillouin zone at the Fermi energy ($E = 0$) for an anisotropic marginal Fermi liquid. Here $T = 100$ K and $\beta = 0.2$ (see Eq. 5.20 for the functional form of the self-energy).

bare Fermi surface, but with considerable nodal-antinodal anisotropy the antinodal regions become blurred and nonzero spectral weight results. This is even more apparent if we take momentum-distribution curves along the nodal and antinodal directions, as plotted in Fig. 5.19: the MDCs along the nodal direction are quite sharp, while those along the antinodal directions are far more incoherent, although traces of peaks remain—in good agreement with ARPES experiments, which still find these antinodal peaks present in MDCs, albeit in a far weaker state compared to those at the nodes.

The LDOS power spectrum in Fig. 5.18 has a number of interesting features worth commenting upon. First, there is a very fuzzy square-shaped central plateau which is formed from small-momenta scattering processes between antinodal portions of the Fermi surface. Because the broadening is very large at the antinodal points, the scattering wavevectors appearing in $P(\mathbf{q}, \omega)$ consequently are severely broadened as well. Second, there is a set of very sharp features near $(\pm\pi, \pm\pi)$ which arise from internodal scattering. Recall that the spectral function remains sharp and well-defined near the nodes. As such, scatter-

ing wavevectors between near-nodal regions remain sharp in $P(\mathbf{q}, \omega)$, unlike those from antinodal-antinodal scattering. If one traces the caustics extending beyond the central plateau carefully, accounting for the backfolding of the spectra, one can make out that they decrease in width as one moves from the antinodal scattering wavevectors to the nodal ones. Linecuts along the nodal and antinodal directions are perhaps even more illuminating. The linecuts along the antinodal directions are featureless, save for the aforementioned plateau region, while the nodal linecuts show a sharp peak near the Brillouin zone boundary corresponding to nodal-nodal scattering. The contrast with the isotropic marginal Fermi liquid is quite striking, as the isotropic case (Fig. 5.12) features a central plateau which is still fairly sharply defined, while the caustics which appear beyond the plateau are of uniform width. With random disorder and thermal smearing, the resulting spectra appear very noisy—owing in part to the large intrinsic broadening at the antinodes. The central plateau visible in the single-impurity results is no longer easily seen, but there do remain sharp peaks near the zone diagonals corresponding to nodal-nodal scattering wavevectors, visible even when finite-temperature smearing is included.

Finally we note that because the frequency-dependence of the self-energy in this case is similar to the isotropic marginal Fermi liquid case considered earlier, the widths of the LDOS power spectra here should behave in the same way. This can be seen in plots of the spectral function and the LDOS power spectra at 100 K as frequency is varied, as seen in Fig. 5.20. As frequency is increased, the spectral function broadens throughout momentum space, and the resulting caustics in the LDOS power spectrum similarly broaden as well. This increased broadening at large frequencies contributes to the loss of signal in the disordered and thermally broadened data.




Figure 5.20: Frequency-dependence at $T = 100$ K of the spectra of an anisotropic marginal Fermi liquid. Shown are plots of the spectral function $A(\mathbf{k}, \omega)$ (upper row); the LDOS power spectrum with a single pointlike scatterer without thermal smearing (middle row); and the LDOS power spectrum with both a 0.5% concentration of pointlike scatterers and thermal smearing (bottom row). Note that the scales used for plotting the LDOS power spectra are the same for all frequencies. For ease of visualization, the scale used here is smaller than that used in the Fermi liquid and isotropic marginal Fermi liquid plots.

We caution the reader that because the anisotropic part of the self-energy here is presumably due to elastic scattering off of off-plane impurities, there is the danger that unless disorder is carefully taken into account, “double-counting” may ensue. As QPI is an intrinsically disorder-driven effect, one has to take care in these simulations that the same disorder producing QPI does not contribute additionally to the anisotropic elastic self-energy. We have taken care to use only pointlike impurities in our simulations of QPI, and the effect of the off-plane impurities is incorporated in the anisotropic elastic scattering rate. A single weak pointlike impurity represents a very small perturbation to the system whose overall effect is negligible, while a dilute ensemble of pointlike scatterers would presumably contribute to an *isotropic* scattering rate, adding only a momentum-independent constant into the full self-energy upon disorder averaging.

5.6 DISCUSSION AND CONCLUSION

We have provided in this chapter a comprehensive overview of the effects of self-energies on quasiparticle scattering interference, and have applied much of this insight to situations of relevance to the copper-oxide superconductors. While self-energies have been well-understood from the perspective of ARPES experiments, their effects on STS experiments have not been as similarly understood and are largely unexplored. A consistent result seen in the many scenarios we considered in this chapter is the destruction of the QPI signal as broadening is increased, even when thermal smearing is ignored. In many ways, this is not an unexpected result. The physics underlying the phenomenological octet model of QPI in the superconducting state of the cuprates relies on the existence of coherent quantum-mechanical waves, which

scatter elastically against impurities present in these materials. If these quasiparticles have a short lifetime—as seen in ARPES experiments in the strange metal, or even at the superconducting state near T_c —then the rather simple picture suggested by the octet model becomes complicated by “off-shell” contributions to the full, disordered Green’s function. That is, in the presence of large broadening, states living away from the contours of constant energy *do* contribute towards the scattering processes which determine the structure of the LDOS and its power spectrum. These effects are in fact already visible in the spectral function itself. We have seen that the contours of constant energy in both the normal and superconducting states turn from sharp, well-defined structures in momentum space into broad, incoherent entities. The effects of this broadening are particularly dramatic in the *d*-wave superconducting state, where we see that the sharp banana-shaped contours seen in the spectral function turn into incoherent arc-like streaks once the quasiparticle scattering rate is of the same order of magnitude as the superconducting gap. The loss of the sharpness in the contours of constant energy translates directly into the smearing and progressive destruction of the octet-model peaks as the scattering rate is increased.

The normal-state LDOS spectra feature no such peak-like structures, and instead what appears is a set of caustics which are continuous and whose broadening as a function of position *on* the caustics directly reflects the degree of coherence of the quasiparticles of the underlying Fermi surface. As such, in the normal state the LDOS power spectrum is far less sensitive to broadening than in the superconducting state. One can see that the main feature differentiating the marginal Fermi liquid from the ordinary Fermi liquid is the amount of broadening present in both the spectral function and the LDOS power spectrum—the marginal Fermi liquid, by virtue of the fact that the imaginary

part of its self-energy is much larger than that of the ordinary Fermi liquid at the same temperature and frequency, shows much more intrinsic smearing in its spectra—and how this broadening depends on temperature and frequency. Another measurable effect is the renormalization of the dispersions, due to the Kramers-Kronig relations, which can in principle be measured directly. Nevertheless this is a rather subtle effect—the bare band structure needs to be known in order for this renormalization to be detected—and given the difficulty experimentalists are sure to face in attempting to observe this effect in STS experiments, the main signal of interest is the width of the measured power spectra.

It is worth explaining further in this section the limitations of our explicitly phenomenological approach. Our starting point consists of mean-field models of the normal and superconducting states, which are then “dressed” by self-energies which have a nontrivial dependence on temperature, frequency, momentum, or some combination of these. The predictions we make in this chapter for STS—and, for that matter, ARPES as well—are sensible only if the actual strongly correlated phases seen in the cuprates can be adequately described by these dressed mean-field models. Much work on the two-dimensional Hubbard model, using dynamical mean-field theory, has shown that this “dressed” picture, involving a single-particle propagator augmented by a nontrivial self-energy, provides a reasonably accurate picture of the physics in some phases of relevance to the cuprates [29, 57, 147]. If such a picture were to hold, then QPI will exist in some form or another. For instance, a model of the pseudogap involving a broadened *d*-wave superconductor will show QPI with the octet-model peaks decohering; nevertheless, despite the absence of sharp peaks, the power spectrum should still consist of wavevectors describing the relevant scattering processes. As another example, the marginal Fermi liquid is

an exotic phase of matter without any low-temperature quasiparticle-like excitations—we remind the reader that its quasiparticle weight vanishes at the Fermi surface at $T = 0$ —but still features ARPES and STS spectra that, at face value, are similar to those of an ordinary Fermi liquid.

Having said all of this, if the phase of matter is *not* describable at all by this dressed mean-field picture, there is no sense in which any of our predictions should hold. In particular, if STS were to show no evidence of these caustics in the strange-metal phase of the cuprates, then that would be one extremely convincing piece of evidence to suggest that the strange metal phase is beyond even the marginal-Fermi-liquid description. Hints of this have in fact been seen in STS studies deep inside the superconducting state: at energies larger than the superconducting gap, no well-defined caustics are seen, and instead the most dominant features are peaks corresponding to charge ordering [90, 99, 50]. In such a scenario, the appropriate theory is a strongly interacting phase of matter whose low-energy excitations are very unlike the Landau quasiparticles of the Fermi liquid. A paradigmatic example of this is the Luttinger liquid in one spatial dimension [60], whose decidedly non-quasiparticle-like excitations result in ARPES and STS spectra considerably different from those of an ordinary Fermi liquid [127, 87]. In addition, numerous examples of these phases have been constructed using holographic methods, and are known to result in physics very different from that of the ordinary Fermi liquid [79, 35]. We end by noting that what high-temperature STS experiments can eventually find in the strange metal phase and in the transition to the superconducting state at optimal doping will undoubtedly be very interesting. The insights that can be gleaned from such future experiments will no doubt go a long way in illuminating the strange physics of the cuprate superconductors.

CONCLUSIONS AND OUTLOOK

This thesis is devoted to an examination of the various ways in which disorder impacts the electronic properties of the high-temperature superconductors. It is worth stepping back to collect many of our results here at one place and to situate them within the broader understanding of the cuprates. As we have noted in the introduction, disorder is a prominent actor in a number of mysterious phenomena in these materials. To ensure that one properly interprets experimental data, it is essential that all the possible effects of disorder are taken into account. Our results demonstrate that what is seen in experiment is in many ways surprising and more subtle than the simplest and most analytically tractable models can predict.

6.1 THE UNREASONABLE EFFECTIVENESS OF QPI

In Chapter 3 we discussed at length how quasiparticle scattering interference is generated by distributed disorder—taking a massive step beyond the simple single-impurity models that have consistently formed the basis for much of the theoretical work done on QPI. We discover the seemingly paradoxical result that our simulations of STS, using various realistic models of disorder, fail to replicate the well-defined

peaks seen in the experimental STS power spectrum. In addition to peaks, our simulations show that the streaky patterns characteristic of the single-impurity spectra should also be seen in the macroscopically-disordered systems we have studied—and the fact that these streaks are much less visible in the actual experimental data is highly unusual. On one hand, the fact that we do see the same dispersing peaks as experiment does mean that the physics of QPI as presently understood—the scattering of coherent d -wave quasiparticles against disorder of some sort—is fundamentally correct. On the other hand, the fact that the real-world signals are far sharper than anything we can find in numerics is a sign that more details have to be added to this scattering picture of QPI. It is not clear at first glance why these peaks should be the dominant signal.

