



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

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

3.1 INTRODUCTION

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

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

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




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

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

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




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

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

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




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

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

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

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

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

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

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

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

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

3.2 MODEL AND METHODS

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

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

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

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

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

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

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

where

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

and

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

and

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

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

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

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

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

3.2.2 Modeling the Measurement Process

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

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

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

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

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

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

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

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

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

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




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

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

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

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

3.3 POINTLIKE SCATTERERS

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

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

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

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




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

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

3.3.1 Single Weak Pointlike Impurity

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




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

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

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

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

3.3.2 *Multiple Weak Pointlike Impurities*

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




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

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

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




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

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

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




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

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

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

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




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

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

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




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

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

3.3.3 Multiple Unitary Pointlike Impurities

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




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

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




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

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

3.3.4 Dependence of the Power Spectrum on the Impurity Strength

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

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

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

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

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

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




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

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

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

3.4 SMOOTH DISORDER

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




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

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

port and single-particle lifetimes; the former, which depends heavily on large-momentum scattering, is much smaller than the latter, and hence any scattering that occurs is argued to be forward (*i.e.*, small-momentum) scattering due to impurities located off the CuO₂ planes [2, 172, 3].

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

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

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

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

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

3.4.1 Single Smooth Scatterer

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

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

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

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

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




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

3.4.2 Multiple Smooth Scatterers

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

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




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

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

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




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

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

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

3.4.3 *Quantifying the Range of the Potential*

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

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

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

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

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

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




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

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

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

3.5 SPATIALLY RANDOM ON-SITE ENERGIES




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

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

3.5 SPATIALLY RANDOM ON-SITE ENERGIES

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




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

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

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

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

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

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

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

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

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

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

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

where

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

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

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

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

3.6 SPATIALLY RANDOM SUPERCONDUCTING GAP




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

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

3.6 SPATIALLY RANDOM SUPERCONDUCTING GAP

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




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

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

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

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

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

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

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

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

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

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

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

3.7 DISCUSSION AND CONCLUSION

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

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

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

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

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

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

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

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

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

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

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

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

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

3.A APPENDIX: SINGLE UNITARY POINTLIKE SCATTERER

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




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

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

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

3.A APPENDIX: SINGLE UNITARY POINTLIKE SCATTERER




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

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

