

LOCAL MEASUREMENTS OF THE SUPERCONDUCTING
PENETRATION DEPTH

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Tom Lippman
December 2013

Abstract

Measurements of the superconducting penetration depth, which is closely related to the superfluid density, are among the first characterizations of new superconductors. These measurements give insight into the gap structure, a key stepping stone to a microscopic understanding of the system.

I will review several techniques for measuring the penetration depth, and describe how we can use scanning SQUID microscopy to measure it locally. These local measurements often reveal spatial inhomogeneity in the penetration depth. Spatial inhomogeneity is a problem, since every experimental probe of the penetration depth assumes uniformity. This discrepancy motivated us to study penetration depth inhomogeneity in two ways: by creating and measuring a model system with controlled local penetration depth, and by calculation.

Because most unconventional superconducting materials are created by doping a non-superconducting parent compound, we expect there to be nanoscale inhomogeneity in the superconducting properties, including the penetration depth. In this situation, what will be the result of a measurement at longer length scales? I designed and analyzed a theoretically tractable model, and showed that the effective (measured) penetration depth is different from the average penetration depth and is sensitive to the details of the disorder.

As a controllable model system for experiments, we chose arrays of superconducting dots on a normal metal, which create a proximity effect Josephson junction array. By using 200 nm diameter niobium dots on gold, the SQUID response is dominated by screening currents flowing in the proximitized layer. No signature of the diamagnetism of the individual dots is seen, so we effectively have an artificial thin-film superconductor. The strength of the superconductivity is set by the dot separation, so by creating disordered arrays we

have our model system of a non-uniform penetration depth. I will show measurements on three such arrays, each with a different disorder correlation length but identical average density and disorder.

Acknowledgements

Completing this thesis has been a long and difficult journey, and I want to thank everyone who helped me along the way.

Working with Kam and the rest of the Moler lab has been a great experience. I had a lot of freedom to choose my own path. In fact, one of the first things Kam ever said to me was that her goal is to help me finish my thesis, rather than have me join her research program. This helped me develop independence and scientific judgment, and I'm grateful for that.

The Moler lab has always been a great community to learn from and work with. There is always someone willing to lend a hand when you need help or give advice when you're stuck.

In my time at Stanford I have been able to work with many great collaborators. Early on I worked closely with Jim Analytis, who made most of the pnictide samples that our lab studied. More recently I worked with Malcolm Durkin and Nadya Mason, who fabricated the superconducting arrays in Chapter 5.

I have been very lucky to meet many wonderful friends and all-around awesome people at Stanford. We suffered through classes and the qual together, and we've gone on ski trips, bike rides, and enjoyed some good beer. My friends have been an indispensable network of support and advice as we went through grad school and navigated the transition to whatever comes next.

Many thanks go to my family, for encouraging the curiosity that led me to grad school and teaching me the persistence I needed to eventually finish.

And to Nicole, for all of the ups and downs we shared. You were always there for me and you always believed in me, even when I didn't. I would not have made it without your love and support.

Contents

Abstract	iv
Acknowledgements	vi
1 Introduction to the Penetration Depth	1
1.1 Superconductivity	1
1.2 The Penetration Depth	2
1.3 Tunnel Diode Resonator	4
1.3.1 Analysis	5
1.3.2 Absolute Value of λ	8
1.4 Muons	9
1.4.1 Vortex Lattice	10
1.4.2 Low Energy	12
1.5 Local Probes	14
1.5.1 Vortex Profiles	17
1.5.2 Touchdown Curves	17
2 Calculating the Response Fields of Superconductors	19
2.1 SQUID Background	19
2.2 Meissner Response	23
2.2.1 Internal Solution	29
2.3 Two-Coil Model	31
2.3.1 Useful Limits of the SQUID Signal	32
2.3.2 Supercurrents in a Film	35

3	Random Superfluid Density and λ	41
3.1	Introduction	42
3.2	Stochastic London Equation	43
3.3	Results	46
3.3.1	Evaluating the Averages	46
3.3.2	Correlation Function	50
3.3.3	Squared Exponential Correlations	50
3.3.4	General Matérn Correlations	53
3.4	Discussion	57
4	Designing Disordered Dot Arrays	60
4.1	Random Variables	60
4.2	Correlations	62
4.3	Generating Correlated Random Numbers	66
4.3.1	The Characteristic Function	68
4.4	Matérn Derivations	70
4.4.1	Exponential	71
4.4.2	Squared Exponential	71
4.4.3	Spectral Function	71
4.4.4	Filter Kernel	72
4.5	Array Structure and Superconductivity	73
4.5.1	Nearest Neighbor Distance	74
4.5.2	Matching Local Properties	78
5	SQUID Measurements of Dot Arrays	81
5.1	Methods	81
5.2	Results and Discussion	86
5.3	Future Possibilities	94
A	Capacitive Touchdown Detection	98
A.1	Introduction	98
A.2	Capacitance	101

A.2.1	Sensitivity	104
A.3	Reducing the Extent (of the damage)	106
A.3.1	Minimum Achievable Extent	106
A.3.2	Cantilever Improvements	107
A.3.3	Data-Taking Procedure	109
A.3.4	Detection Algorithms	110
A.3.5	Goddamnit, MATLAB	113

List of Tables

4.1 Mean and standard deviation of the nearest neighbor distance. 78

5.1 Fitted Pearl length of four dot arrays. 94

List of Figures

1.1	Geometry of cavity measurements.	6
1.2	Touchdown curve data taking procedure.	18
2.1	Summary of scanning SQUID devices.	22
2.2	SQUID derivation geometry.	24
2.3	Calculated touchdown curves in the two-coil model.	33
2.4	Supercurrent response of a thin film, two-coil model.	36
2.5	The f -critical radius of a thin film, two-coil model.	37
2.6	Differential SQUID flux response of a thin film, two-coil model.	38
2.7	The f -critical flux radius of a thin film, two-coil model.	40
2.8	Comparison of the flux and current f -critical radii in a thin film, two-coil model.	40
3.1	Effect of variance and correlation length on random variable	51
3.2	λ_{eff} for squared exponential correlations	53
3.3	Linecuts of λ_{eff} for squared exponential correlations	54
3.4	λ_{eff} for Matérn correlations, $\nu = 2$	57
4.1	Three example disordered dot arrays.	63
4.2	Illustrating the effect of the Matérn smoothness parameter on a random field.	64
4.3	The effect of the Matérn smoothness parameter on the dot arrays.	66
4.4	Diagram of dot displacements.	75
4.5	The effect of the Matérn smoothness parameter on the dot arrays, after matching the local structure.	80

5.1	Optical photographs of the five arrays.	82
5.2	Scanning electron microscopy images of two disordered dot arrays.	83
5.3	Structure analysis of the three disordered arrays.	85
5.4	SQUID susceptibility images of all five dot arrays.	87
5.5	Comparison of two arrays with correlated disorder.	89
5.6	Scanning temperature dependence of DS3.	90
5.7	Isolated superconducting islands near T_c in DS3.	92
5.8	Touchdown curves at base temperature.	93
5.9	Structure analysis of the proposed disorder series with correlated disorder.	96
5.10	Structure analysis of the proposed disorder series with uncorrelated disorder.	97
A.1	Capacitive touchdown overview.	101
A.2	Capacitance offset	102
A.3	Two capacitive touchdown curves and their noise.	103
A.4	Running cutoff touchdown detection.	110
A.5	The perils of believing the MATLAB documentation.	117

Chapter 1

Introduction to the Penetration Depth

In this chapter I will briefly introduce superconductivity, in particular the penetration depth and how it can contribute to our understanding of the underlying superconducting state. Then I will describe several of the most common experimental methods for determining the penetration depth. This will give context for the advantages and disadvantages of the local probe methods I used in this thesis.

1.1 Superconductivity

If you had some time and liquid helium on your hands, you could cool a piece of lead down to 4.2 K, and you would observe its electrical resistance drop to zero. This is because at 7.2 K,[1] the electrons in lead go through a phase transition, from a metal to a superconductor. The transition occurs because the electrons in the lead begin to form bound pairs, which then condense into a quantum mechanical ground state.[2, 3] This ground state has a well-defined phase, which is the origin of the two properties that define superconductivity: zero resistance and perfect diamagnetism.

The electrical resistance of a superconductor is well and truly zero, not just “very small”. The best evidence for this are measurements of the decay of the supercurrent in a superconducting ring. The current flowing through the ring can be measured inductively, so there are no leads to interrupt the supercurrent and introduce dissipation. The ring forms an LR circuit, so current in the ring will decay exponentially with a decay time given by

the geometric self-inductance divided by the resistance. So zero resistance implies infinite decay time. The largest decay time observed to date is over 10^5 years.[4] Even then, the authors pointed out that the observed field decay could be attributed to expansion of the solenoid due to radial magnetic pressure. Specifically, an increase of the solenoid radius by 2 microinches (51 nm) would be sufficient. (The nominal coil diameter was 4 inches.)

The majority of the applications of superconductivity make use of the zero resistance property. Most common is the use of superconducting wire to make extremely powerful magnets, thus enabling Magnetic Resonance Imaging, particle accelerators like the Large Hadron Collider, and maglev trains. Another use is in radiofrequency applications, to make cavities with ludicrously high quality factors. These find use in cell phone towers, where microwave filters with demanding specifications are needed. Of course, the most tantalizing potential application of the zero resistance of superconductors is in the electric grid, for power transmission.

The drawback in all of these applications is the need to keep the superconductor below its transition temperature. This requires cryogenic temperatures, substantially adding to cost and complexity. In metallic elements and their alloys, used to make common superconducting wire, the transition temperature tops out around 23 K. Exotic materials exist with higher transitions, most notably the cuprates. Many compounds in this family are superconducting at the boiling point of liquid nitrogen (77 K), which is much cheaper than the liquid helium needed for the metallic superconductors. Under special conditions, the transition can go as high as 138 K, not quite a cold day in Antarctica.[5] Unfortunately, other materials properties have prevented the cuprates from being widely deployed.

We have a detailed understanding of how and why elemental metals and their alloys are superconducting,[3] but the origin of superconductivity in the cuprates remains a mystery. The goal of basic research in superconductivity is to understand these novel materials, so that understanding can guide us towards more and better materials for broader application.

1.2 The Penetration Depth

The penetration depth comes from the second defining property of superconductivity — perfect diamagnetism. If you put a metal sphere into a background magnetic field, the field

lines will simply go through the sphere unperturbed. But if that sphere became superconducting, perhaps because it is made of lead and you lowered the temperature from 10 K to 4.2 K, now the field lines go *around* the sphere. The field has been expelled from the interior of the sample.

To achieve this, the superconductor sets up screening currents flowing near its edge. The field resulting from the screening currents cancels the applied field inside the sample, and modifies the total field outside the sample so that it smoothly flows around the sphere. It turns out that the magnitude of both the total field and the screening currents decay exponentially with distance as you travel into the sample. The decay length of the field and screening currents defines the superconducting penetration depth, denoted λ .

We are primarily interested in the temperature dependence of λ , because it probes the energy required to break apart a bound pair of electrons. As temperature increases, thermal energy can break apart more and more electron pairs, which kicks those electrons out of the superconducting ground state. This reduces the ability of the superconductor to screen fields, which increases λ . Loosely speaking, we are just counting the number of thermally excited quasiparticles. This makes penetration depth measurements complementary to measurements of the heat capacity and thermal conductivity, because these also probe thermally excited quasiparticles but differ in the details, giving them each advantages and disadvantages relative to one another.

At the simplest level, measurements of the temperature-induced change in λ , denoted $\Delta\lambda(T)$, can discriminate between conventional and unconventional superconducting states based solely on “the shape of the curve”. A conventional superconductor has essentially no change in its penetration depth at temperatures small compared to the transition temperature. Specifically, it has an exponentially-activated form, with

$$\Delta\lambda(T) \propto \sqrt{\Delta_0/T} \exp(-\Delta_0/T),$$

where Δ_0 is the zero-temperature value of the energy gap, or half of the energy needed to break apart a bound pair of electrons. In contrast, the cuprates have a linear slope in the

penetration depth, observation of which was among the first evidence for d-wave superconductivity.[6] It is possible to extract more detailed, quantitative information from measurements of the penetration depth, but this requires an absolute, rather than a differential, measurement of λ . The absolute magnitude of λ allows determination of the normalized superfluid density, $\rho_s = (\lambda(0)/\lambda(T))^2$, which is easier to access theoretically.