Perhaps the most sensible explanation is that the tunneling process between the STM tip and the CuO_2 plane—known to be highly nontrivial—is responsible for the enhancement of the sharpness of the peak-like features in the power spectra. We had in fact considered a simple model of the tunneling process in Chapter 3 and found no notable enhancement of the peaks over those seen in the trivial-tunneling case. However there is good reason to suspect that a fuller, more microscopic model of the tunneling process, derived from first-principles considerations, could explain some of these anomalies away. Single-impurity simulations of STS spectra show that the characteristic real-space pattern of the differential conductance around a zinc impurity in BSCCO-2212 is reproduced very well if one incorporates a very accurate model of tip-plane tunneling into the calculation of the LDOS. An interesting direction for future work along these lines is the incorporation of this realistic tunneling process in the simulation of a macroscopically-disordered d -wave superconductor to see if sharp peaks appear in the speckled many-impurity power spectrum. We had

seen that random disorder of whatever form results in speckle in the power spectrum, degrading the sharpness of the signal. As distributed disorder is an intrinsic feature of the cuprates, it is important to know if the sharpness of the octet-model peaks can ever be recovered in the disordered case even with a realistic model of STM tunneling.

Another result we have obtained is that of all the models of disorder we have considered, two models in particular—a random array of weak pointlike scattering centers and randomly-distributed but uncorrelated on-site energies—reproduce with the most success the experimentally-observed QPI power spectrum. This is not entirely unexpected, as it is known that the cuprates studied by STS generally do not feature strong in-plane impurities unless doped by zinc or nickel, so any disorder is necessarily weak. However, what *is* surprising is that the most plausible model of disorder, at least from a chemical standpoint—smooth disorder sourced by off-plane dopants—does not give rise to the requisite large-momentum octet-model peaks. Rather, the disorder required to reproduce on a phenomenological level the octet-model peaks is pointlike in nature, suggesting that QPI as seen in experiments is likely not due to off-plane disorder. However, the origin of this weak in-plane disorder is not immediately obvious. The copper-oxide planes can feature defects whose effect should mimic that of a local impurity, and deformations of the lattice could also lead to the large-momentum peaks. Nevertheless these are not clear-cut isolated impurities. This stands in contrast to the impurities which lead to QPI in many non-cuprate materials such as the surfaces of three-dimensional topological insulators, which can be visualized clearly with STS.

Taking these two results together, it seems almost miraculous in hindsight that QPI has even been seen in the cuprates! The success of QPI in revealing the details of the momentum-space electronic struc-

ture of the cuprates notwithstanding, there is much about QPI that still eludes understanding. There is a shroud of microscopic details surrounding the microscopic tunneling process which prevents one from naively matching the bare LDOS calculations to experiment without plenty of caveats. One can only hope that more knowledge of these microscopic details can resolve the paradoxes present in the theory of QPI in the cuprates. As we have noted earlier, nothing invalidates the basic picture of QPI as quantum-mechanical waves forming ripples as they pass through a disordered medium and interfere with each other, from which the basis for the octet model can be formed, so QPI as we know it remains a very good probe of electronic structure in the cuprates. The fact remains, however, that beyond the scattering picture (which remains accurate) it is still a black-box-like experimental method—much of what we can see through STS is affected by nonuniversal microscopic details. (Note that we have not even touched on the issue of QPI extinction—it remains its own distinct can of worms!)

6.2 DISORDER: OLD DOG, NEW TRICKS

Chapter 4 takes a look at disorder and its effect on the low-energy quasiparticle density of states, a quantity which can be directly measured by specific heat experiments. As with Chapter 3, we take the perspective of revisiting this old, ostensibly well-understood problem, with the advantage that we now can put in realistic forms of disorder without being constrained by analytical tractability. A key motivation is to examine the impact of smooth disorder, which, as noted earlier, should be ubiquitous in the cuprates due to the presence of dopants within the buffer layers of these materials and whose impact on the density of states has not been studied in prior theoretical work.

The most important result we find in that chapter is that, in the presence of smooth disorder due to off-plane impurities, a finite DOS at the Fermi energy pops up naturally without fundamentally reshaping the structure of the DOS at higher energies. This is in stark contrast to in-plane pointlike disorder, which is seen to affect the higher-energy DOS far more strongly than smooth disorder (even at high off-plane dopings) does. The spectral-weight transfer is concentrated within the Fermi energy to an unusual degree: as the amount of disorder is increased, a sharp resonance forms at $E = 0$, which quickly dies off with increasing energy, all while leaving much of the rest of the quasiparticle excitation spectrum unaltered.

It is important to note that smooth disorder has not been typically considered as an explanation for the finite DOS at the Fermi energy as seen in specific heat experiments on YBCO. The well-known “dirty d -wave” model assumes that the disorder potential is pointlike; however, with smooth disorder, the T -matrix approximation used for point disorder fails, and one thus has to solve the Bogoliubov-de Gennes Hamiltonian exactly to obtain the density of states. Some of the earliest works on disordered d -wave superconductors have assumed that the relatively weak disorder potentials due to off-plane impurities can be modeled in the Born approximation, but it is clear that the Born approximation fails to describe the small-angle scattering that is the dominant feature of such smooth potentials, and as such cannot replicate the gentler spectral-weight transfers due to smooth disorder which are seen in our numerics.

Given that we have seen that smooth disorder is a plausible explanation for the low-energy excitations seen in specific heat experiments, it is worth asking if other experimental results which have been explained in terms of pointlike models of disorder could be retroactively explained using a smooth-disorder model instead. In this light, per-

haps the most interesting future direction is the study of transport properties in the presence of smooth disorder. The thermal and optical conductivities of the cuprates in the superconducting state have been studied quite exhaustively in experiment, but a theoretical understanding of these results based on a smooth-disorder paradigm is limited at the moment. It is also worth considering just how strong the pair-breaking effect of smooth disorder is. The single-particle quantities we have studied in Chapter 4 suggest that smooth disorder has a much softer imprint than pointlike disorder does—witness just how well-preserved the coherence peaks are even at large dopings—which makes it quite likely that off-plane disorder leads to far less pair-breaking. Finally, it is very interesting to see the extent to which the NMR Knight shift experiments alluded to in Chapter 4—which measure the distribution of the DOS at $E = 0$ —can be explained by off-plane disorder.

We finally end this section by noting that on a semantic level, what constitutes “disorder” in the cuprates is surprisingly tricky. Many recently synthesized YBCO samples have been described as “clean.” However it is generally not appreciated that any dopant situated off the planes *will* generate disorder—it just acts on the copper-oxide planes in an indirect manner by means of a Coulomb potential, and a sufficiently large number of them creates enough smooth disorder in the plane to perhaps generate the effects discussed in Chapter 4. Even the ordered phases of YBCO (*e.g.*, ortho-II, where half of the copper-oxygen chains are filled and half are empty) feature off-plane oxygen dopants in both chains and buffer layers, and the net effect of all of these dopants taken together is to create a random but smooth disorder potential within the copper-oxygen planes. Our results in Chapter 4 provide an object lesson in the way a seemingly invisible form of disorder can still generate enough low-lying quasiparticle excitations

to alter the electronic properties near the Fermi energy, while ensuring that the d -wave superconducting state by and large remains untouched at higher energies—which appears to be the case in the actual cuprates.

6.3 STRETCHING QPI TO ITS BREAKING POINT

In Chapter 5 we studied the impact of self-energies on the QPI power spectrum and the spectral function of both a d -wave superconductor and a metal, with the goal of developing testable predictions for STS experiments on the cuprates at temperatures near and above T_c . For the d -wave superconductor, we contrasted the “gap-filling” phenomenology seen in the cuprates—where the gap becomes filled with low-energy excitations as one nears T_c , with the gap fully closing only at temperatures above T_c —with BCS-like “gap-closing,” in which the gap closes *at* T_c . The general lesson as far as the d -wave superconductor is concerned is that the presence of a large self-energy is highly detrimental to the octet-model peaks, which become smeared to the point of incoherence as the scattering rate is increased. Nevertheless, even though the octet-model peaks are lost once scattering rates are large enough, the QPI power spectrum retains “memories” of the d -wave gap—the QPI power spectrum of a broadened d -wave superconductor is manifestly different from that of a normal metal, even if their spectral functions begin to resemble each other more and more.

This is an instance in which STS provides a way of sharply distinguishing a finite superconducting gap from a vanishing one above T_c . ARPES experiments have some difficulty resolving this because at the high temperatures being considered here ($T \sim T_c$), the quasiparticle scattering rate is large, and consequently the peaks seen in EDCs become blurred and incoherent, and thus one can misidentify a phase

with a finite superconducting gap as a phase whose gap vanishes if the peaks marking the locations of the quasiparticle excitations are no longer discernible. As noted earlier, one hypothesis for the existence of Fermi arcs in the pseudogap is that these are simply d -wave nodes that have been smeared thanks to the very large scattering rate inherent in this phase. The large scattering rates near T_c appear to be a consistent feature of the underdoped and optimally-doped cuprates. QPI, because of its sensitivity to coherence factors, does see this difference: when one has a broadened d -wave superconductor, the QPI power spectrum is highly anisotropic, with spectral weight spread nonuniformly across the caustics.

Chapter 4 shows through explicit calculations that QPI does have a “breaking point” as far as its usefulness as a probe of the momentum-space structure is concerned. When the lifetimes of the Bogoliubov quasiparticles become short, these do not yield sharp peaks, with only broad and incoherent features remaining in the power spectrum. This makes sense if one invokes the simple picture of QPI as quantum-mechanical waves interfering with each other. If the quasiparticle lifetime is short, then there is no sense in which these excitations can be treated as coherent propagating waves—the quasiparticles die before any interference ensues. As the main signal of QPI in a d -wave superconductor consists of peaks, the elimination of this signal at sufficiently high temperatures can be taken as one piece of evidence of the loss of coherence of these particles. The ease with which this signal is degraded further points to the stability of the Bogoliubov quasiparticles deep within the superconducting phase at low temperatures, which yield sharp octet-model peaks and are thus identifiably long-lived excitations. In this way, the breakdown of QPI can itself be argued to be a piece of evidence suggesting that the long-lived nature of the excitations is lost. Viewed through the lens of our results, this bolsters

the interpretation that QPI extinction in the underdoped cuprates is caused by the loss of coherence of the Bogoliubov quasiparticles at the antinodal regime. However, the extent to which this interpretation is consistent with ARPES results is unclear—recall that ARPES sees coherent quasiparticles across the entire Fermi surface. It has to be said, however, that QPI, as it rests on a wave-like picture involving the interference of propagating waves, is more sensitive to the coherence of the excitations than ARPES is.

Having said all this, we find that the marginal Fermi liquid is surprisingly stable insofar as it leads to fairly detectable features in the QPI power spectrum. Compared to the ordinary Fermi liquid, the caustics for the MFL are broadened, but still discernable. The main difference between the MFL and the ordinary Fermi liquid as far as QPI is concerned is quite subtle: the QPI power spectrum of the former has much more broadening than the latter, but despite this and the diminished intensity in the power spectrum, the shape of the caustics can still be discerned clearly. Further differences between these two cases come in the form of a different dependence of the widths of the caustics on temperature and frequency. However, none of these differences are particularly sharp at first glance. If one takes a look at QPI spectra at fixed temperature and frequency for the MFL and the FL, one is hard-pressed to identify the differences between them, unless spectra at other temperatures and frequencies are obtained—and even then, the differences still do not jump out of the page. In fact, these subtle (rather than sharp) differences between the MFL and the FL spectra are already apparent in ARPES results taken in the nodal regime, which have been interpreted as the spectra of an anisotropic marginal Fermi liquid. Because ARPES experiments on the normal state of cuprates are done at high temperatures ($T > 100$ K), the manifestly ω -linear behavior of the widths of the spectral function is appar-

ent only at frequencies $|\omega| > \pi T$. It is in this regime where the MFL is most different from the ordinary FL, for which the widths scale as ω^2 . The lesson here is that at finite temperatures the MFL is only subtly different from the FL, and some of the sharper distinctions one can make between the two phases at $T = 0$ are rounded off, requiring fine analyses of the temperature- and frequency-dependence of the observed spectra to tease out these differences.

