



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

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

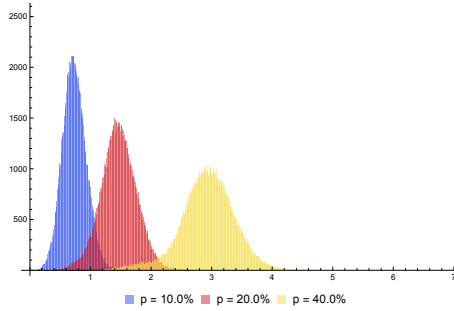


Figure 4.7: Histogram of the values of the disorder potential for smooth disorder with positive net potential for three values of p . The width of each bin is 0.01. Notice that the mean of the disorder potential is nonzero, leading to a shift in the average chemical potential of the overall system.

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

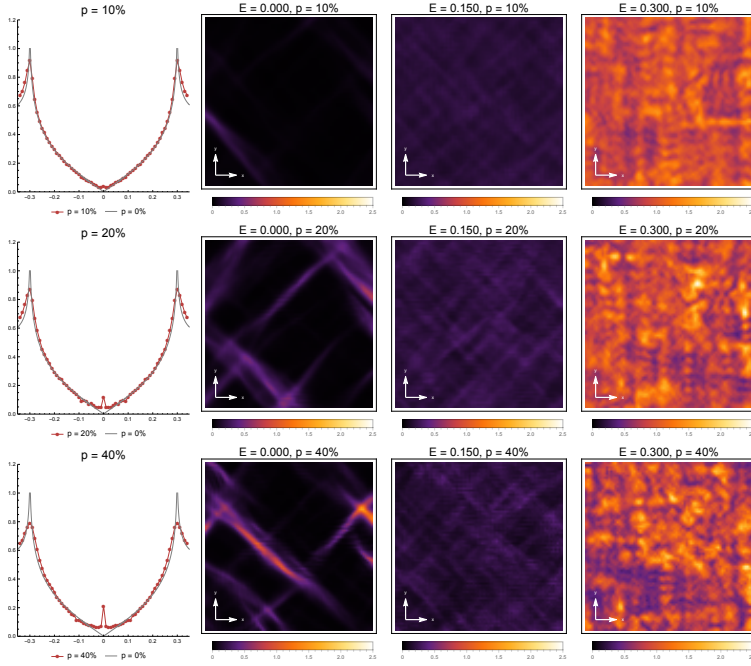


Figure 4.9: Snapshots of the real-space quasiparticle density of states for smooth disorder (with *zero* net potential) with increasing impurity concentration p (top to bottom) and energy E (left to right), extracted from the middlemost 80×80 subset of the full system. The leftmost column shows plots 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

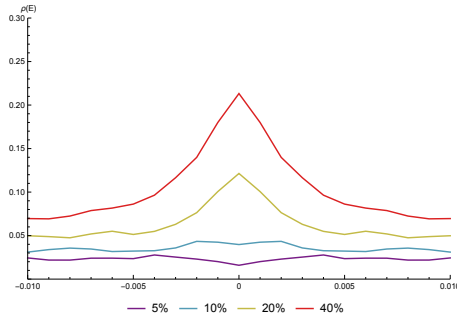


Figure 4.11: Plot of the density of states at and near $E = 0$ for the multiple-smooth-scatterer case for various impurity concentrations p . For the $p = 20\%$ and $p = 40\%$ the resonance is seen to have a width of approximately 0.006.

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

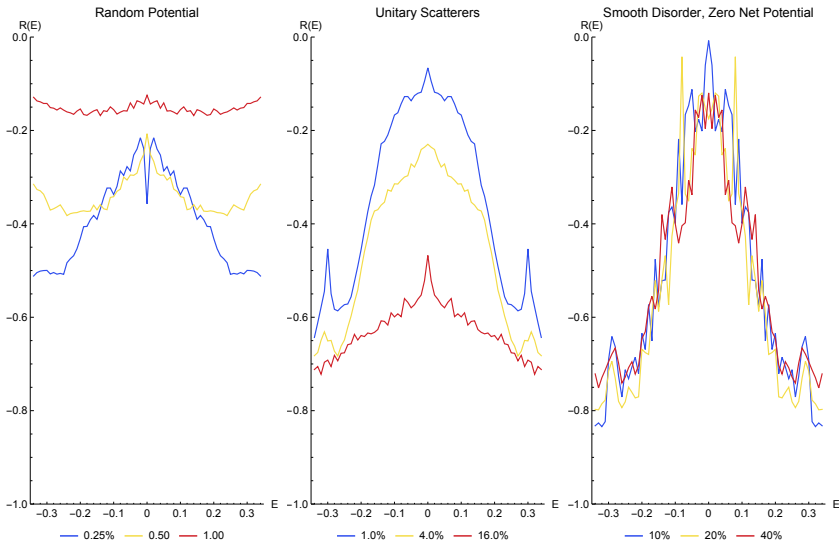


Figure 4.14: Plot of the correlation coefficient R between the local density of states in the middlemost 80×80 patch of the system and the disorder potential in that region for different types of disorder, for varying disorder strength, as a function of energy. For all three plots the correlation coefficient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