There are many different techniques to measure the penetration depth, each with its own advantages and disadvantages. I will describe three here, but there are many more. The most glaring omission is microwave cavity measurements; a good review is Ref. [7]. Frankly, I do not understand the details particularly well. Some microwave cavity measurements, for example Ref. [6], report a differential measurement of λ , and the analysis is identical to Section 1.3. Others report the absolute magnitude of λ , by extending the measurement into the normal state and using the normal metal skin depth as a yardstick.[7] I believe the difference is whether one is measuring the frequency shift or quality factor. Another omitted technique is two-coil mutual inductance. A thin-film sample is placed between two coaxial coils and modifies the mutual inductance between them.[8, 9] These measurements are similar our Scanning SQUID measurements, which are described later.

1.3 Tunnel Diode Resonator

The Tunnel Diode Resonator (TDR) technique is a cavity perturbation method that relates changes in the resonant frequency of an LC tank circuit to changes in the penetration depth. The tank circuit is typically a copper inductor and lump capacitor, and is driven to self-resonate by the negative differential resistance of the tunnel diode that gives the technique its name. The quality factor is deliberately kept low (~ 100) in order to keep the excitation field small, of order 20 mOe. Resonance frequencies are typically in the 10–15 MHz range. The sample is mounted on a sapphire cold finger, so the resonator can be maintained at a fixed temperature (by immersing it in liquid helium) for improved stability. Another benefit of the sapphire cold finger is that it allows the sample to be removed from the resonator *in situ*, which provides a calibration for the measurement. The low noise and reliability of this technique have led to its widespread use. A veritable industry has arisen to measure $\Delta\lambda$ in every novel superconductor under the sun.[10]

1.3.1 Analysis

The directly measured quantity is the change in the resonance frequency of the TDR. With no sample present, the TDR runs at frequency $2\pi f_0 = 1/\sqrt{LC}$. Introducing the sample into the cavity causes a change in inductance ΔL , which causes a fractional change in resonant frequency:

$$\frac{\Delta f}{f_0} = \frac{-\Delta L}{2L}. \quad (1.1)$$

The simplest estimate of the inductance change will be the total magnetic field excluded from the sample volume. For example, consider a perfect diamagnet placed in a region of uniform field H_0 . To the extent we can ignore the field lines that will bend around the sample, the change in cavity inductance will be $-H_0 V_s$, where V_s is the volume of the sample. The unperturbed inductance will scale as $H_0 V_c$, where V_c is the effective volume of the cavity, giving:

$$\frac{\Delta f}{f_0} = \frac{V_s}{2V_c}. \quad (1.2)$$

To add a nonzero penetration depth, we can approximate the field in the sample as H_0 for the first λ of thickness on each side of the sample, and 0 beyond that. Then for a sample of thickness $2c$, the penetrated volume is reduced by a factor of $1 - \lambda/c$, so we estimate:

$$\frac{\Delta f}{f_0} = \frac{V_s}{2V_c} \left(1 - \frac{\lambda}{c}\right). \quad (1.3)$$

With this formula, we're assuming that $\lambda \ll c$ and that the geometry is as pictured in Fig. 1.1(a). A platelet sample, with two dimensions much larger than the third, is in a region of uniform applied field with the field applied parallel to the broad face. We've neglected the contributions to the screening from the thin faces and the field lines bending around the sample at the top and bottom. Under these assumptions, the sample's internal field distribution is identical to that of a semi-infinite slab of the same thickness.

Since we're already approximating the sample response as that of a semi-infinite slab, we can use the exact solution of that problem to relax the assumption $\lambda \ll c$. Suppose a slab of superconductor with penetration depth λ fills the space from $x = 0$ to $x = 2c$ with infinite extent in the yz plane. If the applied field is H_0 on both sides, then the internal field

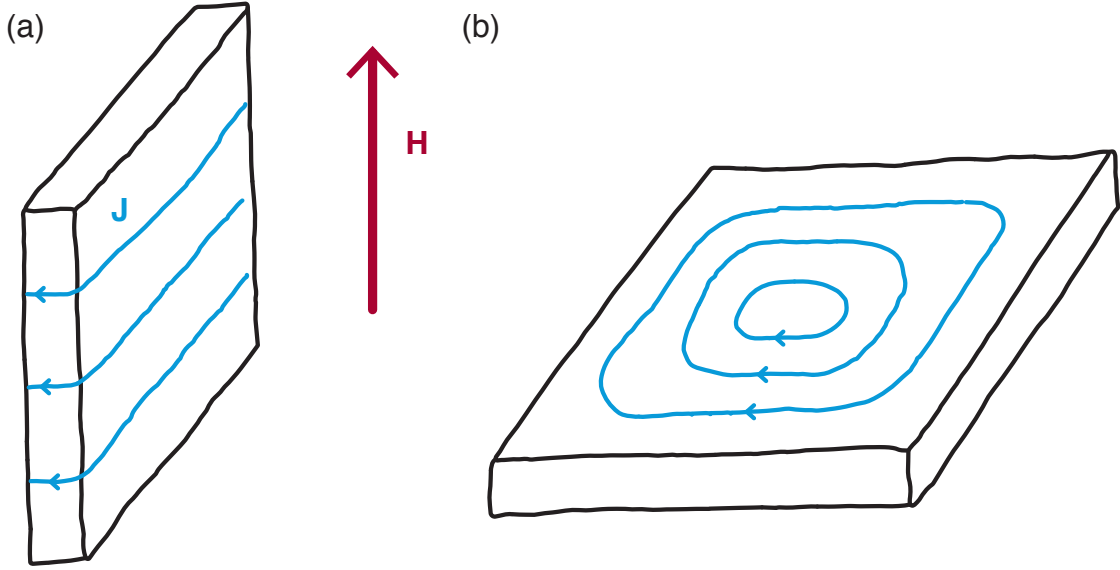


Figure 1.1: Sketch of the screening currents in a thin sample in the longitudinal (a) and transverse (b) field configurations.

solution in the Meissner state is:

$$h_{\text{int}} = H_0 \frac{\cosh((x-c)\lambda)}{\cosh(c/\lambda)}. \quad (1.4)$$

The field penetrating the region per unit cross sectional length is:

$$\begin{aligned} \int_0^{2c} h_{\text{int}} dx &= \frac{H_0}{\cosh(c/\lambda)} \int_0^{2c} \cosh\left(\frac{x-c}{\lambda}\right) dx, \\ &= 2\lambda H_0 \tanh \frac{c}{\lambda}. \end{aligned} \quad (1.5)$$

Multiplying by the cross-sectional area of the sample's broad faces, we find the frequency change due to the sample is:

$$\frac{\Delta f}{f_0} = \frac{V_s}{2V_c} \left(1 - \frac{\lambda}{c} \tanh \frac{c}{\lambda}\right) \quad (1.6)$$

(because $2\lambda A = V_s \lambda / c$). Pleasingly, this matches Eq. 1.3 in the usual limit that $\lambda \ll c$. But

the conversion from frequency shift to penetration depth requires knowing the sample and cavity volume, as well as the thickness of the sample. These are difficult to measure directly, which is why TDR is typically used as a differential measurement of the penetration depth.

To calibrate the geometrical effects in the frequency shift, define Δf_0 to be the frequency change on inserting the sample at the reference temperature T_0 :

$$\Delta f_0 = f_0 \frac{V_s}{2V_c} \left(1 - \frac{\lambda(T_0)}{c} \tanh \frac{c}{\lambda(T_0)} \right). \quad (1.7)$$

Now define δf , which is the change in frequency due to changing the sample temperature from T_0 .

$$\delta f(T) = f_0 \frac{V_s}{2V_c} \left(\frac{\lambda(T_0)}{c} \tanh \frac{c}{\lambda(T_0)} - \frac{\lambda(T)}{c} \tanh \frac{c}{\lambda(T)} \right). \quad (1.8)$$

Dividing Eq. 1.8 by Eq. 1.7 provides our calibration of the sample and cavity volumes. Taking the limit of $\lambda \ll c$ and re-arranging, we get:

$$\Delta \lambda(T) = -\delta f(T) \frac{c}{\Delta f_0}. \quad (1.9)$$

So to determine $\Delta \lambda(T)$, you need to measure two frequency shifts — one at base temperature on inserting the sample and sapphire cold finger into the resonator (Δf_0), the second as the temperature of the sample is slowly warmed and cooled ($\delta f(T)$) — and one length, the sample thickness $2c$. My hunch is that the dominant error is the uncertainty on measuring the sample thickness, because it would re-scale the inferred $\Delta \lambda$. In addition, thermal expansion can make c a function of temperature, which would contaminate the measured $\Delta \lambda(T)$. [6, 7]

The preceding analysis has neglected any effect of screening currents crossing the thin faces of the sample. This is fine if the penetration depth is isotropic (between the three crystal axes) and the sample aspect ratio is large. If the penetration depth is not isotropic this assumption may break down. Many materials of interest, such as the cuprates and pnictides, are layered compounds that grow as thin platelets with the c -axis perpendicular to the broad face. In this geometry the screening currents flowing along the broad faces probe λ_{ab} , and the currents flowing across the thin faces allow λ_c to mix into the measurement. Because λ_c is typically much larger than λ_{ab} , the ends can even come to dominate the

measurement. One way to get a handle on this is to measure multiple samples with different aspect ratios, or to measure one sample then cut pieces off of it and measure those.[6, 7, 11]

The more elegant approach is to measure in the transverse field geometry, sketched in Fig. 1.1(b). With the field applied along the c -axis, screening currents only flow in the ab plane and there is no contribution from λ_c . The drawback to this approach is that there is no solution to the London equation in this geometry, so the sample's contribution to the frequency shift is not precisely known. Prozorov and collaborators worked around this by combining numerical solutions of London's equation in two dimensions with approximate analytical expressions.[12] It's a fairly confusing paper, but the upshot is an approximate expression for the frequency shift due to the sample:

$$\frac{\Delta f}{f_0} = \frac{V_s}{2V_c(1-N)} \left(1 - \frac{\lambda}{R} \tanh\left(\frac{R}{\lambda}\right) \right), \quad (1.10)$$

where now R is an effective sample dimension and N is an effective demagnetization factor. Now the same procedure of inserting the sample at base temperature and changing the temperature will give you an expression just like Eq. 1.9, with c replaced by R . As a practical matter, R is typically taken to be $\sqrt{ab}/5$. I've yet to find any discussion of the effect of thermal expansion on measurements taken in the transverse geometry. Perhaps the ab plane has a much smaller coefficient of thermal expansion than the c -axis, and it is a non-issue.

1.3.2 Absolute Value of λ

An interesting modification of the TDR technique is to first coat the sample in a thin film of a second superconductor, typically aluminum, with a lower transition temperature than the bulk sample of interest.[10, 13, 14] The film thickness and penetration depth provide reference lengths, allowing one to infer the absolute value of the penetration depth in the sample of interest. Below the transition temperature of the aluminum, the aluminum screens the applied field, effectively increasing the linear dimensions of the sample by the film thickness. Once the aluminum becomes non-superconducting it does no screening, since the film thickness is chosen to be much smaller than the skin depth at the operating frequency. So the aluminum film transition causes a shift in the cavity frequency, providing another handhold on the calibration and allowing the absolute value of λ in the sample of interest

be recovered.

The drawbacks I see are that now λ is related to *three* other lengths: the aluminum thickness, the sample geometry, and the penetration depth of the aluminum film. It is also rather difficult to coat a three-dimensional sample with a film of uniform thickness. The penetration depth of the aluminum film is taken from the literature. This strikes me as a bad assumption, given that the *transition temperature* of aluminum can change depending on details of the film growth.[15]

1.4 Muons

The technique of muon spin rotation (μ SR) is unique in that it directly measures the magnetic field *inside* of the sample. A positive muon is implanted in the sample, where it precesses about the local magnetic field. So a measurement of its Larmor frequency, $\omega_\mu = \gamma_\mu B(r)$, gives the local magnetic field. The proportionality is the muon gyromagnetic ratio, $\gamma_\mu/2\pi = 135.5342$ MHz/T. But how do we measure its Larmor frequency? The poor muon is trapped in the sample. Fortunately, it decays with a 2.2 μ s half-life, emitting a positron and two neutrinos. All three decay products flee the scene with a quickness, and the positron is detected by one of the positron counters surrounding the sample.[16]

This is a frequency measurement in the time domain, and there are more than a few technical details that make it possible. Briefly, the positive muons are produced from the decay of a pion at rest, $\pi^+ \rightarrow \mu^+ + \nu_\mu$. The muons produced are spin-polarized, with spin opposite to momentum, and have a kinetic energy of 4.119 MeV. The spin polarization is essential for success of the technique. It comes courtesy of the fixed helicity of the neutrino and the spinlessness of the pion. The muons travel to the sample, where they implant at random. Before they impact the sample, the muons pass through a thin scintillation detector, which initializes the timing information.[16] This way the detected positron can be assigned to the proper time bin.