Finally, we note that an exciting outcome might be—ironically—a *negative* result disproving our predictions for QPI in a marginal Fermi liquid. It is highly likely that the normal state of the cuprates is ultimately described by a theory which is as different from Fermi-liquid theory as one can get. It is hard to imagine what the spectra for such a state should look like, but as we have noted in Chapter 2, a real-world paradigm of a non-Fermi liquid without any quasiparticle-like excitations is the Luttinger liquid in one spatial dimension. It has been shown to lead to a number of anomalous features: MDCs which are sharp, even with strong interactions; EDCs which broaden as the interaction strength is increased, with no quasiparticle peaks visible; and QPI spectra which show dispersing features due to spin-charge-separated excitations collectively formed from the underlying electrons. The Luttinger liquid in this sense goes even further than the MFL in violating the basic features of the noninteracting Fermi gas. The MFL at its core is a dressed Fermi gas—no long-lived excitations exist, but its correlation functions can nevertheless be understood in terms of the electronic excitations forming the Fermi gas which are renormalized away by the MFL self-energy. In a Luttinger liquid, on the other hand, the basic excitations are collective modes which separately carry charge and spin—about as far removed from the electrons as possible!

While we are agnostic as to what the ultimate theory of the normal state is, if one considers the very existence of the Luttinger liquid—a state of matter which has been seen in experiments to be almost unrecognizably different from the weakly-interacting Fermi liquid (and which happens to be the most easily solvable of all models at that!)—it is clear that one has to square oneself with the possibility that QPI for the normal state of the cuprates is dominated by features not at all related to the noninteracting Fermi gas, but instead by some set of exotic, hitherto unknown excitations which reflect the quantum-critical nature of the system. It has been noted early on in this thesis that examples of these finite-density quantum-critical states of matter have been constructed using nonperturbative methods—the most prominent of which is holography. Thus, if one does not see the caustics we predicted will form the QPI spectra of the MFL, this negative experimental result is a strong indication of the highly quantum-critical nature of the normal state—so quantum-critical that there is no sense in it being described even by the fairly radical MFL theory! We can only speculate what this theory will look like, but one remains hopeful that the features of this strange beast can be understood, or at least sketched out roughly, by STS experiments high up in the strange-metal phase.

BIBLIOGRAPHY

- [1] E. Abrahams. The normal state of high- T_c superconductors and the marginal Fermi liquid. 2001.
- [2] E. Abrahams and C. M. Varma. What angle-resolved photoemission experiments tell about the microscopic theory for high-temperature superconductors. *Proceedings of the National Academy of Sciences*, 97(11):5714–5716, 2000.
- [3] E. Abrahams and C. M. Varma. Hall effect in the marginal Fermi liquid regime of high- T_c superconductors. *Phys. Rev. B*, 68(9):094502, 2003.
- [4] A. Allais and T. Senthil. Loop current order and d-wave superconductivity: Some observable consequences. *Phys. Rev. B*, 86(4):045118, 2012.
- [5] M. P. Allan, T. Chuang, F. Massee, Y. Xie, N. Ni, S. L. Budko, G. S. Boebinger, Q. Wang, D. S. Dessau, P. C. Canfield, M. S. Golden, and J. C. Davis. Anisotropic impurity states, quasiparticle scattering and nematic transport in underdoped $\text{Ca}(\text{Fe}_{1-x}\text{Co}_x)_2\text{As}_2$. *Nature Physics*, 9(4):220–224, 2013.
- [6] M. P. Allan, K. Lee, A. W. Rost, M. H. Fischer, F. Massee, K. Kihou, C.-H. Lee, A. Iyo, H. Eisaki, T.-M. Chuang, J. C. Davis, and E. A. Kim. Identifying the ‘fingerprint’ of antiferromagnetic spin fluctuations in iron pnictide superconductors. *Nature Physics*, 11(2):177–182, 2015.

Bibliography

- [7] M. P. Allan, F. Massee, D. K. Morr, J. Van Dyke, A. W. Rost, A. P. Mackenzie, C. Petrovic, and J. C. Davis. Imaging Cooper pairing of heavy fermions in CeCoIn_5 . *Nature Physics*, 9(8):468–473, 2013.
- [8] M. P. Allan, A. W. Rost, A. P. Mackenzie, Y. Xie, J. C. Davis, K. Kihou, C. H. Lee, A. Iyo, H. Eisaki, and T.-M. Chuang. Anisotropic energy gaps of iron-based superconductivity from intraband quasiparticle interference in LiFeAs . *Science*, 336(6081):563–567, 2012.
- [9] J. W. Alldredge, J. Lee, K. McElroy, M. Wang, K. Fujita, Y. Kohsaka, C. Taylor, H. Eisaki, S. Uchida, P. J. Hirschfeld, and J. C. Davis. Evolution of the electronic excitation spectrum with strongly diminishing hole density in superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+x}$. *Nature Physics*, 4(4):319, 2008.
- [10] A. Altland, B. D. Simons, and M. R. Zirnbauer. Theories of low-energy quasi-particle states in disordered d-wave superconductors. *Physics Reports*, 359(4):283–354, 2002.
- [11] B. M. Andersen and P. J. Hirschfeld. Extinction of quasiparticle interference in underdoped cuprates with coexisting order. *Phys. Rev. B*, 79(14):144515, 2009.
- [12] P. W. Anderson. Theory of dirty superconductors. *Journal of Physics and Chemistry of Solids*, 11(1-2):26–30, 1959.
- [13] W. A. Atkinson, P. J. Hirschfeld, and A. H. MacDonald. Gap inhomogeneities and the density of states in disordered d-wave superconductors. *Phys. Rev. Lett.*, 85(18):3922, 2000.
- [14] W. A. Atkinson, P. J. Hirschfeld, A. H. MacDonald, and K. Ziegler. Details of disorder matter in $2d$ d-wave superconductors. *Phys. Rev. Lett.*, 85(18):3926, 2000.

- [15] W. A. Atkinson, P. J. Hirschfeld, and L. Zhu. Quantum interference in nested d-wave superconductors: A real-space perspective. *Phys. Rev. B*, 68(5):054501, 2003.
- [16] A. V. Balatsky, M. I. Salkola, and A. Rosengren. Impurity-induced virtual bound states in d-wave superconductors. *Phys. Rev. B*, 51(21):15547, 1995.
- [17] A. V. Balatsky, I. Vekhter, and J.-X. Zhu. Impurity-induced states in conventional and unconventional superconductors. *Rev. Mod. Phys.*, 78(2):373, 2006.
- [18] J. G. Bednorz and K. A. Müller. Possible high T_c superconductivity in the Ba-La-Cu-O system. *Zeitschrift für Physik B Condensed Matter*, 64(2):189–193, 1986.
- [19] C. Bena, S. Chakravarty, J. Hu, and C. Nayak. Quasiparticle scattering and local density of states in the d-density-wave phase. *Phys. Rev. B*, 69(13):134517, 2004.
- [20] C. Bena and S. A. Kivelson. Quasiparticle scattering and local density of states in graphite. *Phys. Rev. B*, 72(12):125432, 2005.
- [21] E. Berg, C.-C. Chen, and S. A. Kivelson. Stability of nodal quasiparticles in superconductors with coexisting orders. *Phys. Rev. Lett.*, 100(2):027003, 2008.
- [22] J. S. Bobowski, J. C. Baglo, J. Day, L. Semple, P. Dosanjh, P. J. Turner, R. Harris, R. Liang, D. A. Bonn, and W. N. Hardy. Oxygen chain disorder as the weak scattering source in $\text{YBa}_2\text{Cu}_3\text{O}_{6.50}$. *Phys. Rev. B*, 82(13):134526, 2010.

Bibliography

- [23] B. R. Bulka, B. Kramer, and A. MacKinnon. Mobility edge in the three dimensional Anderson model. *Zeitschrift für Physik B Condensed Matter*, 60(1):13–17, 1985.
- [24] J. C. Campuzano, M. R. Norman, and M. Randeria. Photoemission in the High- T_c Superconductors. In *The Physics of Superconductors*, pages 167–273. Springer, 2004.
- [25] L. Capriotti, D. J. Scalapino, and R. D. Sedgewick. Wave-vector power spectrum of the local tunneling density of states: Ripples in a d-wave sea. *Phys. Rev. B*, 68(1):014508, 2003.
- [26] E. W. Carlson, D. Orgad, S. A. Kivelson, and V. J. Emery. Dimensional crossover in quasi-one-dimensional and high- T_c superconductors. *Phys. Rev. B*, 62(5):3422, 2000.
- [27] S. Chakravarty, C. Nayak, and S. Tewari. Angle-resolved photoemission spectra in the cuprates from the d-density wave theory. *Phys. Rev. B*, 68(10):100504, 2003.
- [28] A. V. Chubukov, M. R. Norman, A. J. Millis, and E. Abrahams. Gapless pairing and the Fermi arc in the cuprates. *Phys. Rev. B*, 76(18):180501, 2007.
- [29] M. Civelli, M. Capone, A. Georges, K. Haule, O. Parcollet, T. D. Stanescu, and G. Kotliar. Nodal-antinodal dichotomy and the two gaps of a superconducting doped Mott insulator. *Phys. Rev. Lett.*, 100(4):046402, 2008.
- [30] S. J. Collocott and R. Driver. Specific heat of the high- T_c superconductor $(\text{Bi, Pb})_2\text{Sr}_2\text{Ca}_2\text{Cu}_3\text{O}_{10}$, and related phases Ca_2CuO_3 and $(\text{Ca}_{0.86}\text{Sr}_{0.14})\text{CuO}_2$ from 0.4 to 20 K. *Physica C: Superconductivity*, 167(5-6):598–608, 1990.

- [31] M. F. Crommie, C. P. Lutz, and D. M. Eigler. Imaging standing waves in a two-dimensional electron gas. *Nature*, 363(6429):524–527, 1993.
- [32] E. G. Dalla Torre, D. Benjamin, Y. He, D. Dentelski, and E. Demler. Friedel oscillations as a probe of fermionic quasiparticles. *Phys. Rev. B*, 93(20):205117, 2016.
- [33] A. Damascelli. Probing the electronic structure of complex systems by ARPES. *Physica Scripta*, 2004(T109):61, 2004.
- [34] A. Damascelli, Z. Hussain, and Z.-X. Shen. Angle-resolved photoemission studies of the cuprate superconductors. *Rev. Mod. Phys.*, 75(2):473, 2003.
- [35] R. A. Davison, K. Schalm, and J. Zaanen. Holographic duality and the resistivity of strange metals. *Phys. Rev. B*, 89(24):245116, 2014.
- [36] H. Ding, M. R. Norman, J. C. Campuzano, M. Randeria, A. F. Bellman, T. Yokoya, T. Takahashi, T. Mochiku, and K. Kadowaki. Angle-resolved photoemission spectroscopy study of the superconducting gap anisotropy in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+x}$. *Phys. Rev. B*, 54(14):R9678, 1996.
- [37] N. Doiron-Leyraud, C. Proust, D. LeBoeuf, J. Levallois, J.-B. Bonnemaison, R. Liang, D. A. Bonn, W. N. Hardy, and L. Taillefer. Quantum oscillations and the Fermi surface in an underdoped high- T_c superconductor. *Nature*, 447(7144):565, 2007.
- [38] P. S. Drouvelis, P. Schmelcher, and P. Bastian. Parallel implementation of the recursive Greens function method. *Journal of Computational Physics*, 215(2):741–756, 2006.

