



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

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 CuO_2 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 CuO_2 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 $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ (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 CuO_2 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 muddled 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.