Once the muon enters the sample, it thermalizes with the lattice and comes to rest somewhere in the crystal. Typically it occupies an interstitial site, due to electrostatic repulsion. Importantly, the spin orientation is maintained through this process.[17] Once in the sample, the muon precesses around the local magnetic field and eventually decays as

$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$. The positron momentum is preferentially along the spin direction of the muon. This asymmetry allows the experimenter to track the time evolution of the muon spin, giving them the Larmor frequency and with it the local magnetic field.

Of course, what is actually measured is the superposition of decays from many muon events, each one at a random location in the sample. Which means the measurement is of the ensemble-averaged spin direction, usually called the muon spin polarization $P(t)$. [16] The spin polarization signal is analogous to the free induction decay signal of NMR, and its Fourier transform is the spectrum of local magnetic fields, the μ SR line shape. I have omitted the details of how raw histogram counts are converted into the time dependent muon polarization, see Ref. [16].

Since they are related by a Fourier transform, both the line shape and the spin polarization contain the information about λ that we're interested in. In practice, the data is fit in the time domain and the penetration depth extracted from the fit parameters. There are two distinct methods for determining the penetration depth from μ SR. The oldest and most common is to study the field profile in the vortex state. Here the vortex spacing provides a reference lengthscale and the penetration depth is a parameter describing the field distribution around the vortices. A newer technique uses low energy muons that stop within a few hundred nanometers of the surface and directly probe the field screening in the Meissner state.

1.4.1 Vortex Lattice

A superconducting vortex is a defect in the superconducting state wherein a bundle of magnetic field lines passes through the sample. This creates screening currents flowing in the superconductor, and the lengthscale over which the screening currents (and magnetic field) decay is just λ . If we could selectively implant muons certain distances from the vortex core, this would determine the decay length of the field, and be a measurement of the penetration depth. Current muon sources don't have this capability, so instead the technique requires applying a large (roughly 1 Tesla) to the sample, creating a vortex lattice. The spacing between vortices in the lattice is set by the field intensity, which is why it can serve as a reference lengthscale. As the penetration depth changes, changing the field profile between

vortices, the vortex spacing itself doesn't change.

To extract λ from this we will need a model of the field distribution. The simplest one to start with is the London model. We treat the vortices as one-dimensional sources in the London equation, and solve for \mathbf{h} :

$$\mathbf{h}(\mathbf{r}) + \lambda^2 \nabla \times \nabla \times \mathbf{h}(\mathbf{r}) = \Phi_0 \sum_n \delta(\mathbf{r} - \mathbf{r}_n) \hat{\mathbf{z}}, \quad (1.11)$$

where the \mathbf{r}_n are locations of the vortices and $\delta(\mathbf{r})$ is a two-dimensional delta function. The vortices are located on a lattice of some kind, so it's easiest to solve in Fourier modes. For a perfect vortex lattice, the field far from any edges will be entirely along $\hat{\mathbf{z}}$. The solution is:

$$\mathbf{h}(\mathbf{r}) = \sum_{\mathbf{K}} \frac{h_0 e^{-i\mathbf{K} \cdot \mathbf{r}}}{1 + K^2 \lambda^2} \hat{\mathbf{z}}, \quad (1.12)$$

where the sum is over all reciprocal lattice vectors \mathbf{K} , and h_0 is the average field. It is also the applied field, which sets the inter-vortex spacing. In the London model, the field diverges at the core of each vortex, which is unphysical. This can be fixed by rolling off the higher modes with a cutoff factor such as $\exp(-K^2 \xi^2/2)$, which does unfortunately introduce an additional fitting parameter.[16]

The internal field solution determines the muon lineshape, $n(B)$, essentially the histogram of internal magnetic fields. The penetration depth is determined by fitting the data to the model in the time domain, by calculating the complex muon polarization:[16]

$$P_x(t) + iP_y(t) = G(t) \int n(B) e^{i\gamma_\mu B t} \mathbf{dB}, \quad (1.13)$$

where $G(t)$ is a *depolarization function* that exists to account for additional sources of inhomogeneous magnetic fields. This can include nuclear dipole moments and inhomogeneity in the vortex lattice. The depolarization function is almost always Gaussian, usually with several sources, such as:

$$G(t) = e^{-(\sigma_d^2 + \sigma_f^2)t^2/2}, \quad (1.14)$$

used in Ref. [18]. In that work, σ_d determined by data above T_c and kept fixed. However the authors did find a correlation between σ_f and $1/\lambda^2$, which lead to a small increase in

claimed systematic uncertainty on the measurement. They attributed the effects of σ_f to pinning-induced disorder in the vortex lattice causing approximately five percent deviations from the ideal lattice positions.[18]

Many early μ SR studies of the cuprates did not fit to a complete model like the one described, but instead assumed a Gaussian lineshape and fit for a relaxation rate like σ in Eq. 1.14. The relaxation rate was then equated to the second moment of the field distribution found in the London model.[19–22] The second moment of the field distribution is calculated by adding the squares of all the Fourier modes in Eq. 1.12 with $K \neq 0$. This works out to:[23]

$$\langle h^2 \rangle - \langle h \rangle^2 = 0.00371 \Phi_0^2 \lambda^{-4}. \quad (1.15)$$

The drawbacks to this analysis technique were known at the time; indeed, they were pointed out by Brandt in Ref. [23]. I'm just going to quote from the abstract:

Fitting of the highly asymmetric $n(B)$ by a Gaussian or Lorentzian, as done in existing μ SR experiments, yields for a perfect vortex lattice λ values which are too large, and for a strongly distorted (due to pinning) lattice λ values which are too small. An improved evaluation of μ SR data is suggested.[23]

Even so, a simple analysis based on Eq. 1.15 is still used from time to time.[24–27]

1.4.2 Low Energy

The vortex lattice method is commonly used to infer the penetration depth in newly discovered superconducting systems. But the technique relies on the existence of a highly-ordered vortex lattice, which may not exist at all in some materials, making interpretation of these experiments problematic.[28] But the vortex lattice is needed, for without it there would be no field for the muons to precess around. This is because the incoming muons have a large kinetic energy (4.2 MeV) which implants them deep into the sample bulk.

The alternate technique of low energy μ SR uses low energy muons with stopping depths less than 100 nm. The stopping depths are tunable, allowing direct measurement of the magnetic field as a function of depth below the sample surface in the Meissner state. This is quite special, as every other technique to measure the penetration depth assumes the

London equation as a starting point. In addition, it is an absolute measurement, since the stopping depth profile is a known function of beam energy.

The technique starts with your standard intense beam of 4.2 MeV muons, which are moderated by (implanted in) a layer of Van der Waals ice such as Ar, Ne, or N₂. Something happens in the ice, and out pops epithermal muons with kinetic energies of around 15 eV. These can then be electrostatically accelerated, manipulated, and lensed, so the outcome of all of this is a beam of muons, with energy tunable from 15 eV to 30 keV, headed for the sample.[29] The spin polarization is maintained through all this, so standard μ SR techniques can be used from there.

The muons have a range of stopping distances, depending on the beam energy. This can be determined by Monte Carlo calculations, resulting in an implantation depth probability density, $n(z, E)$. So running the beam at constant energy still gives a superposition of decay frequencies, depending on the field depth profile, $B(z)$. The measured data are still polarization functions in the time domain, but now we can consider the data as a superposition of field (Eq. 1.13) or as a superposition of implantation depths:

$$P_x(t) + iP_y(t) = G(t) \int n(z, E) e^{i\gamma_\mu B(z)t} \mathbf{d}z. \quad (1.16)$$

From here, there are several approaches to the analysis. The simplest is to assume a model for $B(z)$, whether a vortex lattice near the surface,[29] or an exponential decay such as you get from the London equation.[30, 31] With a model assumed for $B(z)$ and $n(z, E)$ known, the data can be fit in the time domain just as in the vortex lattice technique. This is a perfectly reasonable thing to do, and has given some very nice results. But it is not a *direct* measurement of the penetration depth, no matter what the title of the paper says.[30]

The second and third analysis techniques in the literature to date first invert Eq. 1.13, either by Fourier transform or maximum entropy methods, to get $n(B)$. Now we have in hand the calculated depth probability, $n(z, E)$, and the measured field probability, $n(B)$. It sure seems like that should be enough to get the field profile, $B(z)$, doesn't it? As a first approximation, we can pick off the peak field and plot it against the peak depth.[32] A second approximation would be to take first moments of field and depth, to plot average field versus average implantation depth. This will only equal the true $B(z)$ if the field profile

is linear, due to Jensen's inequality, explained in Section 4.5.1.

For a more detailed comparison, the equality of Eqs. 1.13 and 1.16 implies:

$$n(z, E)dz = n(B)dB. \quad (1.17)$$

If we assume that $B(z)$ is a monotonically decaying function of z , then we can use Eq. 1.17 to get an integral equation for the field at a specified z :

$$\int_0^z n(z', E)dz' = \int_{B(z)}^{\infty} n(B', E)dB'. \quad (1.18)$$

This equation can then be used to retrieve the *directly-measured* field profile $B(z)$. This method of analysis was developed and exploited to observe nonlocal effects in lead.[33, 34] In those papers the authors showed that $B(z)$ deviates from exponential over a decade of decay, and is consistent with the Pippard nonlocal electrodynamics.

One drawback to this technique is the maximum stopping distance of the muons, currently around 200 nm. This makes it difficult to convincingly probe the field screening in superconductors with penetration depths larger than 100 nm, which includes both the cuprates and pnictides. An additional problem is the necessity of including some kind of “dead layer” in the data analysis. There are very good reasons for it to exist, particularly surface roughness mixing λ_c into the screening and the surface potentially suppressing the superfluid density, but currently it is an additional fit parameter.

1.5 Local Probes

By “local probes”, I mean any technique that measures magnetic fields on lengthscales much smaller than the sample size. This opens the possibility of measuring spatial variation in the sample's magnetic response. The measurement probe itself can be any number of things, whether Hall probe, scanning SQUID, or MFM. I'll describe a few techniques in more detail below. The common thread among all local probes techniques is that the probe is moved over the sample surface. Then the position dependence can be used for analysis of the sample properties.

Local probe measurements of the penetration depth fall into two categories, regardless of the experimental technique used. The first category, described in Section 1.5.1, uses the field profile of a superconducting vortex to determine the penetration depth. The second category, described in Section 1.5.2, applies a small excitation field to the sample and measures the field generated by the sample's Meissner response. Determining the penetration depth from this requires a model of the probe field and taking touchdown curves by changing the sample-probe separation.

The scanning Hall probe is the simplest local magnetic probe. Typically an active region is lithographically defined in a GaAs-AlGaAs heterostructure which supports a two-dimensional electron gas (2DEG) at the interface. Leads to the active region allow the experimenter to measure the Hall effect in the active region, giving a direct measurement of the average magnetic field perpendicular to the 2DEG in the region. I want to emphasize the importance of this, as it is a significant advantage of the Hall probe technique. We will soon see that quantitative extraction of the penetration depth from scanning SQUID or MFM is hampered by the need to model probe response. In contrast, the Hall probe directly measures the magnetic field perpendicular to the 2DEG. If that field varies slowly enough, we don't need to know the size of the active region at all. Even for rapidly changing fields, where the detected feature will be resolution-limited, the effect of the device resolution is a simple area average.

Other advantages of the Hall probe are its ability to be used at high temperature and high magnetic field. Disadvantages include a high level of low-frequency electronic noise, which becomes substantially worse on reducing the size of the active area.[35] Another disadvantage is the fragility of GaAs. Devices are more susceptible electrostatic discharge and breaking due to impact. Such impact is more common than you might think. Step one of scanning a microprobe is finding the sample surface. Magnetic probes have no built-in way to do this, unlike a force microscope or tunnel probe. Our lab has historically solved this by putting the probe on a metal cantilever that forms half of a capacitor. When the tip of the probe touches the sample, the cantilever deflects and changes the capacitance. For more, see Appendix A. Point being, touching the sample with the probe happens a lot. This can kill delicate GaAs Hall probes.

Similar in spirit to the scanning Hall probe is the scanning Superconducting QUantum

Interference Device (SQUID). A SQUID is a superconducting loop interrupted by one or two weak links. The weak link(s) make the loop's critical current a sensitive function of the magnetic flux in the loop. Measuring the critical current can happen one of several ways, see Section 2.1, but the end result is an indirect measurement of the magnetic field in the sensor active area. The primary advantage of the SQUID is its incredible sensitivity combined with low noise. Another advantage of the SQUIDS used in our lab is that they are very, very robust. They take a lickin' and keep on tickin'. I have taken good data with a device that was dropped on the lab floor. Disadvantages are primarily the limited field and temperature ranges the devices can be used in. This doesn't rule out measuring high temperature samples, but it does make it more difficult. Another disadvantage is the indirect nature of the field measurement alluded to previously.