Bibliography

- [39] A. C. Durst and P. A. Lee. Impurity-induced quasiparticle transport and universal-limit Wiedemann-Franz violation in d-wave superconductors. *Phys. Rev. B*, 62(2):1270, 2000.
- [40] A. C. Durst and P. A. Lee. Microwave conductivity due to scattering from extended linear defects in d-wave superconductors. *Phys. Rev. B*, 65(9):094501, 2002.
- [41] A. Einstein. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der physik*, 322(6):132–148, 1905.
- [42] H. Eisaki, N. Kaneko, D. L. Feng, A. Damascelli, P. K. Mang, K. M. Shen, Z.-X. Shen, and M. Greven. Effect of chemical inhomogeneity in bismuth-based copper oxide superconductors. *Phys. Rev. B*, 69(6):064512, 2004.
- [43] V. J. Emery, E. Fradkin, S. A. Kivelson, and T. C. Lubensky. Quantum theory of the smectic metal state in stripe phases. *Phys. Rev. Lett.*, 85(10):2160, 2000.
- [44] V. J. Emery and S. A. Kivelson. Importance of phase fluctuations in superconductors with small superfluid density. *Nature*, 374(6521):434–437, 1995.
- [45] A. Fang, C. Howald, N. Kaneko, M. Greven, and A. Kapitulnik. Periodic coherence-peak height modulations in superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. B*, 70(21):214514, 2004.
- [46] D. L. Feng, D. H. Lu, K. M. Shen, C. Kim, H. Eisaki, A. Damascelli, R. Yoshizaki, J.-i. Shimoyama, K. Kishio, G. D. Gu, S. Oh, A. Andrus, J. O'Donnell, J. N. Eckstein, and Z.-X. Shen. Signature of superfluid density in the single-particle excitation spectrum of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 289(5477):277–281, 2000.

- [47] M. Franz, C. Kallin, and A. J. Berlinsky. Impurity scattering and localization in d-wave superconductors. *Phys. Rev. B*, 54(10):R6897, 1996.
- [48] M. Franz, C. Kallin, A. J. Berlinsky, and M. I. Salkola. Critical temperature and superfluid density suppression in disordered high- T_c cuprate superconductors. *Phys. Rev. B*, 56(13):7882, 1997.
- [49] L. Fu. Hexagonal warping effects in the surface states of the topological insulator Bi_2Te_3 . *Phys. Rev. Lett.*, 103(26):266801, 2009.
- [50] K. Fujita, C. K. Kim, I. Lee, J. Lee, M. H. Hamidian, I. A. Firmo, S. Mukhopadhyay, H. Eisaki, S. Uchida, M. J. Lawler, E.-A. Kim, and J. C. Davis. Simultaneous transitions in cuprate momentum-space topology and electronic symmetry breaking. *Science*, 344(6184):612–616, 2014.
- [51] K. Fujita, A. R. Schmidt, E.-A. Kim, M. J. Lawler, D. Hai Lee, J. C. Davis, H. Eisaki, and S.-i. Uchida. Spectroscopic imaging scanning tunneling microscopy studies of electronic structure in the superconducting and pseudogap phases of cuprate high- T_c superconductors. *Journal of the Physical Society of Japan*, 81(1):011005, 2011.
- [52] E. M. Godfrin. A method to compute the inverse of an n -block tridiagonal quasi-Hermitian matrix. *Journal of Physics: Condensed Matter*, 3(40):7843, 1991.
- [53] K. K. Gomes, A. N. Pasupathy, A. Pushp, S. Ono, Y. Ando, and A. Yazdani. Visualizing pair formation on the atomic scale in the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 447(7144):569, 2007.

Bibliography

- [54] L. P. Gorkov and P. A. Kalugin. Defects and an unusual superconductivity. *JETP Letters*, 41:253–255, 1985.
- [55] U. Gottwick, R. Held, G. Sparn, F. Steglich, H. Rietschel, D. Ewert, B. Renker, W. Bauhofer, S. Von Molnar, M. Wilhelm, and H. E. Hoenig. Transport properties of $\text{YBa}_2\text{Cu}_3\text{O}_7$: resistivity, thermal conductivity, thermopower and Hall effect. *EPL (Europhysics Letters)*, 4(10):1183, 1987.
- [56] S. S. Gubser, I. R. Klebanov, and A. M. Polyakov. Gauge theory correlators from non-critical string theory. *Physics Letters B*, 428(1-2):105–114, 1998.
- [57] E. Gull and A. J. Millis. Quasiparticle properties of the superconducting state of the two-dimensional Hubbard model. *Phys. Rev. B*, 91(8):085116, 2015.
- [58] H.-M. Guo and M. Franz. Theory of quasiparticle interference on the surface of a strong topological insulator. *Phys. Rev. B*, 81(4):041102, 2010.
- [59] M. Gurvitch and A. T. Fiory. Resistivity of $\text{La}_{1.825}\text{Sr}_{0.175}\text{CuO}_4$ and $\text{YBa}_2\text{Cu}_3\text{O}_7$ to 1100 K: Absence of saturation and its implications. *Phys. Rev. Lett.*, 59(12):1337, 1987.
- [60] F. D. M. Haldane. ‘Luttinger liquid theory’ of one-dimensional quantum fluids. I. Properties of the Luttinger model and their extension to the general 1D interacting spinless Fermi gas. *Journal of Physics C: Solid State Physics*, 14(19):2585, 1981.
- [61] T. Hanaguri, Y. Kohsaka, J. C. Davis, C. Lupien, I. Yamada, M. Azuma, M. Takano, K. Ohishi, M. Ono, and H. Takagi. Quasiparticle interference and superconducting gap in $\text{Ca}_{2-x}\text{Na}_x\text{CuO}_2\text{Cl}_2$. *Nature Physics*, 3(12):865–871, 2007.

- [62] T. Hanaguri, C. Lupien, Y. Kohsaka, D.-H. Lee, M. Azuma, M. Takano, H. Takagi, and J. C. Davis. A checkerboard electronic crystal state in lightly hole-doped $\text{Ca}_{2-x}\text{Na}_x\text{CuO}_2\text{Cl}_2$. *Nature*, 430(7003):1001–1005, 2004.
- [63] S. A. Hartnoll, A. Lucas, and S. Sachdev. *Holographic Quantum Matter*. MIT Press, 2018.
- [64] M. Hashimoto, I. M. Vishik, R.-H. He, T. P. Devereaux, and Z.-X. Shen. Energy gaps in high-transition-temperature cuprate superconductors. *Nature Physics*, 10(7):483, 2014.
- [65] H. Hertz. Ueber einen Einfluss des ultravioletten Lichtes auf die elektrische Entladung. *Annalen der Physik*, 267(8):983–1000, 1887.
- [66] R. W. Hill, C. Lupien, M. Sutherland, E. Boaknin, D. G. Hawthorn, C. Proust, F. Ronning, L. Taillefer, R. Liang, D. A. Bonn, and W. N. Hardy. Transport in Ultraclean $\text{YBa}_2\text{Cu}_3\text{O}_7$: Neither Unitary nor Born Impurity Scattering. *Phys. Rev. Lett.*, 92(2):027001, 2004.
- [67] P. J. Hirschfeld and W. A. Atkinson. Π à la Node: Disordered d-wave superconductors in two dimensions for the random masses. *Journal of Low Temperature Physics*, 126(3-4):881–900, 2002.
- [68] P. J. Hirschfeld, P. Wölfle, and D. Einzel. Consequences of resonant impurity scattering in anisotropic superconductors: Thermal and spin relaxation properties. *Phys. Rev. B*, 37(1):83, 1988.
- [69] O. Hod, J. E. Peralta, and G. E. Scuseria. First-principles electronic transport calculations in finite elongated systems: A divide and conquer approach. *The Journal of Chemical Physics*, 125(11):114704, 2006.

Bibliography

- [70] J. E. Hoffman, K. McElroy, D.-H. Lee, K. M. Lang, H. Eisaki, S. Uchida, and J. C. Davis. Imaging quasiparticle interference in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 297(5584):1148, 2002.
- [71] P. Hofmann, B. G. Briner, M. Doering, H.-P. Rust, E. W. Plummer, and A. M. Bradshaw. Anisotropic two-dimensional Friedel oscillations. *Phys. Rev. Lett.*, 79(2):265, 1997.
- [72] C. C. Homes, S. V. Dordevic, M. Strongin, D. A. Bonn, R. Liang, W. N. Hardy, S. Komiya, Y. Ando, G. Yu, N. Kaneko, X. Zhao, M. Greven, D. N. Basov, and T. Timusk. A universal scaling relation in high-temperature superconductors. *Nature*, 430(6999):539, 2004.
- [73] C. Howald, H. Eisaki, N. Kaneko, M. Greven, and A. Kapitulnik. Periodic density-of-states modulations in superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. B*, 67(1):014533, 2003.
- [74] 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):9705–9709, 2003.
- [75] C. Howald, P. Fournier, and A. Kapitulnik. Inherent inhomogeneities in tunneling spectra of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8-x}$ crystals in the superconducting state. *Phys. Rev. B*, 64(10):100504, 2001.
- [76] B. Huckestein and A. Altland. Quasi-Particle density of states of disordered d-wave superconductors, 2000.
- [77] E. W. Hudson, K. M. Lang, V. Madhavan, S. H. Pan, H. Eisaki, S. Uchida, and J. C. Davis. Interplay of magnetism and high- T_c superconductivity at individual Ni impurity atoms in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 411(6840):920–924, 2001.

- [78] N. Iqbal, H. Liu, and M. Mezei. Lectures on holographic non-Fermi liquids and quantum phase transitions. In *String Theory And Its Applications: TASI 2010 From meV to the Planck Scale*, pages 707–815. World Scientific, 2012.
- [79] N. Iqbal, H. Liu, and M. Mezei. Semi-local quantum liquids. *Journal of High Energy Physics*, 2012(4):86, 2012.
- [80] A. Junod, K.-Q. Wang, T. Tsukamoto, G. Triscone, B. Revaz, E. Walker, and J. Muller. Specific heat up to 14 tesla and magnetization of a $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_8$ single crystal thermodynamics of a 2D superconductor. *Physica C: Superconductivity*, 229(3-4):209–230, 1994.
- [81] A. Kaminski, H. M. Fretwell, M. R. Norman, M. Randeria, S. Rosenkranz, U. Chatterjee, J. C. Campuzano, J. Mesot, T. Sato, T. Takahashi, T. Terashima, M. Takano, K. Kadowaki, Z. Z. Li, and H. Raffy. Momentum anisotropy of the scattering rate in cuprate superconductors. *Phys. Rev. B*, 71(1):014517, 2005.
- [82] A. Kaminski, J. Mesot, H. Fretwell, J. C. Campuzano, M. R. Norman, M. Randeria, H. Ding, T. Sato, T. Takahashi, T. Mochiku, K. Kadowaki, and H. Hoehst. Quasiparticles in the superconducting state of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. Lett.*, 84(8):1788, 2000.
- [83] A. Kanigel, M. R. Norman, M. Randeria, U. Chatterjee, S. Souma, A. Kaminski, H. M. Fretwell, S. Rosenkranz, M. Shi, T. Sato, T. Takahashi, Z. Z. Li, H. Raffy, K. Kadowaki, D. Hinks, L. Ozyuzer, and J. C. Campuzano. Evolution of the pseudogap from Fermi arcs to the nodal liquid. *Nature Physics*, 2(7):447, 2006.