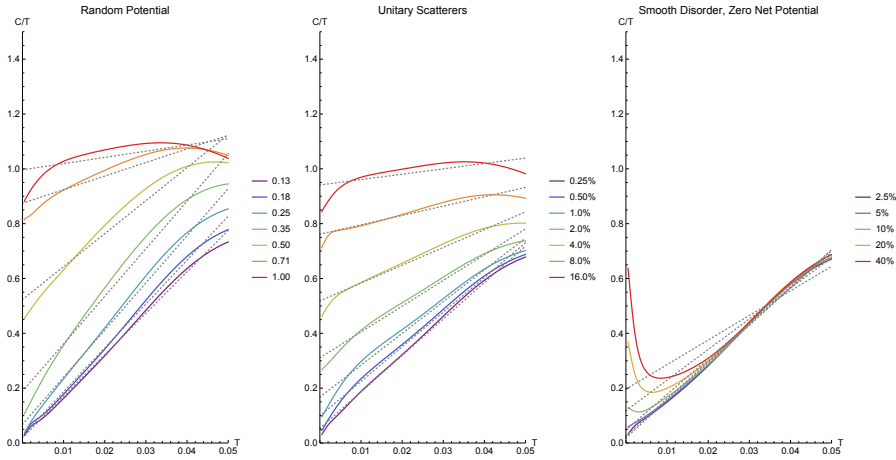


Figure 4.17: Plots of C/T as a function of temperature T for different types of disorder. Gray dotted lines indicate fits of the C/T curves to the form $C/T = \gamma_0 + \alpha T$, the scaling expected from d -wave quasiparticles with a nonzero DOS at $E = 0$. The numerically-obtained C/T exhibits visible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

4.8 QUASIPARTICLE LOCALIZATION

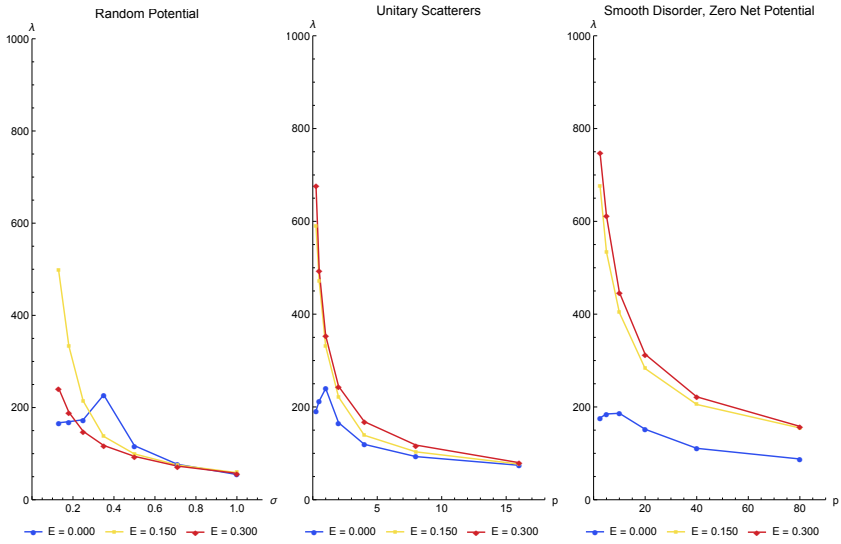


Figure 4.18: Plots of the localization length λ (in units where the lattice constant $a = 1$) for different types of disorder, taken at three different energies E . The x -axis shows the disorder strength parameter, given by σ for random-potential disorder (leftmost plot) and the impurity concentration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

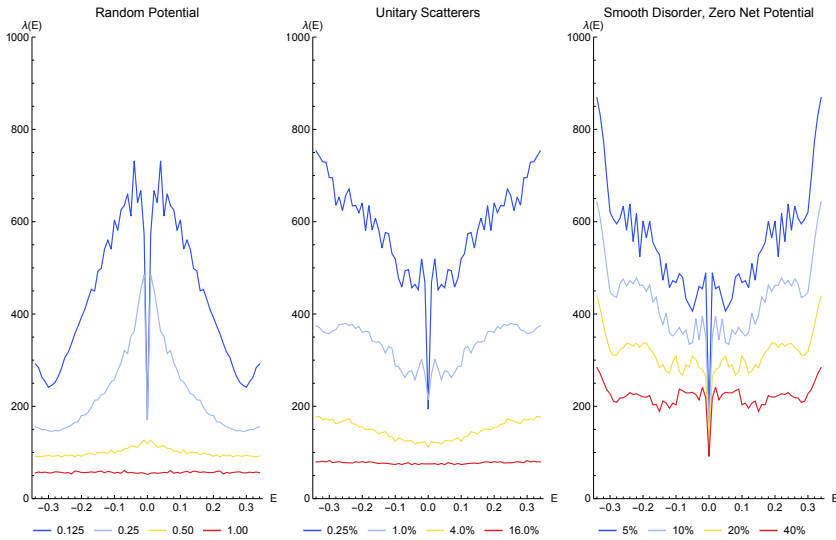


Figure 4.19: Plot of the localization length λ as a function of energy E for different types of disorder.

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.