Rather different in spirit from SQUIDS and Hall probes is the Magnetic Force Microscope (MFM). In force microscopy, a flexible silicon cantilever is brought close to the sample. At the end of the cantilever is a large pyramid which interacts with the sample. A changing interaction force between the sample and the pyramid deflects the cantilever. This deflection is usually detected optically, but there are piezoresistive cantilevers in use. The experimenter can choose to measure the static deflection of the cantilever, giving a measurement of the interaction force, or they can drive the cantilever at resonance and measure changes in the resonance frequency or quality factor, measuring the spatial gradient of the interaction force.

The tip-sample interaction force can arise from a number of physical effects, and in fact all are present at any given time. What distinguishes force microscopies is which interaction force is dominant. When the tip is very close to the sample, the Van der Waals force dominates, and the measurement can resolve purely the topography of the sample. Atomic resolution is even possible. By coating the tip in a magnetic layer the magnetic features of the sample can be measured, whether it is a hard drive or a superconductor. But for the magnetic interaction force to be dominant, all other interactions must be minimized. The Van der Waals interaction is extremely short-ranged; scanning 5 nm above the surface makes this negligible. The electrostatic interaction is more troublesome, but can be minimized by controlling the tip-sample voltage difference.

Like the Hall probe, the MFM can be used over a wide range of temperature and applied

field. Also like the Hall probe, the MFM has worse noise performance than the SQUID, although this is a bit unfair. SQUIDs are among the lowest noise sensors available to modern science. The unique advantage of the MFM is the very small lengthscales accessible. It is also possible to manipulate vortices directly using the MFM, which opens up other experimental opportunities that are entirely irrelevant to us here. Another disadvantage of the MFM is the fear of crashing the tip into the sample, which usually ends the cooldown.

1.5.1 Vortex Profiles

The field from an isolated superconducting vortex is approximately that of a magnetic monopole located λ below the surface.[36–38] So one method to extract the penetration depth from local probe measurements is to image isolated vortices and determine λ from a fit parameter.[38] The fit is typically to line cuts through the data, rather than the entire scanned image. The fitted λ depends strongly on the scan height above the sample surface and other geometric parameters, including the probe geometry and scanner calibration constants. The offset height in particular is difficult to control accurately, and is the main source of systematic uncertainty in these measurements. It is difficult to properly account for this sort of uncertainty.

1.5.2 Touchdown Curves

In the second category of experiments, the probe is now a source of magnetic field. The probe is moved towards the sample, which generates an increasingly stronger response field due to the Meissner effect. The probe response during this process is called a touchdown curve. The experimenter can then extract λ from the touchdown curve by fitting it to a model of the probe source field and response function. The process is illustrated in Fig. 1.2, showing the movement of the probe towards the sample alongside the touchdown curve so generated.




Figure 1.2: Touchdown curves, shown in red, are slightly different for the SQUID (a) and MFM (b). The SQUID is only sensitive to magnetic response, and the signal increases until touchdown and is mostly constant as the sensor pushes past the touchdown point. The MFM responds to *all* forces, so the signal peaks then decreases rapidly due to Van der Waals forces.

Chapter 2

Calculating the Response Fields of Superconductors

In the last chapter I described several different local magnetic probes and the two techniques of using local probes to determine the penetration depth. Aside from a brief dalliance with the MFM, I spent my graduate career using scanning SQUIDs to measure the penetration depth, so the rest of this thesis will focus only on SQUID measurements. In this chapter I will work through the equations and model we use to analyze the measurements. Then I will spend some time considering the induced current distribution in thin films and what it can tell us about how the measurement might respond to inhomogeneity in the penetration depth.

2.1 SQUID Background

A SQUID, or Superconducting QUantum Interference Device, is a very sensitive detector of magnetic flux. Our devices are dc SQUIDs, so at the core there is a superconducting loop interrupted by two weak links. The loop also has leads, and is sketched in Fig. 2.1(a). The critical current through the loop (across the two leads) is periodic in the magnetic flux threading the loop. When the SQUID is driven above its critical current it is in the voltage state, and only time-dependent solutions to the Josephson equations are allowed.[2] These Josephson currents circulate through the SQUID loop and occur at microwave frequencies

(10 GHz). I'm not interested in them, although they can be a significant source of noise and heating in measurements of mesoscopic structures.[39] In the voltage state there is an average (dc) voltage across the two terminals of the SQUID, and it is periodic in the applied flux.[2] There are many more details to how SQUIDs really work;[40] I understand them at the level that a current-biased SQUID functions as a flux to voltage transducer.

But it can't even be that simple. A current-biased superconducting device is unstable against heating. The device goes normal, the resistance gets larger, which causes more heating, and perhaps you fry it. Much better to voltage bias the SQUID. Now if it goes normal, the current drops so there's less Joule heating, then the temperature goes down and the device returns to the superconducting state. We voltage bias our devices with an alkaline battery and a very small (0.1Ω) bypass resistor. It's still operating in the voltage state, but now the current through the SQUID is periodic in the applied flux. We inductively couple the SQUID bias current to a series array of 100 more SQUIDs, which serves as a low-noise preamplifier and impedance match to room temperature electronics.[41]

But wait, there's more! Our devices are optimized for scanning in several ways, illustrated in Fig. 2.1. The SQUID loop is gradiometric (crosses over itself), stretched out, and shielded everywhere except for the two pickup loops. In addition there is a pair of modulation coils for coupling flux into the SQUID loop. This allows us to operate in a flux-locked loop, linearizing the sensor.[40] Finally, there is a set of one-turn field coils around each pickup loop. A center tap allows precise nulling of the total flux into the SQUID loop due to the field coils.

To measure the penetration depth I applied an alternating current to the field coils, and used a lock-in amplifier to record the component of the SQUID signal caused by the excitation. Typically I would apply the current with a voltage source and a $1 \text{ k}\Omega$ resistor into the front field coil lead (at the scanning tip), float the center tap, and ground the rear lead. This doesn't give a perfect null signal, but it's a simple configuration and one that's impossible to change in some subtle way. Other folks in the lab prefer to use voltage-controlled bidirectional current sources with a tunable second output for the center tap. This gives a good null signal and can drive more current through the field coil (important for studying mesoscopic rings[39]), and as a current source gives more certainty in the excitation level

than my method. Although using the center tap introduces uncertainty in the current flowing around the rear field coil, since the center tap can either source or sink current. This could be an issue if configured differently — with the current source connected to the rear field coil, center tap connected, and the front field coil lead grounded. Which you might choose in order to have the SQUID corner nearest the sample “closer to ground”, if you’re concerned about electrostatic effects.

An alternate configuration that I found useful when studying bulk superconductors is to source current into the front field coil, float the rear lead, and ground the center tap. Then the signal is large far from the sample, but small close to it. So if I was going to park near the sample to measure $\Delta\lambda(T)$, I could use more gain.

The takeaway is that the SQUID driver has some flexibility in the offset of the susceptibility signal, by choosing to use or ignore the rear field coil and center tap. And, because the lockin has a phase setting, the driver can also change the sign of the signal. Me, I always preferred the data to go up when I got closer to the sample, so I kept the phase set to 180 degrees, meaning my measurements are of $\Phi/(-I)$. Really, a susceptibility signal due to diamagnetism should be negative. But then, our measurements aren’t *really* measurements of susceptibility, meaning M/H . The “susceptibility” measurement is a mutual inductance measurement. For example, with current only going through the front field coil, the measured signal is proportional to the mutual inductance between the front field coil and front pickup loop. With no sample nearby, the mutual inductance is a positive number, around 2 pH. With a strongly screening superconductor nearby, the mutual inductance is less, possibly even zero. In the balanced configuration, the measurement is of the mutual inductance between both field coils and the entire SQUID loop. Far from the sample, this is close to zero, because the SQUID loop is gradiometric. A superconductor near the front field coil *increases* the measured signal by canceling the flux into the front pickup loop. I think there’s an ambiguity here, whether the signal will be positive or negative.

To understand the measurement and interpret the data, we need to model the system. To do that, I consider an arbitrary magnetic field source above a superconducting slab of thickness t and infinite lateral extent. The superconductor has an isotropic penetration depth. In Section 2.2, I work through the solution of this problem, including the internal solution in the superconductor. Quantitative comparison with data requires a model of the




Figure 2.1: Summary of scanning SQUID devices. (a) Schematic of a dc SQUID, the crosses are Josephson junctions. (b) Schematic of scanning SQUIDs, showing the gradiometric SQUID loop (black), covered in Niobium shielding (gray) everywhere except the two pickup loops and a region near the modulation coils (green). Also shown are the field coils (blue). (c) Picture of a SQUID taken with an optical microscope. The rear pickup loop is roughly 1 mm from the scanning tip, and is usually not sensitive to the sample. (d) Another microscope image of the SQUID, at higher magnification, showing the front pickup loop and field coil.

field from the SQUID field coil and how the SQUID signal reacts to the superconducting response fields. The simplest model for this is the two-coil model, developed in Section 2.3. Most of these results have previously been derived in one form or another, a few examples are Refs. [42–46]. I’ll follow the formalism of Ref. [43].

How reasonable is this model? What assumptions have I made? By choosing to solve the London equation I’m assuming that the samples are in the local limit, with the superconducting coherence length much smaller than the penetration depth. The restriction to isotropic λ doesn’t necessarily exclude anisotropic systems; in the common scenario where the sample surface is the ab -plane and $\lambda_a = \lambda_b$, the c -axis penetration depth drops out and the response is identical to an isotropic superconductor with penetration depth λ_{ab} . [43]

Real samples are finite in extent, typically 500 μm on a side or larger. But the fields and currents involved in the measurement should be confined to the much smaller lengthscale of the SQUID, leaving the measurement unaware of the boundary. This assumption could

fail in the dot arrays studied in Chapter 5, where the effective penetration depth can be extremely large and the structures are only 120 μm on a side. I try to get a sense of this in Section 2.3.2, by calculating the supercurrents induced in a thin film by a circular field coil.

An unstated assumption is that the sample surface is smooth and flat. Roughness in the surface could cause (in an ab-oriented surface) λ_c to mix in as the screening current is forced to follow the topography. This is part of the “dead layer” discussed in Section 1.4.2. I would expect the SQUID to be less sensitive to any dead layer than μSR , just due to the very different lengthscales probed. More concerning for the SQUID are terraces and other vertical features. Not atomic terraces like you get from film growth on a miscut substrate. These are visible under the microscope, and quite common on flux-grown single crystals. This could be a way for λ_c to enter the measurement. But even estimating the effect of a vertical feature in an isotropic superconductor can only be approached numerically. So this problem will have to stay under the rug for now.

2.2 Meissner Response

The geometry I’m using is sketched in Fig. 2.2. The SQUID is in vacuum, at a height D above the sample surface. The SQUID is drawn as two circular wires, but for now is just the origin of an arbitrary source field. The sample is an isotropic superconductor, a slab of infinite lateral extent and width t . I’m going to take the top sample surface to be $z = 0$ and the region above the sample, where the SQUID is, is positive z . Below the sample, $z < -t$, is more vacuum.

In the region where $0 < z < D$, the field is given by:

$$\mathbf{h} = \nabla(\varphi_s + \varphi_r), \quad (2.1)$$




Figure 2.2: The geometry for our calculation. The SQUID is represented by two infinitely thin circular wires. The larger field coil (blue) is concentric and coplanar with a smaller pickup loop (red). The SQUID is in vacuum above a superconducting film of arbitrary thickness t .

where φ_s is the magnetic potential from the source and φ_r is the response from the superconductor. They are given in terms of their two-dimensional transforms as:

$$\varphi_s(\mathbf{r}, z) = \int \frac{d\mathbf{k}}{4\pi^2} \varphi_s(\mathbf{k}) e^{kz} e^{i\mathbf{k}\cdot\mathbf{r}} \quad (2.2)$$

$$\varphi_r(\mathbf{r}, z) = \int \frac{d\mathbf{k}}{4\pi^2} \varphi_r(\mathbf{k}) e^{-kz} e^{i\mathbf{k}\cdot\mathbf{r}}. \quad (2.3)$$

For $z < -t$, the field is given by:

$$\mathbf{h} = \nabla \varphi_t, \quad (2.4)$$

and φ_t is given by:

$$\varphi_t(\mathbf{r}, z) = \int \frac{d\mathbf{k}}{4\pi^2} \varphi_t(\mathbf{k}) e^{kz} e^{i\mathbf{k}\cdot\mathbf{r}}. \quad (2.5)$$

In the region $-t < z < 0$, the field is a linear combination of both possible solutions to the London equation:

$$\mathbf{h} = \int \frac{d\mathbf{k}}{4\pi^2} \left[\mathbf{h}_1(\mathbf{k}) e^{Qz} + \mathbf{h}_2(\mathbf{k}) e^{-Qz} \right] e^{i\mathbf{k}\cdot\mathbf{r}}, \quad (2.6)$$

where $Q^2 = k^2 + \lambda^{-2}$. The goal is to determine φ_r in terms of φ_s . There are 8 unknowns - φ_r , φ_t , and three components for each of \mathbf{h}_1 and \mathbf{h}_2 . We require continuity of the field at both boundaries, which gives us six equations. The final two equations come from requiring that \mathbf{h}_1 and \mathbf{h}_2 are both divergenceless. We require this of each field separately, since they're both valid magnetic fields.