Bibliography

- [84] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen. From quantum matter to high-temperature superconductivity in copper oxides. *Nature*, 518(7538):179, 2015.
- [85] E.-A. Kim and M. J. Lawler. Interference of nematic quantum critical quasiparticles: A route to the octet model. *Phys. Rev. B*, 81(13):132501, 2010.
- [86] E.-A. Kim, M. J. Lawler, P. Oreto, S. Sachdev, E. Fradkin, and S. A. Kivelson. Theory of the nodal nematic quantum phase transition in superconductors. *Phys. Rev. B*, 77(18):184514, 2008.
- [87] S. A. Kivelson, I. P. Bindloss, E. Fradkin, V. Oganesyan, J. M. Tranquada, A. Kapitulnik, and C. Howald. How to detect fluctuating stripes in the high-temperature superconductors. *Rev. Mod. Phys.*, 75(4):1201, 2003.
- [88] S. A. Kivelson and C. M. Varma. Fermi pockets in a d-wave superconductor with coexisting loop-current order. *arXiv preprint arXiv:1208.6498*, 2012.
- [89] Y. Kohsaka, C. Taylor, K. Fujita, A. 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):1380–1385, 2007.
- [90] Y. Kohsaka, C. Taylor, P. Wahl, A. Schmidt, J. Lee, K. Fujita, J. W. Alldredge, K. McElroy, J. 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, 2008.

- [91] J. D. Koralek, J. F. Douglas, N. C. Plumb, Z. Sun, A. V. Federov, M. M. Murnane, H. C. Kapteyn, S. T. Cundiff, Y. Aiura, K. Oka, H. Eisaki, and D. S. Dessau. Laser based angle-resolved photoemission, the sudden approximation, and quasiparticle-like spectral peaks in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. Lett.*, 96(1):017005, 2006.
- [92] B. Kramer and A. MacKinnon. Localization: theory and experiment. *Reports on Progress in Physics*, 56(12):1469, 1993.
- [93] A. Kreisel, P. Choubey, T. Berlijn, W. Ku, B. M. Andersen, and P. J. Hirschfeld. Interpretation of scanning tunneling quasiparticle interference and impurity states in cuprates. *Phys. Rev. Lett.*, 114(21):217002, 2015.
- [94] A. Kuzmin, M. Luisier, and O. Schenk. Fast methods for computing selected elements of the Greens function in massively parallel nanoelectronic device simulations. In *European Conference on Parallel Processing*, pages 533–544. Springer, 2013.
- [95] K. M. Lang, V. Madhavan, J. E. Hoffman, E. W. Hudson, H. Eisaki, S. Uchida, and J. C. Davis. Imaging the granular nature of high- T_c superconductivity in underdoped $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 415(6870):412–6, 2002.
- [96] M. J. Lawler, K. Fujita, J. Lee, A. R. Schmidt, Y. Kohsaka, C. K. Kim, H. Eisaki, S. Uchida, J. C. Davis, J. P. Sethna, and E.-A. Kim. Intra-unit-cell electronic nematicity of the high- T_c copper-oxide pseudogap states. *Nature*, 466(7304):347–351, 2010.
- [97] S. Lederer, Y. Schattner, E. Berg, and S. A. Kivelson. Superconductivity and non-Fermi liquid behavior near a nematic quan-

Bibliography

- tum critical point. *Proceedings of the National Academy of Sciences*, 114(19):4905–4910, 2017.
- [98] J. Lee, M. P. Allan, M. A. Wang, J. Farrell, S. A. Grigera, F. Baumberger, J. C. Davis, and A. P. Mackenzie. Heavy d-electron quasiparticle interference and real-space electronic structure of $\text{Sr}_3\text{Ru}_2\text{O}_7$. *Nature Physics*, 5(11):800–804, 2009.
- [99] J. Lee, K. Fujita, A. R. Schmidt, C. K. Kim, H. Eisaki, S. Uchida, and J. C. Davis. Spectroscopic fingerprint of phase-incoherent superconductivity in the underdoped $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 325(5944):1099–1103, 2009.
- [100] P. A. Lee. Localized states in a d-wave superconductor. *Phys. Rev. Lett.*, 71(12):1887, 1993.
- [101] S.-S. Lee. Low-energy effective theory of Fermi surface coupled with $U(1)$ gauge field in $2+1$ dimensions. *Phys. Rev. B*, 80(16):165102, 2009.
- [102] W. S. Lee, I. M. Vishik, K. Tanaka, D. H. Lu, T. Sasagawa, N. Nagaosa, T. P. Devereaux, Z. Hussain, and Z.-X. Shen. Abrupt onset of a second energy gap at the superconducting transition of underdoped Bi_2Tl_2 . *Nature*, 450(7166):81, 2007.
- [103] S. Li and E. Darve. Extension and optimization of the FIND algorithm: Computing Greens and less-than Greens functions. *Journal of Computational Physics*, 231(4):1121–1139, 2012.
- [104] S. Li, W. Wu, and E. Darve. A fast algorithm for sparse matrix computations related to inversion. *Journal of Computational Physics*, 242:915–945, 2013.

- [105] J. M. Luttinger. An Exactly Soluble Model of a Many-Fermion System. *Journal of Mathematical Physics*, 4:1154–1162, 1963.
- [106] A. MacKinnon and B. Kramer. One-parameter scaling of localization length and conductance in disordered systems. *Phys. Rev. Lett.*, 47(21):1546, 1981.
- [107] G. D. Mahan. *Many-Particle Physics*. Springer Science & Business Media, 2013.
- [108] J. Maldacena. The Large-N Limit of Superconformal Field Theories and Supergravity. *International Journal of Theoretical Physics*, 38(4):1113–1133, 1999.
- [109] I. Martin, A. V. Balatsky, and J. Zaanen. Impurity states and interlayer tunneling in high temperature superconductors. *Phys. Rev. Lett.*, 88(9):097003, 2002.
- [110] D. C. Mattis and E. H. Lieb. Exact Solution of a Many-Fermion System and Its Associated Boson Field. *Journal of Mathematical Physics*, 6:304–312, 1965.
- [111] K. McElroy, J. Lee, J. 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):1048–1052, 2005.
- [112] K. McElroy, R. W. Simmonds, J. E. Hoffman, D.-H. Lee, J. Orenstein, H. Eisaki, S. Uchida, and J. C. Davis. Relating atomic-scale electronic phenomena to wave-like quasiparticle states in superconducting $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 422(6932):592–596, 2003.

Bibliography

- [113] M. A. Metlitski and S. Sachdev. Quantum phase transitions of metals in two spatial dimensions. I. Ising-nematic order. *Phys. Rev. B*, 82(7):075127, 2010.
- [114] S. Misra, M. Vershinin, P. Phillips, and A. Yazdani. Failure of scattering interference in the pseudogap state of cuprate superconductors. *Phys. Rev. B*, 70(22):220503, 2004.
- [115] K. A. Moler, D. J. Baar, J. S. Urbach, R. Liang, W. N. Hardy, and A. Kapitulnik. Magnetic Field Dependence of the Density of States of $\text{YBa}_2\text{Cu}_3\text{O}_{6.95}$ as Determined from the Specific Heat. *Phys. Rev. Lett.*, 73(20):2744, 1994.
- [116] K. A. Moler, D. L. Sisson, J. S. Urbach, M. R. Beasley, A. Kapitulnik, D. J. Baar, R. Liang, and W. N. Hardy. Specific heat of $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$. *Phys. Rev. B*, 55(6):3954, 1997.
- [117] R. Mukhopadhyay, C. L. Kane, and T. C. Lubensky. Sliding Luttinger liquid phases. *Phys. Rev. B*, 64(4):045120, 2001.
- [118] A. A. Nersisyan, A. M. Tsvelik, and F. Wenger. Disorder effects in two-dimensional d-wave superconductors. *Phys. Rev. Lett.*, 72(16):2628, 1994.
- [119] M. R. Norman, H. Ding, M. Randeria, J. C. Campuzano, T. Yokoya, T. Takeuchi, T. Takahashi, T. Mochiku, K. Kadowaki, P. Guptasarma, and D. G. Hinks. Destruction of the Fermi surface in underdoped high- T_c superconductors. *Nature*, 392(6672):157, 1998.
- [120] M. R. Norman, A. Kanigel, M. Randeria, U. Chatterjee, and J. C. Campuzano. Modeling the Fermi arc in underdoped cuprates. *Phys. Rev. B*, 76(17):174501, 2007.

- [121] M. R. Norman, D. Pines, and C. Kallin. The pseudogap: friend or foe of high T_c ? *Advances in Physics*, 54(8):715–733, 2005.
- [122] M. R. Norman, M. Randeria, H. Ding, and J. C. Campuzano. Phenomenological models for the gap anisotropy of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_8$ as measured by angle-resolved photoemission spectroscopy. *Phys. Rev. B*, 52(1):615, 1995.
- [123] M. R. Norman, M. Randeria, H. Ding, and J. C. Campuzano. Phenomenology of the low-energy spectral function in high- T_c superconductors. *Phys. Rev. B*, 57(18):R11093, 1998.
- [124] T. S. Nunner, B. M. Andersen, A. Melikyan, and P. J. Hirschfeld. Dopant-modulated pair interaction in cuprate superconductors. *Phys. Rev. Lett.*, 95(17):177003, 2005.
- [125] T. S. Nunner, W. Chen, B. M. Andersen, A. Melikyan, and P. J. Hirschfeld. Fourier transform spectroscopy of d-wave quasiparticles in the presence of atomic scale pairing disorder. *Phys. Rev. B*, 73(10):104511, 2006.
- [126] T. S. Nunner and P. J. Hirschfeld. Microwave conductivity of d-wave superconductors with extended impurities. *Phys. Rev. B*, 72(1):014514, 2005.
- [127] D. Orgad, S. A. Kivelson, E. W. Carlson, V. J. Emery, X. J. Zhou, and Z. X. Shen. Evidence of electron fractionalization from photoemission spectra in the high temperature superconductors. *Phys. Rev. Lett.*, 86(19):4362, 2001.
- [128] S. Ouazi, J. Bobroff, H. Alloul, M. Le Tacon, N. Blanchard, G. Collin, M. H. Julien, M. Horvatić, and C. Berthier. Impurity-induced local magnetism and density of states in the superconducting state of $\text{YBa}_2\text{Cu}_3\text{O}_7$. *Phys. Rev. Lett.*, 96(12):127005, 2006.

