



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

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.