Continuity at the $z = 0$ interface requires:

$$\begin{aligned} h_{1x} + h_{2x} &= ik_x(\varphi_s + \varphi_r) \\ h_{1y} + h_{2y} &= ik_y(\varphi_s + \varphi_r) \\ h_{1z} + h_{2z} &= k(\varphi_s - \varphi_r) \end{aligned} \quad (2.7)$$

Continuity at $z = -t$ gives:

$$\begin{aligned} h_{1x}e^{-Qt} + h_{2x}e^{Qt} &= ik_x\varphi_t \\ h_{1y}e^{-Qt} + h_{2y}e^{Qt} &= ik_y\varphi_t \\ h_{1z}e^{-Qt} + h_{2z}e^{Qt} &= k\varphi_t \end{aligned} \quad (2.8)$$

And the two divergenceless conditions gives:

$$ik_x h_{1x} + ik_y h_{1y} + Q h_{1z} = 0 \quad (2.9)$$

$$ik_x h_{2x} + ik_y h_{2y} - Q h_{2z} = 0. \quad (2.10)$$

I found it easiest to solve the system with a matrix notation. Define:

$$\boldsymbol{\varphi} = \begin{bmatrix} \varphi_s \\ \varphi_r \\ \varphi_t \end{bmatrix}. \quad (2.11)$$

We can write the equations for the $z = 0$ interface as:

$$\mathbf{h}_1 + \mathbf{h}_2 = \begin{bmatrix} ik_x & ik_x & 0 \\ ik_y & ik_y & 0 \\ k & -k & 0 \end{bmatrix} \boldsymbol{\varphi} \equiv M_0 \boldsymbol{\varphi}. \quad (2.12)$$

Similarly for the $z = -t$ interface:

$$e^{-Qt} \mathbf{h}_1 + e^{Qt} \mathbf{h}_2 = \begin{bmatrix} 0 & 0 & ik_x \\ 0 & 0 & ik_y \\ 0 & 0 & k \end{bmatrix} \boldsymbol{\varphi} \equiv M_t \boldsymbol{\varphi}. \quad (2.13)$$

We will also define the row vectors A and B :

$$A = [ik_x \quad ik_y \quad Q] \quad (2.14)$$

$$B = [ik_x \quad ik_y \quad -Q], \quad (2.15)$$

so that the divergenceless conditions become:

$$A\mathbf{h}_1 = 0 \quad (2.16)$$

$$B\mathbf{h}_2 = 0. \quad (2.17)$$

Now we can combine Eqs. 2.12, 2.13, 2.16, and 2.17 to get the equations:

$$AM_0 \boldsymbol{\varphi} = e^{-Qt} AM_t \boldsymbol{\varphi} \quad (2.18)$$

$$BM_0 \boldsymbol{\varphi} = e^{Qt} BM_t \boldsymbol{\varphi}. \quad (2.19)$$

Once we evaluate the matrix products, Eq. 2.18 becomes:

$$(Q - k)\varphi_s - (Q + k)\varphi_r = e^{-Qt}(Q - k)\varphi_t. \quad (2.20)$$

Similarly, Eq. 2.19 becomes:

$$-(Q + k)\varphi_s + (Q - k)\varphi_r = -e^{Qt}(Q + k)\varphi_t. \quad (2.21)$$

We can solve the system (Eqs. 2.20 and 2.21) to get:

$$\varphi_r = \varphi_s \frac{(Q^2 - k^2) \sinh(Qt)}{(Q^2 + k^2) \sinh(Qt) + 2kQ \cosh(Qt)} \quad (2.22)$$

$$\varphi_t = \varphi_s \frac{2kQ}{(Q^2 + k^2) \sinh(Qt) + 2kQ \cosh(Qt)}. \quad (2.23)$$

For convenience later, I'm going to define two kernel functions:

$$\mathcal{K}_r \equiv \frac{\varphi_r}{\varphi_s}, \quad \text{and} \quad \mathcal{K}_t \equiv \frac{\varphi_t}{\varphi_s}.$$

We're mostly concerned with the thin film ($t \rightarrow 0$) and bulk ($t \rightarrow \infty$) limits of the response field. In the thin film limit:

$$\begin{aligned} \mathcal{K}_r &\approx \frac{(Q^2 - k^2)Qt}{(Q^2 + k^2)Qt + 2kQ} \\ &= \frac{\lambda^{-2}Qt}{(2k^2 + \lambda^{-2})Qt + 2kQ} \\ &= \frac{1}{2k^2\lambda^2 + 1 + 2k\lambda^2t^{-1}} \\ &= \frac{1}{1 + \Lambda k + \Lambda tk^2} \approx \frac{1}{1 + \Lambda k}. \end{aligned} \quad (2.24)$$

Here, I've introduced the Pearl length, $\Lambda = \frac{2\lambda^2}{t}$. In the opposite limit of a superconducting

half space we have:

$$\begin{aligned}
\mathcal{K}_r &\approx \frac{(Q^2 - k^2)e^{Qt}}{(Q^2 + k^2)e^{Qt} + 2kQe^{Qt}} \\
&= \frac{Q^2 - k^2}{(Q + k)^2} \\
&= \frac{Q - k}{Q + k} = \frac{\sqrt{1 + \lambda^2 k^2} - \lambda k}{\sqrt{1 + \lambda^2 k^2} + \lambda k}.
\end{aligned} \tag{2.25}$$

Since I'll need it later (Section 2.3.2), the thin film limit of the transmitted response is:

$$\begin{aligned}
\mathcal{K}_t &\approx \frac{2kQ}{(Q^2 - k^2)Qt + 2kQ} \\
&= \frac{2kQ}{(2k^2 + \lambda^{-2})Qt + 2kQ} \\
&= \frac{1}{1 + (2k^2\lambda^{-2})t/2k} \\
&= \frac{\Lambda k}{1 + \Lambda k + \Lambda t k^2} \approx \frac{\Lambda k}{1 + \Lambda k}.
\end{aligned} \tag{2.26}$$

The bulk limit of \mathcal{K}_t is zero, naturally.

With Eq. 2.22 in hand, we can determine the superconducting response field arising from an arbitrary local source field. We still need to relate this to the experimental response. The obvious way forward is to make some reasonable assumption about the geometry of the probe. For a SQUID, a reasonable starting point is the two-coil model, where the pickup loop and field coil are infinitely thin circular coils. The field coil creates a source field, and the measured response is given by the flux of \mathbf{h}_z through the pickup loop. This describes the data reasonably well, so we'll develop it further in Section 2.3.

The drawback here is that this is an approximation to the true behavior. Worse, it is an uncontrolled approximation — we have no *a priori* way to know the effect any given model geometry will have on the inferred penetration depth. We may be forced to consider a set of geometrical models, a seriously disheartening proposition. It might be worth looking at what can be said without reference to the detailed geometry of the scanning probe. This could mean analysis of the response kernel \mathcal{K}_r in real space, because the form of Eq. 2.22

tells us that the (real space) response field is a two-dimensional convolution of the source field and the response kernel. Specifically, you can show that:

$$\mathbf{h}_r \times \hat{\mathbf{z}} = \mathcal{K}_r \circ (\mathbf{h}_s \times \hat{\mathbf{z}}) \quad (2.27)$$

$$\mathbf{h}_r \cdot \hat{\mathbf{z}} = -\mathcal{K}_r \circ (\mathbf{h}_s \cdot \hat{\mathbf{z}}). \quad (2.28)$$

An artificial example would be if in some limit \mathcal{K}_r is only a function of the penetration depth and the sensor height. Then the response field at the sensor would be proportional to the source field with the z -component flipped. Then the signal is proportional to this function of λ and z times minus the mutual inductance between the front field coil and the pickup loop. This assumes that the SQUID response is linear in the response field, so it can't account for flux focusing due to the diamagnetism of the niobium pickup loop.

2.2.1 Internal Solution

We can also solve for the magnetic field inside the superconductor, which lets us calculate the screening currents. This might be useful when we try to understand the spatial resolution of local penetration depth measurements, or what the response to a non-uniform penetration depth might be. For example, what effect does the finite sample size have in various limits (thin film, bulk, weak screening, strong screening)?

Anyhow, Eqs. 2.22 and 2.23 give the reflected and transmitted potentials as proportional to the source potential φ_s . Then we can use the system defined by Eqs. 2.12 and 2.13:

$$\mathbf{h}_1 + \mathbf{h}_2 = M_0 \varphi \quad (2.29)$$

$$e^{-Qt} \mathbf{h}_1 + e^{Qt} \mathbf{h}_2 = M_t \varphi, \quad (2.30)$$

to solve for the internal field, since φ is known. The internal magnetic field is related to \mathbf{h}_1 and \mathbf{h}_2 via Eq. 2.6. Solving the system gives us:

$$\mathbf{h}_1 = \frac{1}{2 \sinh(Qt)} (M_0 e^{Qt} - M_t) \varphi, \quad (2.31)$$

$$\mathbf{h}_2 = \frac{1}{2 \sinh(Qt)} (M_t - M_0 e^{-Qt}) \varphi. \quad (2.32)$$

Of course, we aren't actually interested in \mathbf{h}_1 or \mathbf{h}_2 , but the total field:

$$\mathbf{h}(\mathbf{k}) = \mathbf{h}_1 e^{Qz} + \mathbf{h}_2 e^{-Qz}. \quad (2.33)$$

Plugging Eqs. 2.31 and 2.32 into Eq. 2.33 and simplifying gives us:

$$\mathbf{h}(\mathbf{k}) = \left[\frac{\sinh[Q(z+t)]}{\sinh(Qt)} M_0 - \frac{\sinh(Qz)}{\sinh(Qt)} M_t \right] \boldsymbol{\varphi}. \quad (2.34)$$

This is going to get worse before it gets better. Using the definitions of M_0 and M_t (Eqs. 2.12 and 2.13) we have:

$$M_0 \boldsymbol{\varphi} = \begin{bmatrix} ik_x(\varphi_s + \varphi_r) \\ ik_y(\varphi_s + \varphi_r) \\ k(\varphi_s - \varphi_r) \end{bmatrix}, \quad (2.35)$$

and:

$$M_t \boldsymbol{\varphi} = \begin{bmatrix} ik_x \varphi_t \\ ik_y \varphi_t \\ k \varphi_t \end{bmatrix}. \quad (2.36)$$

Collecting everything and simplifying, we get:

$$\mathbf{h}(\mathbf{k}) = \left\{ f_{xy}(k, \lambda, t, z) [ik_x \hat{\mathbf{x}} + ik_y \hat{\mathbf{y}}] + f_z(k, \lambda, t, z) k \hat{\mathbf{z}} \right\} \boldsymbol{\varphi}_s(\mathbf{k}), \quad (2.37)$$

with f_{xy} and f_z defined by:

$$f_{xy} = \frac{Q^2 [\cosh[Q(z+2t)] - \cosh(Qz)] + kQ [\sinh[Q(z+2t)] - \sinh(Qz)]}{\frac{1}{2}(Q^2 + k^2)(\cosh(2Qt) - 1) + kQ \sinh(2Qt)}, \quad (2.38)$$

and:

$$f_z = \frac{k^2 [\cosh[Q(z+2t)] - \cosh(Qz)] + kQ [\sinh[Q(z+2t)] - \sinh(Qz)]}{\frac{1}{2}(Q^2 + k^2)(\cosh(2Qt) - 1) + kQ \sinh(2Qt)}. \quad (2.39)$$

Recall that $Q^2 = k^2 + \lambda^{-2}$ and $k = \sqrt{k_x^2 + k_y^2}$. From here, we can calculate the screening

currents from $\mu_0 \mathbf{J} = \nabla \times \mathbf{h}$. This won't work for the thin film limit, where there is no internal solution; the superconductor simply provides a boundary condition for the field. The bulk limit of Eqs. 2.38 and 2.39 is:

$$f_{xy} \rightarrow \frac{2Q}{Q+k} e^{Qz} \quad (2.40)$$

$$f_z \rightarrow \frac{2k}{Q+k} e^{Qz}. \quad (2.41)$$

2.3 Two-Coil Model

In the two-coil model the SQUID is reduced to two coaxial and coplanar circles, like in Fig. 2.2. The outer circle represents the field coil, and it carries current I and has no thickness. The inner circle represents the pickup loop. The SQUID signal is the total flux of the z -component of the response field inside this area. That is, I'm leaving out the flux directly coupled into the pickup loop by the field coil, it's been canceled out by the back pickup loop so that the signal is 0 far from the sample.