Bibliography

- [129] S. H. Pan, E. W. Hudson, K. M. Lang, H. Eisaki, S. Uchida, and J. C. Davis. Imaging the effects of individual zinc impurity atoms on superconductivity in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 403(6771):746–750, 2000.
- [130] S. H. Pan, J. P. O’Neal, R. L. Badzey, C. Chamon, H. Ding, J. R. Engelbrecht, Z. Wang, H. Eisaki, S. Uchida, A. K. Gupta, K.-W. Ng, E. W. Hudson, K. M. Lang, and J. C. Davis. Microscopic electronic inhomogeneity in the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+x}$. *Nature*, 413(6853):282–285, 2001.
- [131] C. V. Parker, P. Aynajian, E. H. da Silva Neto, A. Pushp, S. Ono, J. Wen, Z. Xu, G. Gu, and A. Yazdani. Fluctuating stripes at the onset of the pseudogap in the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature*, 468(7324):677, 2010.
- [132] A. N. Pasupathy, A. Pushp, K. K. Gomes, C. V. Parker, J. Wen, Z. Xu, G. Gu, S. Ono, Y. Ando, and A. Yazdani. Electronic origin of the inhomogeneous pairing interaction in the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 320(5873):196–201, 2008.
- [133] T. Pereg-Barnea and M. Franz. Theory of quasiparticle interference patterns in the pseudogap phase of the cuprate superconductors. *Phys. Rev. B*, 68(18):180506, 2003.
- [134] T. Pereg-Barnea and M. Franz. Quasiparticle interference patterns as a test for the nature of the pseudogap phase in the cuprate superconductors. *International Journal of Modern Physics B*, 19(04):731–761, 2005.
- [135] T. Pereg-Barnea and M. Franz. Magnetic-field dependence of quasiparticle interference peaks in a d-wave superconductor with weak disorder. *Phys. Rev. B*, 78(2):020509, 2008.

- [136] D. E. Petersen, H. H. B. Sørensen, P. C. Hansen, S. Skelboe, and K. Stokbro. Block tridiagonal matrix inversion and fast transmission calculations. *Journal of Computational Physics*, 227(6):3174–3190, 2008.
- [137] L. Petersen, P. T. Sprunger, P. Hofmann, E. Lægsgaard, B. G. Briner, M. Doering, H.-P. Rust, A. M. Bradshaw, F. Besenbacher, and E. W. Plummer. Direct imaging of the two-dimensional Fermi contour: Fourier-transform STM. *Phys. Rev. B*, 57(12):R6858, 1998.
- [138] T. J. Reber, S. Parham, N. C. Plumb, Y. Cao, H. Li, Z. Sun, Q. Wang, H. Iwasawa, M. Arita, J. S. Wen, Z. J. Xu, G. D. Gu, Y. Yoshida, H. Eisaki, G. B. Arnold, and D. S. Dessau. Pairing, pair-breaking, and their roles in setting the T_c of cuprate high temperature superconductors. *arXiv preprint arXiv:1508.06252*, 2015.
- [139] T. J. Reber, S. Parham, N. C. Plumb, Y. Cao, H. Li, Z. Sun, Q. Wang, H. Iwasawa, J. S. Wen, Z. J. Xu, G. Gu, S. Ono, H. Berger, Y. Yoshida, H. Eisaki, Y. Aiura, G. B. Arnold, and D. S. Dessau. Coordination of the energy and temperature scales of pairing across the doping phase diagram of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *arXiv preprint arXiv:1509.01556*, 2015.
- [140] T. J. Reber, N. C. Plumb, Y. Cao, Z. Sun, Q. Wang, K. McElroy, H. Iwasawa, M. Arita, J. S. Wen, Z. J. Xu, G. Gu, Y. Yoshida, H. Eisaki, Y. Aiura, and D. S. Dessau. Preparing and the filling gap in the cuprates from the tomographic density of states. *Phys. Rev. B*, 87(6):060506, 2013.
- [141] T. J. Reber, N. C. Plumb, Z. Sun, Y. Cao, Q. Wang, K. McElroy, H. Iwasawa, M. Arita, J. S. Wen, Z. J. Xu, G. Gu, Y. Yoshida,

Bibliography

- H. Eisaki, Y. Aiura, and D. S. Dessau. The origin and non-quasiparticle nature of Fermi arcs in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature Physics*, 8(8):606, 2012.
- [142] T. J. Reber, X. Zhou, N. C. Plumb, S. Parham, J. A. Waugh, Y. Cao, Z. Sun, H. Li, Q. Wang, J. S. Wen, Z. J. Xu, G. Gu, Y. Yoshida, H. Eisaki, G. B. Arnold, and D. S. Dessau. Power Law Liquid-A Unified Form of Low-Energy Nodal Electronic Interactions in Hole Doped Cuprate Superconductors. *arXiv preprint arXiv:1509.01611*, 2015.
- [143] M. G. Reuter and J. C. Hill. An efficient, block-by-block algorithm for inverting a block tridiagonal, nearly block Toeplitz matrix. *Computational Science & Discovery*, 5(1):014009, 2012.
- [144] S. C. Riggs, O. Vafek, J. B. Kemper, J. B. Betts, A. Migliori, F. F. Balakirev, W. N. Hardy, R. Liang, D. A. Bonn, and G. S. Boebinger. Heat capacity through the magnetic-field-induced resistive transition in an underdoped high-temperature superconductor. *Nature Physics*, 7(4):332–335, 2011.
- [145] P. Roushan, J. Seo, C. V. Parker, Y. S. Hor, D. Hsieh, D. Qian, A. Richardella, M. Z. Hasan, R. J. Cava, and A. Yazdani. Topological surface states protected from backscattering by chiral spin texture. *Nature*, 460(7259):1106–1109, 2009.
- [146] S. Sachdev. *Quantum Phase Transitions*. Wiley Online Library, 2007.
- [147] S. Sakai, M. Civelli, and M. Imada. Hidden-fermion representation of self-energy in pseudogap and superconducting states of the two-dimensional Hubbard model. *Phys. Rev. B*, 94(11):115130, 2016.

- [148] D. J. Scalapino, S. R. White, and S. Zhang. Insulator, metal, or superconductor: The criteria. *Phys. Rev. B*, 47(13):7995, 1993.
- [149] D. J. Scalapino, S. R. White, and S. C. Zhang. Superfluid density and the Drude weight of the Hubbard model. *Phys. Rev. Lett.*, 68(18):2830, 1992.
- [150] A. R. Schmidt, K. Fujita, E.-A. Kim, M. J. Lawler, H. Eisaki, S. Uchida, D.-H. Lee, and J. C. Davis. Electronic structure of the cuprate superconducting and pseudogap phases from spectroscopic imaging STM. *New Journal of Physics*, 13(6):065014, 2011.
- [151] A. R. Schmidt, M. H. Hamidian, P. Wahl, F. Meier, A. V. Balatsky, J. D. Garrett, T. J. Williams, G. M. Luke, and J. C. Davis. Imaging the Fano lattice to hidden ordertransition in URu₂Si₂. *Nature*, 465(7298):570–576, 2010.
- [152] G. Schubert, J. Schleede, K. Byczuk, H. Fehske, and D. Vollhardt. Distribution of the local density of states as a criterion for Anderson localization: Numerically exact results for various lattices in two and three dimensions. *Phys. Rev. B*, 81(15):155106, 2010.
- [153] S. E. Sebastian, N. Harrison, F. F. Balakirev, M. M. Altarawneh, P. A. Goddard, R. Liang, D. A. Bonn, W. N. Hardy, and G. G. Lonzarich. Normal-state nodal electronic structure in underdoped high- T_c copper oxides. *Nature*, 511(7507):61, 2014.
- [154] S. E. Sebastian, N. Harrison, and G. G. Lonzarich. Towards resolution of the Fermi surface in underdoped high- T_c superconductors. *Reports on Progress in Physics*, 75(10):102501, 2012.
- [155] T. Senthil and M. P. A. Fisher. Quasiparticle density of states in dirty high- T_c superconductors. *Phys. Rev. B*, 60(9):6893, 1999.

Bibliography

- [156] T. Senthil, M. P. A. Fisher, L. Balents, and C. Nayak. Quasiparticle transport and localization in high- T_c superconductors. *Phys. Rev. Lett.*, 81(21):4704, 1998.
- [157] Z.-X. Shen, D. S. Dessau, B. O. Wells, D. M. King, W. E. Spicer, A. J. Arko, D. Marshall, L. W. Lombardo, A. Kapitulnik, P. Dickinson, S. Doniach, J. DiCarlo, T. Loeser, and C. H. Park. Anomalous large gap anisotropy in the a-b plane of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. Lett.*, 70(10):1553, 1993.
- [158] S. L. Sondhi, S. M. Girvin, J. P. Carini, and D. Shahar. Continuous quantum phase transitions. *Rev. Mod. Phys.*, 69(1):315, 1997.
- [159] P. T. Sprunger, L. Petersen, E. W. Plummer, E. Lægsgaard, and F. Besenbacher. Giant Friedel oscillations on the beryllium (0001) surface. *Science*, 275(5307):1764–1767, 1997.
- [160] S. Sugai. Phonon Raman scattering in $(\text{La}_{1-x}\text{Sr}_x)_2\text{CuO}_4$ single crystals. *Phys. Rev. B*, 39(7):4306, 1989.
- [161] M. A. Sulangi, M. P. Allan, and J. Zaanen. Revisiting quasiparticle scattering interference in high-temperature superconductors: The problem of narrow peaks. *Phys. Rev. B*, 96(13):134507, 2017.
- [162] M. A. Sulangi and J. Zaanen. Quasiparticle density of states, localization, and distributed disorder in the cuprate superconductors. *Phys. Rev. B*, 97(14):144512, 2018.
- [163] K. Tanaka, W. S. Lee, D. H. Lu, A. Fujimori, T. Fujii, I. Terasaki, D. J. Scalapino, T. P. Devereaux, Z. Hussain, and Z.-X. Shen. Distinct Fermi-momentum-dependent energy gaps in deeply underdoped Bi_2Tl_2 . *Science*, 314(5807):1910–1913, 2006.

- [164] T. Timusk and B. Statt. The pseudogap in high-temperature superconductors: an experimental survey. *Reports on Progress in Physics*, 62(1):61, 1999.
- [165] S. Tomonaga. Remarks on Bloch's Method of Sound Waves applied to Many-Fermion Problems. *Progress of Theoretical Physics*, 5:544–569, 1950.
- [166] Y. J. Uemura, G. M. Luke, B. J. Sternlieb, J. H. Brewer, J. F. Carolan, W. N. Hardy, R. Kadono, J. R. Kempton, R. F. Kiefl, S. R. Kreitzman, P. Mulhern, T. M. Riseman, D. L. Williams, B. X. Yang, S. Uchida, H. Takagi, J. Gopalakrishnan, A. W. Sleight, M. A. Subramanian, C. L. Chien, M. Z. Cieplak, G. Xiao, V. Y. Lee, B. W. Statt, C. E. Stronach, W. J. Kossler, and X. H. Yu. Universal Correlations between T_c and $\frac{n_s}{m^*}$ (Carrier Density over Effective Mass) in High- T_c Cuprate Superconductors. *Phys. Rev. Lett.*, 62(19):2317, 1989.
- [167] J. S. Urbach, D. B. Mitzi, A. Kapitulnik, J. Y. T. Wei, and D. E. Morris. Low-temperature specific heat of single-crystal $\text{Bi}_2\text{CaSr}_2\text{Cu}_2\text{O}_8$ and $\text{Tl}_2\text{Ca}_2\text{Ba}_2\text{Cu}_3\text{O}_{10}$. *Phys. Rev. B*, 39(16):12391, 1989.
- [168] T. Valla, A. V. Fedorov, P. D. Johnson, Q. Li, G. D. Gu, and N. Koshizuka. Temperature dependent scattering rates at the Fermi surface of optimally doped $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. Lett.*, 85(4):828, 2000.
- [169] T. Valla, A. V. Fedorov, P. D. Johnson, B. O. Wells, S. L. Hulbert, Q. Li, G. D. Gu, and N. Koshizuka. Evidence for quantum critical behavior in the optimally doped cuprate $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 285(5436):2110–2113, 1999.