Since we're approximating the field coil as an infinitely thin circular current of radius a at a height D above the sample surface, the source potential is given by:

$$\varphi_s(k) = \frac{\mu_0 \pi I a}{k} e^{-kD} J_1(ka), \quad (2.42)$$

where J_1 is a Bessel function of the first kind, I is the current in the field coil, and μ_0 is the vacuum permeability. The SQUID signal is the flux of the z -component of the response field through the pickup loop, which has radius p , or:

$$\Phi = 2\pi \int_0^p \mathbf{d}r r \partial_z \varphi_r(\mathbf{r}, z) \Big|_{z=D}. \quad (2.43)$$

Inserting Eq. 2.3 into Eq. 2.43, this becomes:

$$\Phi = -2\pi \int_0^p r \mathbf{d}r \int \frac{\mathbf{d}\mathbf{k}}{4\pi^2} \varphi_r(\mathbf{k}) k e^{-kD} e^{i\mathbf{k}\cdot\mathbf{r}}. \quad (2.44)$$

Everything is cylindrically symmetric, so we can integrate over the orientation of k as:

$$\int_0^{2\pi} d\varphi e^{ikr \cos \varphi} = 2\pi J_0(kr). \quad (2.45)$$

Now the SQUID signal is:

$$\Phi = - \int_0^p r dr \int_0^\infty k^2 dk \mathcal{K}_r \varphi_s(k) J_0(kr) e^{-kD}. \quad (2.46)$$

Using the identity

$$\int_0^u u' J_0(u') du' = u J_1(u), \quad (2.47)$$

and inserting Eq. 2.42 this simplifies to:

$$\Phi = -\mu_0 \pi I a p \int_0^\infty dk \mathcal{K}_r J_1(ka) J_1(kp) e^{-2kD}, \quad (2.48)$$

where \mathcal{K}_r is the response kernel defined by Eq. 2.22. In Fig. 2.3 I plot the SQUID signal as a function of height above the sample (D) in the bulk and thin film limits at a few values of the relevant screening length (λ or Λ). The pickup loop radius (a) is 7.58 μm , and the field coil radius (p) is 2.55 μm . These values come from fitting sets of touchdown curves, and don't perfectly match the actual geometry of the SQUID. The curves are normalized by the mutual inductance between the field coil and pickup loop, $\mathcal{M}(a, p, 0)$, with \mathcal{M} is defined by Eq. 2.49.

2.3.1 Useful Limits of the SQUID Signal

I'm going to define the function $\mathcal{M}(a, p, h)$ to be the mutual inductance between coaxial rings of radius a and p , separated by a vertical height h . I can look up in Jackson the vector potential of a circular current-carrying loop, it's on page 182 in my copy.[47] From there it's pretty easy to determine that

$$\mathcal{M}(a, p, h) = \mu_0 \frac{\sqrt{ap}}{m} \left[(2 - m^2)K(m) - 2E(m) \right], \quad (2.49)$$

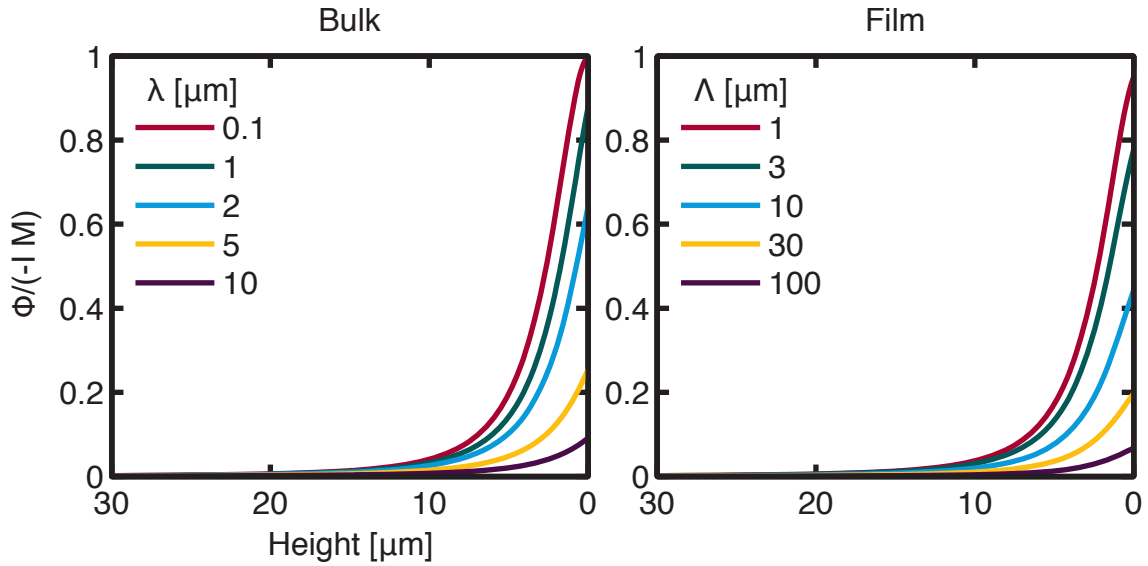


Figure 2.3: Calculated touchdown curves in the two-coil model, for the bulk (left) and thin film (right) limiting cases. The curves use the best-fit values of the pickup loop and field coil radii, 7.58 and 2.55 μm respectively. The normalization (M) is the mutual inductance between field coil and pickup loop.

where K , and E are the complete elliptic integrals of the first and second kind, respectively. The elliptic modulus m is given by:

$$m^2 = \frac{4ap}{h^2 + (a+p)^2}. \quad (2.50)$$

This seems like a good time to mention the integral relation

$$\int_0^\infty dx e^{-xh} J_1(xa) J_1(xp) = \frac{1}{\pi m \sqrt{ap}} \left[(2 - m^2) K(m) - 2E(m) \right], \quad (2.51)$$

which is equation 2.5.4 in Ref. [48]. I distinctly recall a homework problem where you prove that the mutual inductance of two coaxial rings had a form with the integral of two Bessel functions. I've always wondered how to prove the equivalence of the two forms for the mutual inductance without making reference to circular current loops. The best I could do was find a reference.

Why am I telling you this? There are a few useful limits of Eq. 2.48 that can be expressed in terms of \mathcal{M} , and complete elliptic integrals can be evaluated using the arithmetic-geometric mean, which is *much* faster than numerical quadrature of an oscillatory integrand. The simplest case is when the penetration depth is zero, then \mathcal{K}_r is one, and the SQUID signal is:

$$\Phi = -I\mathcal{M}(a, p, 2D). \quad (2.52)$$

In this case, the SQUID signal is from a time-reversed image field coil, with the image plane at the sample surface.

The second case is a bulk sample with small penetration depth. Specifically, when λ is small compared to the SQUID height D , the exponential term limits weight in the integral to wavevectors that satisfy $k\lambda \ll 1$, so we can expand \mathcal{K}_r :

$$\mathcal{K}_r = \frac{Q - k}{Q + k} \approx e^{-2k\lambda} \left(1 + \mathcal{O}\left(\frac{\lambda}{h}\right)^3\right). \quad (2.53)$$

Then over a bulk sample, when the SQUID height is large compared to the penetration depth, the signal is a modified image problem:

$$\Phi = -I\mathcal{M}(a, p, 2(D + \lambda)), \quad (2.54)$$

where the mirror plane is λ below the surface. For most bulk superconductors, the penetration depth is at most a few hundred nanometers. And the center of the pickup loop is roughly one micron above the sample surface even when the SQUID tip is touching the sample, so we are almost always in this limit. A nice feature of this limit is that height changes are equivalent to penetration depth changes, so it is possible to determine $\Delta\lambda$ without reference to the SQUID geometry; only the z-bender calibration enters.

The last relation I want to mention is in the thin film limit, Eq. 2.24. Using the Laplace transform pair

$$\frac{1}{1 + k\Lambda} = \frac{1}{\Lambda} \int_0^\infty e^{-t/\Lambda} e^{-kt} dt,$$

you can re-write the SQUID response as:

$$\Phi = \frac{-I}{\Lambda} \int_0^\infty dt e^{-t/\Lambda} \mathcal{M}(a, p, t + 2D). \quad (2.55)$$

The advantage to this is that \mathcal{M} is a monotonically decreasing function of h , so numerical integration is faster and more accurate.

2.3.2 Supercurrents in a Film

I want to investigate the distribution of screening currents in the SQUID measurement, to better understand where the assumptions break down in terms of sample extent and what might be the imaging kernel of the susceptibility measurement. I'll work it out in the case of the thin film, since it has one less variable (depth) to plot and worry about. Maxwell's equations tell us that the screening current in a thin film will be determined by the discontinuity in the transverse component of the field, or:

$$\hat{z} \times (\mathbf{h}_+ - \mathbf{h}_-) = \mu_0 \mathbf{J}, \quad (2.56)$$

where \mathbf{h}_\pm indicates the field above and below the film:

$$\mathbf{h}_+ = \nabla(\varphi_r + \varphi_s), \quad (2.57)$$

$$\mathbf{h}_- = \nabla\varphi_t. \quad (2.58)$$

The cylindrical symmetry guarantees that there is only an azimuthal component to the current, and that we can use Eq. 2.45 to get:

$$\mu_0 J_\phi = - \int_0^\infty \frac{k^2 dk}{2\pi} J_1(kr) \left[\frac{2\varphi_s(k)}{1 + k\Lambda} \right], \quad (2.59)$$

which with Eq. 2.42 for φ_s gives:

$$J_\phi = -Ia \int_0^\infty \frac{k dk}{1 + k\Lambda} J_1(kr) J_1(ka) e^{-kD}. \quad (2.60)$$

It's going to help to make things unitless, by defining all lengths relative to the field coil radius. Defining $z = ka$, $d = D/a$, $l = \Lambda/a$, and $x = r/a$, Eq. 2.60 becomes:

$$J_\phi = \frac{-I}{a} \int_0^\infty \frac{z dz}{1 + zl} J_1(xz) J_1(z) e^{-zd}. \quad (2.61)$$




Figure 2.4: The supercurrent density (top row) and cumulative supercurrent (bottom row) as a function of radius. Each column is a different sensor height, normalized by the field coil radius. Colors correspond to different values of the Pearl length.

The upper row of Fig. 2.4 shows the supercurrent density, normalized by the total supercurrent, as a function of radius for three different heights (columns) and five values of the Pearl length (curves). The total induced supercurrent can be found by integrating Eq. 2.61 over x :

$$I_{sc} = -I \int_0^\infty \frac{dz}{1 + zl} J_1(z) e^{-zd}. \quad (2.62)$$

The lower row of Fig. 2.4 shows the cumulative current as a function of radius. As the name

implies, the cumulative current is the integral of the supercurrent density:

$$I_{\text{cum}}(R) = \int_0^R J_\phi dr, \quad (2.63)$$

$$= -I \int_0^\infty \frac{z dz}{1 + z l} J_1(z) e^{-zd} \int_0^X J_1(xz) dx, \quad (2.64)$$

$$= -I \int_0^\infty \frac{dz}{1 + z l} [1 - J_0(Xz)] J_1(z) e^{-zd}. \quad (2.65)$$

where $X = R/a$. To me, the biggest take-away from Fig. 2.4 is that the response is much more a function of SQUID height than it is of the penetration depth. At $d = 0.1$, slightly lower than the typical offset height, the supercurrent is sharply peaked at the field coil radius, no matter what the Pearl length is. At $d = 1$ the picture is similar, although the tail of the current distribution is more prominent and the maximum has moved out slightly. At $d = 3$ the maximum current density is more clearly a function of the Pearl length.




Figure 2.5: The f -critical radius as a function of reduced Pearl length, when $f = 0.5$ (left) and $f = 0.9$ (right). Colors correspond to different sensor heights.

As mentioned, there are two things I'd like to learn about by solving for the current distribution. The first is to understand when the finite size of the sample might become

important. The second is the imaging kernel of the SQUID susceptibility measurement, or the way in which the measurement will average an inhomogeneous penetration depth. We can address both of these by defining the f -critical radius, R_f . This is the radius at which $I_{\text{cum}}(R_f) = f I_{\text{sc}}$. So when $f = 0.5$, this is the radius that encloses half of the total supercurrent. Finite size effects can be addressed by considering f close to 1, because if the sample boundary is in a region with a negligible amount of current flow it should be a small perturbation. In Fig. 2.5 I consider $f = 0.9$ (right graph). The imaging kernel will be influenced by the spatial extent of the induced supercurrents. For this we can consider either the peak in the supercurrent density or the fifty percent f -critical radius, shown in the left graph of Fig. 2.5. This shows that sensor heights less than roughly $7 \mu\text{m}$, the lengthscale of the supercurrent response is the field coil radius and is essentially independent of Λ . Even at $d = 3$, $R_{0.5}$ is less than 3 at small Pearl length and slowly grows to 5 at $\Lambda/a = 100$.




Figure 2.6: The differential flux (top row) and cumulative flux (bottom row) as a function of radius. Each column is a different sensor height, normalized by the field coil radius. Colors correspond to different values of the Pearl length.

So that's the currents. What about the SQUID signal? Well, in the two coil model the signal is the flux through the pickup loop. So I can define the differential flux or SQUID signal as the flux through the pickup loop due to the current flowing at radius r :

$$d\Phi = J_\phi(r)\mathcal{M}(p, r, D). \quad (2.66)$$

There's little point writing out the equations for this, since it is the product of two improper integrals. Of course, it makes sense to normalize by the total SQUID signal Φ . The importance of the differential SQUID signal is that while there can be screening currents flowing over a large area, the sensor is not sensitive to them. By considering the pickup loop I've introduced another length; I'll use the same values for a and p as before, 7.58 and 2.55 μm respectively. The differential and cumulative SQUID signal are shown in Fig. 2.6, in exact analogy to Fig. 2.4. I've also calculated f -critical radii, shown in Fig. 2.7, just like Fig. 2.5.

It turns out that the f -critical flux radii are both smaller and larger than the corresponding values for the current. At small Pearl length and low height the flux $R_{0.5}$ is smaller than the current $R_{0.5}$. But at larger height and small Pearl length the flux value is larger than the current value. At larger Pearl length the flux radii are always smaller than the corresponding current values. The flux radii saturate at large values of the Pearl length, in contrast to the currents, which spread out logarithmically with increasing Pearl length. This was expected, since currents flowing far from the pickup loop are not effective at coupling flux into the SQUID. In Fig. 2.8 I've plotted both the flux and current R_f to directly compare the two. The figure's a bit busy, but the crossover is very clear for $R_{0.5}$.

Another fun feature of the differential flux is that at small heights there is a shoulder corresponding to the pickup loop radius. The shoulder seems to not depend on the Pearl length, except when $\Lambda/a = 0.1$. What seems to be happening is that at short Pearl length, the supercurrents are strongly localized at the field coil radius, so the measurement kernel is the size of the field coil. As the Pearl length increases, the supercurrent density spreads out and the coupling to the pickup loop becomes more important. Surprisingly, the resolution at low heights gets *better* as the Pearl length increases. At larger height the resolution ($R_{0.5}$) grows with Λ as expected.




Figure 2.7: The f -critical flux radius as a function of reduced Pearl length, when $f = 0.5$ (left) and $f = 0.9$ (right). Colors correspond to different sensor heights.




Figure 2.8: The f -critical flux (dashed) and current (solid) radius as a function of reduced Pearl length, when $f = 0.5$ (left) and $f = 0.9$ (right). Colors correspond to different sensor heights.

Chapter 3

Calculation of the Effect of Random Superfluid Density on the Temperature Dependence of the Penetration Depth

Thomas M. Lippman, Kathryn A. Moler,

Calculation of the effect of random superfluid density on the temperature dependence of the penetration depth.

Physical Review B **85**, 104529 (2012).

©2012 American Physical Society.

Abstract

Microscopic variations in composition or structure can lead to nanoscale inhomogeneity in superconducting properties such as the magnetic penetration depth, but measurements of these properties are usually made on longer length scales. We solve a generalized London equation with a non-uniform penetration depth, $\lambda(\mathbf{r})$, obtaining an approximate solution for the disorder-averaged Meissner screening. We find that the effective penetration depth is different from the average penetration depth and is sensitive to the details of the disorder. These results indicate the need for caution when interpreting measurements of the penetration depth and its temperature dependence in systems which may be inhomogeneous.

3.1 Introduction

The penetration depth and its temperature dependence are important characteristics of any superconductor and are considered key to determining the momentum space structure of the order parameter.[6, 7, 10] The possibility of disorder in exotic superconductors is well known, but analyses performed to date have concentrated on the effect of disorder-induced scattering on the momentum space structure of the gap.[49–52] This paper is motivated by the possibility that disorder may lead to nanoscale real space variation and the associated need to model the relationship between such spatial variation and properties that are measured on longer length scales. We address how inhomogeneity in the penetration depth may affect bulk measurements of the penetration depth for methods that rely on Meissner screening and can be analyzed by solutions to London's equation. In particular, we show that the measured result is not simply given by the average value of the penetration depth, but is affected by the statistical structure of the spatial variations in the penetration depth.

Many superconductors are created by chemical doping of a non-superconducting parent compound. In these systems the inherent randomness of the doping process may give rise to an inhomogeneous superconducting state. The importance of this effect will be determined by the characteristic length over which the dopant atoms affect the superconductivity. Even in the most ordered material, there will be binomial fluctuations in the total number of dopants in a given region. In general, one does not expect significant spatial variation in materials that are weakly correlated and can be described by a rigid band model. For example, disorder is largely irrelevant in classic metallic superconductors, due to their long coherence lengths and weakly correlated nature.[53] In contrast, the cuprates are doped insulators with a coherence length on the scale of the lattice. They are known to have nanoscale disorder in their superconducting properties, as seen by scanning tunneling microscopy.[54] Similar gap maps have been observed in the iron pnictide family[55–57] and in disordered titanium nitride films close to the superconductor to insulator transition.[58, 59]

Materials with intrinsic disorder present two separate challenges. Understanding how the random doping process gives rise to local superconducting properties, such as the penetration depth or local density of states, requires a microscopic model. But even with such a

model, we still need to make the connection between the local superconducting properties and bulk measurements. The manner in which local superconducting properties relate to the observed properties will differ from experiment to experiment. For instance, a measurement of the heat capacity will return the total heat capacity of the macroscopic sample, so the inferred specific heat capacity will be a volume average over the sample. In contrast, we might expect the thermal conductivity response to be dominated by a percolation path connecting regions with small local gap, $\Delta(\mathbf{r})$, or large local density of states.

Here we focus on the penetration depth, λ , as measured by screening of the magnetic field, including both resonant cavity frequency shift measurements at radio frequencies[10] and the local probes of Magnetic Force Microscopy[60] and Scanning SQUID Susceptometry.[61] These methods measure $\lambda(T)$ by detecting the response magnetic field generated by the superconductor due to an applied field, and can be analyzed using the London equation. Thus, we can model the effect of inhomogeneity by solving the London equation with $\lambda(\mathbf{r})$ as a random function of position \mathbf{r} . Our goal is to find a new equation for the disorder-averaged magnetic field, as this will determine the measured response. Here, we work in the limit of small fluctuations to find an approximate equation for the disorder-averaged magnetic field, as this will determine the measured response.

3.2 Stochastic London Equation

To understand the measured penetration depth when $\lambda(\mathbf{r})$ is a random function of position, we calculate the disorder-averaged magnetic field response to obtain an effective penetration depth. For isotropic and local superconductors in three dimensions, the static magnetic field $\mathbf{h}(\mathbf{r})$ is given by the London equation with $\lambda(\mathbf{r})$ a function of position. The correct form[62] of the London equation when the penetration depth is non-uniform is:

$$\mathbf{h} + \nabla \times [\lambda(\mathbf{r})^2 \nabla \times \mathbf{h}] = 0, \quad (3.1)$$

which is derived from the second Ginzburg–Landau equation in the London limit.¹ We parametrize the penetration depth as the average value plus a fluctuating term:

$$\lambda(\mathbf{r}) = \Lambda [1 + \xi(\mathbf{r})], \quad (3.2)$$

so that $\langle \lambda(\mathbf{r}) \rangle = \Lambda$. Then Eq. 3.1 becomes:

$$(L + M_1 + M_2) \mathbf{h} = 0, \quad (3.3)$$

where:

$$\begin{aligned} L &\equiv 1 - \Lambda^2 \nabla^2 \mathbf{I}, \\ M_1 &\equiv -2\Lambda^2 \xi \nabla^2 \mathbf{I} + 2\Lambda^2 \nabla \xi \times \nabla \times, \quad \text{and} \\ M_2 &\equiv -\Lambda^2 \xi^2 \nabla^2 \mathbf{I} + \Lambda^2 \nabla \xi^2 \times \nabla \times. \end{aligned} \quad (3.4)$$

Here \mathbf{I} is the identity tensor, and the “dangling curl” is understood to operate on a vector to its right. The terms are grouped so that M_1 is first-order in ξ , M_2 is second-order in ξ , and L gives the unperturbed London equation. We will work in the limit of small fluctuations, $\xi(\mathbf{r}) \ll 1$, so that $M_1 + M_2$ is a perturbative term in Eq. 3.3.

Our method of solution comes from the similarity of the Helmholtz and London equations. The Helmholtz equation, which governs wave propagation, becomes the London equation when the wavevector is purely imaginary. Thus our problem is related to the propagation of waves in a random medium, and we can build upon a large and multidisciplinary literature devoted to this challenge.[63, 64] The paper by Karal and Keller [65] is particularly relevant, because it retains the vectorial nature of the problem, rather than simplifying to a scalar wave equation.

We now derive, from Eq. 3.3, an approximate equation for the disorder-averaged field

¹The familiar relation $\mathbf{J} = -\mathbf{A}/\lambda^2$ is not valid when the order parameter is non-uniform. Instead, we start with the Ginzburg–Landau equations, which can describe a non-uniform superconducting state. From the second Ginzburg–Landau equation we derive our starting point, Eq. 3.1, a generalized London equation.

$\langle \mathbf{h} \rangle$. Applying the inverse of L to both sides:

$$[1 + L^{-1}(M_1 + M_2)] \mathbf{h} = \mathbf{h}_0, \quad (3.5)$$

where $L \mathbf{h}_0 = 0$. Solving for \mathbf{h} :

$$\mathbf{h} = [1 + L^{-1}(M_1 + M_2)]^{-1} \mathbf{h}_0, \quad (3.6)$$

assuming the inverse exists. Averaging both sides:

$$\langle \mathbf{h} \rangle = \left\langle [1 + L^{-1}(M_1 + M_2)]^{-1} \right\rangle \mathbf{h}_0, \quad (3.7)$$

where \mathbf{h}_0 comes outside of the average because it is non-random. Solving for \mathbf{h}_0 :

$$\left\langle [1 + L^{-1}(M_1 + M_2)]^{-1} \right\rangle^{-1} \langle \mathbf{h} \rangle = \mathbf{h}_0. \quad (3.8)$$

Since we assume small fluctuations, we can expand the term inside the average:

$$\langle 1 - L^{-1}(M_1 + M_2) + L^{-1}M_1L^{-1}M_1 + \mathcal{O}(\xi^3) \rangle^{-1} \langle \mathbf{h} \rangle = \mathbf{h}_0. \quad (3.9)$$

Averaging and expanding again:

$$(1 - L^{-1} \langle M_1L^{-1}M_1 \rangle + L^{-1} \langle M_2 \rangle) \langle \mathbf{h} \rangle = \mathbf{h}_0, \quad (3.10)$$

since $\langle M_1 \rangle = 0$ due to Eq. 3.2. We then apply L to both sides:

$$(L - \langle M_1L^{-1}M_1 \rangle + \langle M_2 \rangle) \langle \mathbf{h} \rangle = 0, \quad (3.11)$$

which yields the average field to second order in ξ .

3.3 Results

We first evaluate the averages in Eq. 3.11, giving us an equation for $\langle \mathbf{h} \rangle$ in terms of the penetration depth correlation function, $\langle \lambda(\mathbf{r})\lambda(\mathbf{r}') \rangle$. We then consider two specific cases for the correlation function and numerically evaluate the effective penetration depth for a range of parameters.

3.3.1 Evaluating the Averages

We will solve Eq. 3.11 for a single Fourier mode of $\langle \mathbf{h}(\mathbf{r}) \rangle = \mathbf{h} e^{i\mathbf{k}\cdot\mathbf{r}}$, then derive an equation for \mathbf{k} that yields exponentially decaying solutions consistent with Meissner screening.

First, we evaluate $\langle M_2 \rangle$:

$$\langle M_2 \rangle = -\Lambda^2 \langle \xi(\mathbf{r})^2 \rangle \nabla^2 I + \Lambda^2 \langle \nabla \xi(\mathbf{r})^2 \rangle \times \nabla \times. \quad (3.12)$$

We introduce the correlation function $R(\mathbf{r}, \mathbf{r}') = \langle \xi(\mathbf{r})\xi(\mathbf{r}') \rangle$, which is a function only of $|\mathbf{r} - \mathbf{r}'|$ if ξ is stationary and isotropic. Then we see that $\langle \xi^2 \rangle = R(0)$ and $\langle \nabla \xi^2 \rangle = \nabla \langle \xi^2 \rangle = 0$, so

$$\langle M_2 \rangle \langle \mathbf{h} \rangle = \Lambda^2 k^2 R(0) \mathbf{h} e^{i\mathbf{k}\cdot\mathbf{r}}. \quad (3.13)$$

We now evaluate the remaining average, $\langle M_1 L^{-1} M_1 \rangle$, in three stages to derive Eq. 3.21. First we expand the differential operations, then evaluate the disorder average. The last stage is to evaluate the integral. We will then combine this integral with Eq. 3.13 to solve Eq. 3.11.