Bibliography

- [170] T. Valla, T. E. Kidd, J. D. Rameau, H.-J. Noh, G. D. Gu, P. D. Johnson, H.-B. Yang, and H. Ding. Fine details of the nodal electronic excitations in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Phys. Rev. B*, 73(18):184518, 2006.
- [171] J. S. Van Dyke, F. Masee, M. P. Allan, J. C. S. Davis, C. Petrovic, and D. K. Morr. Direct evidence for a magnetic f-electron-mediated pairing mechanism of heavy-fermion superconductivity in CeCoIn_5 . *Proceedings of the National Academy of Sciences*, 111(32):11663–11667, 2014.
- [172] C. M. Varma and E. Abrahams. Effective Lorentz force due to small-angle impurity scattering: magnetotransport in high- T_c superconductors. *Phys. Rev. Lett.*, 86(20):4652, 2001.
- [173] C. M. Varma, P. B. Littlewood, S. Schmitt-Rink, E. Abrahams, and A. E. Ruckenstein. Phenomenology of the normal state of Cu-O high-temperature superconductors. *Phys. Rev. Lett.*, 63(18):1996, 1989.
- [174] M. Vershinin, S. Misra, S. Ono, Y. Abe, Y. Ando, and A. Yazdani. Local ordering in the pseudogap state of the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Science*, 303(5666):1995–1998, 2004.
- [175] I. Vishik. Photoemission perspective on pseudogap, superconducting fluctuations, and charge order: a review of recent progress. *Reports on Progress in Physics*, 2018.
- [176] I. Vishik, E. Nowadnick, W. Lee, Z. Shen, B. Moritz, T. Devereaux, K. Tanaka, T. Sasagawa, and T. Fujii. A momentum-dependent perspective on quasiparticle interference in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature Physics*, 5(10):718, 2009.
- [177] I. M. Vishik, M. Hashimoto, R.-H. He, W.-S. Lee, F. Schmitt, D. Lu, R. G. Moore, C. Zhang, W. Meevasana, T. Sasagawa,

- S. Uchida, K. Fujita, S. Ishida, M. Ishikado, Y. Yoshida, H. Eisaki, Z. Hussain, T. P. Devereaux, and Z.-X. Shen. Phase competition in trisected superconducting dome. *Proceedings of the National Academy of Sciences*, 109(45):18332–18337, 2012.
- [178] I. M. Vishik, W. S. Lee, F. Schmitt, B. Moritz, T. Sasagawa, S. Uchida, K. Fujita, S. Ishida, C. Zhang, T. P. Devereaux, and Z.-X. Shen. Doping-dependent nodal Fermi velocity of the high-temperature superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ revealed using high-resolution angle-resolved photoemission spectroscopy. *Phys. Rev. Lett.*, 104(20):207002, 2010.
- [179] A. Vishwanath and D. Carpentier. Two-dimensional anisotropic non-Fermi-liquid phase of coupled Luttinger liquids. *Phys. Rev. Lett.*, 86(4):676, 2001.
- [180] R. E. Walstedt and W. W. Warren. Nuclear resonance properties of $\text{YBa}_2\text{Cu}_3\text{O}_{6+x}$ superconductors. *Science*, 248(4959):1082–1087, 1990.
- [181] L. Wang and O. Vafek. Quantum oscillations of the specific heat in d-wave superconductors with loop current order. *Phys. Rev. B*, 88(2):024506, 2013.
- [182] Q.-H. Wang and D.-H. Lee. Quasiparticle scattering interference in high-temperature superconductors. *Phys. Rev. B*, 67(2):020511, 2003.
- [183] W. W. Warren Jr, R. Walstedt, G. Brennert, R. Cava, R. Tycko, R. Bell, and G. Dabbagh. Cu spin dynamics and superconducting precursor effects in planes above T_c in $\text{YBa}_2\text{Cu}_3\text{O}_{6.7}$. *Phys. Rev. Lett.*, 62(10):1193, 1989.

Bibliography

- [184] M. Wimmer and K. Richter. Optimal block-tridiagonalization of matrices for coherent charge transport. *Journal of Computational Physics*, 228(23):8548–8565, 2009.
- [185] E. Witten. Anti-de Sitter space and holography. *Advances in Theoretical and Mathematical Physics*, 2:253–291, 1998.
- [186] T. Xiang. Effect of superconducting correlation on the localization of quasiparticles in low dimensions. *Phys. Rev. B*, 52(9):6204, 1995.
- [187] A. G. Yashenkin, W. A. Atkinson, I. V. Gornyi, P. J. Hirschfeld, and D. V. Khveshchenko. Nesting symmetries and diffusion in disordered d-wave superconductors. *Phys. Rev. Lett.*, 86(26):5982, 2001.
- [188] J. Zaanen. High-temperature superconductivity: Quantum salad dressing. *Nature*, 415(6870):379, 2002.
- [189] J. Zaanen. Superconductivity: Pebbles in the nodal pond. *Nature*, 422(6932):569–570, 2003.
- [190] J. Zaanen. Superconductivity: Why the temperature is high. *Nature*, 430(6999):512, 2004.
- [191] J. Zaanen, Y. Liu, Y.-W. Sun, and K. Schalm. *Holographic Duality in Condensed Matter Physics*. Cambridge University Press, 2015.
- [192] R. Zhou, M. Hirata, T. Wu, I. Vinograd, H. Mayaffre, S. Krämer, M. Horvatić, C. Berthier, A. P. Reyes, P. L. Kuhns, R. Liang, W. N. Hardy, D. A. Bonn, and M.-H. Julien. Quasiparticle scattering off defects and possible bound states in charge-ordered $\text{YBa}_2\text{Cu}_3\text{O}_y$. *Phys. Rev. Lett.*, 118(1):017001, 2017.

- [193] R. Zhou, M. Hirata, T. Wu, I. Vinograd, H. Mayaffre, S. Krämer, A. P. Reyes, P. L. Kuhns, R. Liang, W. N. Hardy, D. A. Bonn, and M.-H. Julien. Spin susceptibility of charge-ordered $\text{YBa}_2\text{Cu}_3\text{O}_y$ across the upper critical field. *Proceedings of the National Academy of Sciences*, page 201711445, 2017.
- [194] J.-X. Zhu, D. N. Sheng, and C. S. Ting. Quasiparticle localization in disordered d-wave superconductors. *Phys. Rev. Lett.*, 85(23):4944, 2000.
- [195] L. Zhu, W. A. Atkinson, and P. J. Hirschfeld. Two impurities in a d-wave superconductor: Local density of states. *Phys. Rev. B*, 67(9):094508, 2003.
- [196] L. Zhu, W. A. Atkinson, and P. J. Hirschfeld. Power spectrum of many impurities in a d-wave superconductor. *Phys. Rev. B*, 69(6):060503, 2004.

SAMENVATTING

Dit proefschrift is gewijd aan een gedetailleerd onderzoek naar de verschillende effecten van wanorde in hoge-temperatuur cupraat supergeleiders. Wanorde is alomtegenwoordig in deze materialen en speelt een centrale rol in een aantal verschijnselen die geobserveerd worden in verschillende fasen. Een voorbeeld hiervan is de verstrooiingsinterferentie van quasideeltjes—modulaties van de lokale toestandsdichtheid door de verstrooiing van quantummechanische golven door onzuiverheden—die gebruikt is om de structuur in de impulsruimte van elektronische excitaties in deze materialen in kaart te brengen. Een ander voorbeeld is de waarneming van een eindige toestandsdichtheid bij de Fermi-energie in de supergeleidende fase, dat typisch verklaard wordt door de aanwezigheid van wanorde. Uit deze twee voorbeelden wordt de duale rol van wanorde in deze materialen zichtbaar: als de wanorde zwak is tast het de elektronische excitaties af, maar onder bepaalde omstandigheden kan het ook nieuwe laagenergetische excitaties genereren die in een “schoon” materiaal afwezig zijn.

In dit proefschrift worden enkele verschijnselen die voorkomen in cupraten onder de loep genomen, vanuit het perspectief van wat nu bekend is over de aard van de wanorde in deze materialen. Van de metallische koperoxidevlakken is bekend dat zij geen wanorde bevatten, wat impliceert dat veel van de aanwezige wanorde veroorzaakt wordt door doping defecten die zich in de isolerende bufferlagen bevinden, op afstand van de metallische vlakken. Deze defecten zorgen voor een gedistribueerde Coulombische wanorde, een relatief gladde poten-

tiaal die gevoeld wordt door de elektronen in de metallische vlakken. Deze details zijn cruciaal, omdat in veel van de literatuur aangenomen wordt dat de wanorde zich *in* de metallische vlakken bevindt. Veel van de bestaande theorie over de verstrooiingsinterferentie van quasi-deeltjes draaide om modellen van wanorde met slechts één defect. Op een soortgelijke manier wordt er in veel modellen aangenomen dat de wanorde zich in de koperoxidevlakken bevindt om de eindige toestandsdichtheid bij de Fermi-energie te verklaren. In hoeverre de door wanorde aangedreven verschijnselen in experimenten door realistische, gedistribueerde vormen van wanorde kunnen worden verklaard is nauwelijks bestudeert.

In hoofdstuk 3 richten wij ons op verstrooiingsinterferentie van quasi-deeltjes in cupraten en beschouwen wij hoe de scherpe pieken in de tunnelingspectra van experimenten gereproduceerd kunnen worden door een reeks modellen van gedistribueerde wanorde, inclusief de modellering voor het tunnelproces tussen de koperoxidevlakken en de tip van de scanning tunneling microscoop. Wij concluderen dat zwakke puntvormige wanorde en willekeurige wanorde in de chemische potentiaal het best de experimentele spectra nabootst in zowel de reële ruimte als in de impulsruimte. Tegelijkertijd is gladde wanorde niet afdoende om de experimentele spectra te reproduceren omdat verstrooiing bij groot momentum onderdrukt wordt in dit geval. Even interessant als paradoxaal is het feit dat de pieken geobserveerd in experimenten *scherper* zijn dan de pieken die volgen uit onze simulaties.

Hoofdstuk 4 beschouwt verschillende vormen van wanorde en hun uitwerking op de toestandsdichtheid bij de Fermi-energie met variërende hoeveelheden wanorde. Hoewel sterke puntvormige onzuiverheden en wanorde in de chemische potentiaal wel tot een eindige toestandsdichtheid leiden, moet er onrealistisch veel wanorde aangenomen worden om het experiment te reproduceren. Anderzijds laten we zien dat

wanorde door defecten buiten het vlak tot realistische waarden van de toestandsdichtheid bij de Fermi-energie kan leiden zonder de supergeleidende excitaties bij middelmatige en hogere energieën aan te tasten. Wij vinden scherpe resonanties *bij* de Fermi-energie wanneer er zeer veel dotanten buiten het vlak aanwezig zijn. Wij onderzoeken ook de lokalisatielengte voor verschillende modellen van wanorde, en het wordt aangetoond dat voor alle modellen van wanorde de quasi-deeltjes bij de Fermi-energie gelokaliseerd zijn, terwijl de lokalisatielengte als functie van energie sterk gevoelig is voor het type wanorde dat aanwezig is.

Hoofdstuk 5 beschouwt wederom verstrooiingsinterferentie van quasi-deeltjes, dit keer met het oog op de effecten van zelf-energieën op de lokale toestandsdichtheid, wat toegepast wordt op relevante voorbeelden in de cupraten. In de supergeleidende toestand bestuderen we het fenomeen van *kloofvullen* versus *-dichten* dat in ARPES experimenten wordt waargenomen en wij analyseren in hoeverre scanning tunneling spectroscopie experimenten dit fenomeen ook waar kunnen nemen. De pieken in de supergeleidende spectra worden aangetoond gevoelig te zijn voor de hoeveelheid verbreding die aanwezig is: de pieken worden uitgesmeerd en worden incoherent bij grote zelf-energieën. We besturen ook de spectra in de normale toestand, waar wij aannemen dat het “vreemde metaal” goed beschreven wordt door een marginale Fermivloeistof. Wij vinden dat een belangrijk kenmerk van deze toestand de aanwezigheid van brede patronen is in het spectrum dat de verstrooiing tussen punten op het Fermi-oppervlak beschrijft. Het grootste verschil tussen de marginale Fermivloeistof en de gewone Fermivloeistof blijkt de mate van verbreding te zijn; het spectrum van een marginale Fermivloeistof heeft veel meer verbreding dan dat van een gewone Fermivloeistof.

De meest belangrijkste resultaten van dit proefschrift zijn als volgt: de gedistribueerde vormen van wanorde kan de verstrooiingsinterferentie van quasideeltjes die gezien wordt in experimenten veroorzaken, niettegenstaande de afwezigheid van duidelijk geïsoleerde onzuiverheden; de microscopische details van het tunnelproces geassocieerd met scanning tunneling spectroscopie experimenten zijn van groot belang voor de theoretische interpretatie van experimenten; wanorde buiten de vlakken is een kandidaat mechanisme voor de productie van laagenergetische excitaties die in soortelijke warmte experimenten worden waargenomen; de gevoeligheid van de verstrooiingsinterferentie van quasideeltjes in de supergeleidende toestand voor de lange levensduur van Bogoliubov quasideeltjes; en tenslotte het verrassend subtiele verschil tussen de spectra van een marginale Fermivloeistof en die van zijn conventionele broertje, de Fermivloeistof. In hoofdstuk 6 brengen wij al deze resultaten samen en plaatsen wij deze in de grotere context van wat er over de cupraten bekend is.

ACKNOWLEDGEMENTS

Here I would like to thank everyone I've encountered who in one way or another has enabled me to enjoy life in the past several years as a theoretical physicist.

First, I thank Jan Zaanen for being a wonderful supervisor, and for being singularly responsible for my worldview as a theoretical physicist. In the years I've spent working with him I learned much about condensed matter physics from him, and his endless reserves of knowledge about the vast universe of theories and experiments on the cuprates (among other materials) helped convince me that the down-and-dirty world of high- T_c superconductivity and strongly correlated electrons was where the most interesting questions in physics were. From Jan I also learned to be intellectually fearless and independent—he impressed upon me the need to find my own direction as a scientist and not to simply follow the herd, advice that I've adopted as a life mantra of sorts.

Milan Allan has been a wonderful collaborator and sounding board for many of my ideas, and has been remarkably patient with me in going over some real-life subtleties of experiments that would likely have flown over the heads of theorists like me. I can't thank him enough for giving me a dose of "reality"—nothing really matches the thrill of doing theory that one can actually see in experiments close by!

A considerable part of what I know about condensed matter physics and quantum field theory was a direct result of my interactions, whether in-class or informal, with a terrific set of physicists I've been very lucky

ACKNOWLEDGEMENTS

to have known in my first few years as a graduate student—Peter Armitage, Hans Hilgenkamp, Predrag Nikolic, Tjerk Oosterkamp, Koenraad Schalm, and Oleg Tchernyshyov—and for that I am truly grateful. Discussions with Lewis Bawden, Seamus Davis, Mark Golden, Marc-Henri Julien, and Steven Kivelson proved invaluable in shaping the content of the papers forming this thesis.

I wouldn't have made it this far if it weren't for two physicists in particular who have been profoundly formative in my early training as a physicist and who have, unfortunately, passed on. Germelino Abito introduced me to condensed matter theory early in my undergrad years in Ateneo, and my experience working with him would later prove to be truly invaluable as I moved up to graduate school—I can't thank him enough for affording me a considerable amount of intellectual freedom in my research, which no doubt readied me for the rigors of life yet to come. Zlatko Tesanovic taught me graduate quantum mechanics and took me in as a Ph.D. student in his group at Johns Hopkins, a stint that ended with his sudden and tragic passing; from my interactions with him as a beginning graduate student I became acquainted with the very high standards of theoretical physics he operated by. It makes for an interesting coincidence that this thesis revolves around topics they specialized in (disorder for Germelino, high- T_c superconductivity for Zlatko), and I hope that they would've at the very least found what I wrote here interesting.

I am grateful to the secretaries at the Instituut-Lorentz for all they've done for me and my colleagues. To Marianne van Dun, Manon van Ette, Trudy Geurds, and Fran Ouwerkerk, thank you for all your help, advice, conversations, and good humor, all of which were instrumental in making my life as a Ph.D. student easier and more productive.

I owe a lot to my friends who've made graduate school an enjoyable experience. To all the members of Jan and Koenraad's group I over-

lapped with—Andrey Bagrov, Floris Balm, Bartosz Benenowski, Josko de Boer, Richard Davison, Philippe Sabella-Garnier, Mikhail Goykhman, Saso Grozdanov, Olfa Jaibi, Nikolaos Kaplis, Alexander Krikun, Ke Liu, Balazs Meszena, Jaakko Nissinen, Nick Poovuttikul, Louk Rademaker, Aurelio Romero Bermudez, Petter Saterskog, Vincenzo Scopelitti, and Robert-Jan Slager—thank you for all the wide-ranging discussions and the countless dinners and drinks we’ve had together that have helped make academic life less dreary. I owe a lot in particular to Robert-Jan for being a fantastically helpful academic sibling, to Louk for guiding me upon arriving in Leiden, to Jaakko for sharing his insights and perspectives on matters both physics- and non-physics-related, to Floris and Josko for help in translating part of this thesis into Dutch, and, last but not least, to Nick for the interminable conversations about life and physics (inevitable if you’ve been housemates with this guy for four years). Back at Hopkins, my fellow grad students in condensed matter theory—Damien Benveniste, Wes Fuhrman, Anirban Ghosh, Jian Kang, Imam Makhfudz, James Murray, and Yuan Wan—were generous with their time and helped me find my direction as a beginning graduate student. My experimentalist buddies at Hopkins—Nick Laurita, JT Mlack, and Mike Valentine—were fantastic comrades, and I thank them for all the fun and not-so-fun times spent inside and outside Bloomberg, for the infrequent but memorable reunions, and for the impromptu lessons in experimental methods. Finally I should mention my many friends outside physics whose presence—whether constant or intermittent—has enriched my life in one way or another. Jed Alegado, Toni Ang, Anj Aquino, Herdeline Ardona, Virgil Banta, Euben Concepcion, Uzein Corcuera, Icai Enriquez, Jay-El Esguerra, Paul Francisco, Srikar Ganti, JP Garcia, Mark Hernandez, Guy Hotson, Ehsan Jahangiri, Jeanette Kindipan-Dulawan, Lean Lava, Lex Maglunob, Mitch Sayoc, Joanna San Pedro, Cheska Siongco,

ACKNOWLEDGEMENTS

Jerome Unidad, Paul Villa, and Paulo Zaragoza, thank you guys for the travels, reunions, basketball games, brunches, dinners, cookouts, etc.—you know where you fit in this scheme!—that have made graduate school life as fun as it could be.

Finally, my stint in graduate school would not have been possible had it not been for the persistent love, support, and encouragement from my family. My parents have been tremendously supportive of my career choices and have been constantly patient and understanding. I am very happy to know that they'll always support me, even if they live halfway across the world from me (literally). My sister Thea and I are pretty much on the same wavelength nearly all the time, despite working in vastly different fields, and I am grateful for her advice and support.

LIST OF PUBLICATIONS

- M. A. Sulangi and Q. M. Sugon Jr., “The effect of the geometric potential and an external magnetic field on a charged particle on a helicoid,” arXiv:1211.5699 (2012).
- M. A. Sulangi, M. P. Allan, and J. Zaanen, “Revisiting quasiparticle scattering interference in high-temperature superconductors: The problem of narrow peaks,” *Phys. Rev. B* **96**, 134507 (2017). **(Chapter 3.)**
- M. A. Sulangi and J. Zaanen, “Quasiparticle density of states, localization, and distributed disorder in the cuprate superconductors,” *Phys. Rev. B* **97**, 144512 (2018). **(Chapter 4.)**
- M. A. Sulangi and J. Zaanen, “Self-energies and quasiparticle scattering interference,” to be submitted. **(Chapter 5.)**
- M. A. Sulangi and J. Zaanen, “Superconducting phase fluctuations and gap-filling phenomenology in the underdoped cuprates,” in preparation.

CURRICULUM VITAE

I was born on September 5, 1989 in Manila, Philippines. I spent my childhood in Bataan and Aklan before moving back to Manila to study at Ateneo de Manila University, from which I graduated with a B.S. in Physics, *cum laude*, in 2011. As an undergraduate I did thesis research on the quantum mechanics of charged particles on curved surfaces under the supervision of Germelino Abito and Quirino Sugon, for which I received the Bank of the Philippine Islands-Department of Science and Technology Science Award in 2011. In addition to physics and mathematics coursework, I spent time as a member of the editorial staff of *Heights*, Ateneo's literary journal, and I was lucky to be a part of the Silliman National Writers Workshop in 2011 as a fellow for non-fiction.

After finishing my undergraduate studies, I moved to Baltimore, Maryland to begin graduate school at Johns Hopkins University. I joined Zlatko Tesanovic's research group, having decided to focus on condensed matter theory as a Ph.D. student, but that stint was unfortunately cut short by Zlatko's sudden and tragic passing in 2012. After receiving my M.A. in Physics and Astronomy from Johns Hopkins, I decided to move to Leiden, Netherlands in 2013 to start anew my Ph.D. work, this time in Jan Zaanen's group at Leiden University. I spent the vast majority of my time at Leiden working on high- T_c superconductivity, with much of that research work covered in this very thesis.

Throughout my time in graduate school I have attended several conferences, workshops, and summer schools in various countries—the

CURRICULUM VITAE

Netherlands, Poland, Spain, France, and the United States—and have presented my work a number of times. I have also served as a teaching assistant at Johns Hopkins, where I led discussion sessions for general physics and taught laboratory classes, and at Leiden, where I was a grader for the graduate-level condensed matter theory class.