The average to evaluate has the form:

$$\langle M_1 L^{-1} M_1 \rangle \langle \mathbf{h} \rangle = \int d\mathbf{r}' \langle M_1(\mathbf{r}) G(\mathbf{r} - \mathbf{r}') M_1(\mathbf{r}') \rangle \langle \mathbf{h}(\mathbf{r}') \rangle. \quad (3.14)$$

The Green's function is the solution to $(1 - \Lambda^2 \nabla^2) G(\mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}')$, and is:

$$G(z) = \frac{1}{\Lambda^2} \frac{1}{4\pi z} e^{-z/\Lambda}. \quad (3.15)$$

Here, we have defined $z = \mathbf{r} - \mathbf{r}'$.

We now expand the differential operations in Eq. 3.14. We do this in two segments, first with derivatives at \mathbf{r} , then with derivatives at \mathbf{r}' . The first is:

$$\begin{aligned} M_1(\mathbf{r})G(\mathbf{r} - \mathbf{r}')\mathbf{v}(\mathbf{r}') &= [-2\Lambda^2\xi(\mathbf{r})\nabla_{\mathbf{r}}^2 + 2\Lambda^2\nabla_{\mathbf{r}}\xi(\mathbf{r}) \times \nabla_{\mathbf{r}}\times]G(\mathbf{r} - \mathbf{r}')\mathbf{v}(\mathbf{r}') \\ &= 2\xi\mathbf{v}(\mathbf{r}')[\delta(z) - G(z)] + 2\Lambda^2\nabla\xi \times [\nabla_{\mathbf{r}}G(\mathbf{r} - \mathbf{r}')] \times \mathbf{v}(\mathbf{r}'). \end{aligned} \quad (3.16)$$

The second, which was represented by $\mathbf{v}(\mathbf{r}')$ above, is:

$$\begin{aligned} \mathbf{v}(\mathbf{r}') &= M_1(\mathbf{r}')\langle\mathbf{h}(\mathbf{r}')\rangle = e^{i\mathbf{k}\cdot\mathbf{r}'}\left[2\Lambda^2k^2\xi(\mathbf{r}')\mathbf{h} + 2i\Lambda^2\nabla_{\mathbf{r}'}\xi(\mathbf{r}') \times \mathbf{k} \times \mathbf{h}\right] \\ &= 2\Lambda^2e^{i\mathbf{k}\cdot\mathbf{r}'}\left\{[k^2\xi(\mathbf{r}') - i\nabla_{\mathbf{r}'}\xi(\mathbf{r}')\cdot\mathbf{k}]\mathbf{h} + i[\nabla_{\mathbf{r}'}\xi(\mathbf{r}')\cdot\mathbf{h}]\mathbf{k}\right\}. \end{aligned} \quad (3.17)$$

Combining Eqs. 3.16 and 3.17, we obtain:

$$\begin{aligned} M_1(\mathbf{r})G(\mathbf{r} - \mathbf{r}')M_1(\mathbf{r}')\langle\mathbf{h}(\mathbf{r}')\rangle &= 2\Lambda^2e^{i\mathbf{k}\cdot\mathbf{r}'}\left\{\xi(\mathbf{r})\xi(\mathbf{r}')[\delta(z) - G(z)]2k^2\mathbf{h} \right. \\ &\quad + \xi(\mathbf{r})\nabla_{\mathbf{r}'}\xi(\mathbf{h} \otimes \mathbf{k} - \mathbf{k} \otimes \mathbf{h})2i[\delta(z) - G(z)] \\ &\quad \left. - 2\Lambda^2G(z)(\Lambda^{-1} + z^{-1})\left[\xi(\mathbf{r}')\nabla_{\mathbf{r}}\xi(\mathbf{h} \otimes \hat{\mathbf{z}} - \hat{\mathbf{z}} \otimes \mathbf{h})k^2 \right. \right. \\ &\quad \left. \left. + i(\mathbf{h} \otimes \hat{\mathbf{z}} - \hat{\mathbf{z}} \otimes \mathbf{h})(\nabla_{\mathbf{r}}\xi \otimes \nabla_{\mathbf{r}'}\xi)\mathbf{k} + i(\hat{\mathbf{z}} \otimes \mathbf{k} - \mathbf{k} \otimes \hat{\mathbf{z}})(\nabla_{\mathbf{r}}\xi \otimes \nabla_{\mathbf{r}'}\xi)\mathbf{h}\right]\right\}, \end{aligned} \quad (3.18)$$

where we use \otimes to indicate the tensor product.

To perform the disorder average in the second stage, we need various derivatives of the correlation function $R(z)$:

$$\begin{aligned} \langle\xi(\mathbf{r})\nabla_{\mathbf{r}'}\xi\rangle &= \nabla_{\mathbf{r}'}R(|\mathbf{r} - \mathbf{r}'|) = -\hat{\mathbf{z}}\dot{R}(z), \\ \langle\xi(\mathbf{r}')\nabla_{\mathbf{r}}\xi\rangle &= \nabla_{\mathbf{r}}R(|\mathbf{r} - \mathbf{r}'|) = \hat{\mathbf{z}}\dot{R}(z), \quad \text{and} \\ \langle\nabla_{\mathbf{r}}\xi\nabla_{\mathbf{r}'}\xi\rangle &= \nabla\nabla'R(|\mathbf{r} - \mathbf{r}'|) = -\left[\frac{\dot{R}}{z}\mathbf{I} + \hat{\mathbf{z}} \otimes \hat{\mathbf{z}}\left(\ddot{R} - \frac{\dot{R}}{z}\right)\right], \end{aligned}$$

where the overdot indicates differentiation with respect to z , and \mathbf{I} is the identity tensor.

Then averaging Eq. 3.18 gives:

$$\begin{aligned} \langle M_1(\mathbf{r})G(\mathbf{r} - \mathbf{r}')M_1(\mathbf{r}') \rangle \langle \mathbf{h}(\mathbf{r}') \rangle &= 2\Lambda^2 e^{i\mathbf{k}\cdot\mathbf{r}'} \left\{ [A(z) + 2k^2 R(z)\delta(z)] \mathbf{h} \right. \\ &\quad \left. - [B(z) + 2i\dot{R}(z)\delta(z)] (\mathbf{k} \otimes \mathbf{h} - \mathbf{h} \otimes \mathbf{k}) \hat{\mathbf{z}} - C(z) \mathbf{h} (\hat{\mathbf{z}} \otimes \hat{\mathbf{z}}) \right\}, \end{aligned} \quad (3.19)$$

with the scalars A , B , and C given by:

$$\begin{aligned} A(z) &= 2k^2 \left[\dot{R}(z) \Lambda^2 G(z) (\Lambda^{-1} + z^{-1}) - R(z) G(z) \right], \\ B(z) &= 2i \left[\ddot{R}(z) \Lambda^2 G(z) (\Lambda^{-1} + z^{-1}) - \dot{R}(z) G(z) \right], \quad \text{and} \\ C(z) &= 2\Lambda^2 k^2 \dot{R}(z) G(z) (\Lambda^{-1} + z^{-1}). \end{aligned}$$

The final stage in evaluating Eq. 3.14 is to perform the integral over \mathbf{r}' . We first change variables from \mathbf{r}' to z , then integrate over the orientation of z . Using the relations

$$\begin{aligned} \int d\hat{\mathbf{z}} e^{-i\mathbf{k}\cdot\mathbf{z}} &= 4\pi \frac{\sin(kz)}{kz} \equiv F(k, z), \\ \int d\hat{\mathbf{z}} \hat{\mathbf{z}} e^{-i\mathbf{k}\cdot\mathbf{z}} &= \hat{\mathbf{k}} \frac{i}{z} \partial_k F, \quad \text{and} \\ \int d\hat{\mathbf{z}} \hat{\mathbf{z}} \otimes \hat{\mathbf{z}} e^{-i\mathbf{k}\cdot\mathbf{z}} &= \frac{-1}{z^2} \left[\frac{\partial_k F}{k} \mathbf{I} + \hat{\mathbf{k}} \otimes \hat{\mathbf{k}} \left(\partial_k^2 F - \frac{\partial_k F}{k} \right) \right], \end{aligned} \quad (3.20)$$

we find that Eq. 3.14 evaluates to:

$$\begin{aligned} \int d\mathbf{r}' \langle M_1(\mathbf{r})G(\mathbf{r} - \mathbf{r}')M_1(\mathbf{r}') \rangle \langle \mathbf{h}(\mathbf{r}') \rangle &= \\ &= 4\Lambda^2 k^2 e^{i\mathbf{k}\cdot\mathbf{r}} \left\{ [X + R(0)] \mathbf{h} + Y \hat{\mathbf{k}} (\mathbf{h} \cdot \hat{\mathbf{k}}) \right\}. \end{aligned} \quad (3.21)$$

The functions X and Y are given by:

$$\begin{aligned}
X(k) &= \int_0^\infty dz G(z) \left\{ \Lambda^2 (\Lambda^{-1} + z^{-1}) \left[\dot{R}(z^2 F + k^{-1} \partial_k F) - \ddot{R} z k^{-1} \partial_k F \right] \right. \\
&\quad \left. + \dot{R} z k^{-1} \partial_k F - R z^2 F \right\}, \\
Y(k) &= \int_0^\infty dz G(z) \left\{ \Lambda^2 (\Lambda^{-1} + z^{-1}) \left[\ddot{R} z k^{-1} \partial_k F + \dot{R} (\partial_k^2 F - k^{-1} \partial_k F) \right] - \dot{R} z k^{-1} \partial_k F \right\}.
\end{aligned} \tag{3.22}$$

We require the average magnetic field to have $\nabla \cdot \mathbf{h} = 0$, which means that $\mathbf{k} \cdot \mathbf{h} = 0$. We now collect our results from Eqs. 3.13 and 3.21, and insert them into Eq. 3.11:

$$\mathbf{h} e^{i\mathbf{k} \cdot \mathbf{r}} \left[1 + \Lambda^2 k^2 (1 - 3R(0) - 4X) \right] = 0. \tag{3.23}$$

We are interested in solutions consistent with Meissner screening, so we require that k be positive and purely imaginary. Then the field decays on a length scale $\lambda_{\text{eff}} = \frac{i}{k}$, which we identify as the experimentally measured penetration depth. To calculate λ_{eff} , we will solve the equation:

$$\frac{\lambda_{\text{eff}}^2}{\Lambda^2} = 1 - 3R(0) - 4X. \tag{3.24}$$

Inserting $k = \frac{i}{\lambda_{\text{eff}}}$ into our equation for X , we get:

$$\begin{aligned}
X &= 4\pi \int_0^\infty dz G(z) \sinh\left(\frac{z}{\lambda_{\text{eff}}}\right) \left\{ \Lambda^2 (\Lambda^{-1} + z^{-1}) \left[\dot{R} \lambda_{\text{eff}} z^{-1} (z^2 + \lambda_{\text{eff}}^2) - \ddot{R} \lambda_{\text{eff}}^3 \right] \right. \\
&\quad \left. + \dot{R} \lambda_{\text{eff}}^3 - R z \lambda_{\text{eff}} \right\} \\
&+ 4\pi \int_0^\infty dz G(z) \cosh\left(\frac{z}{\lambda_{\text{eff}}}\right) \left\{ \Lambda^2 (\Lambda^{-1} + z^{-1}) \left[-\dot{R} \lambda_{\text{eff}}^2 + \ddot{R} z \lambda_{\text{eff}}^2 \right] - \dot{R} z \lambda_{\text{eff}}^2 \right\}.
\end{aligned} \tag{3.25}$$

Valid solutions for λ_{eff} will require the integral for X to converge and Eq. 3.24 to have solutions.

3.3.2 Correlation Function

A full solution of the disorder-averaged magnetic field, $\langle \mathbf{h} \rangle$, requires knowledge of the correlation function $R(z)$ and hence requires not only a detailed knowledge of the composition, structure, and disorder of the sample, but also a microscopic model to locally determine the superconducting properties from that structure. Without guidance from microscopic calculations, we will use the Matérn one-parameter family of correlation functions[66] to tune the smoothness, as well as the magnitude and correlation length, of the penetration depth fluctuations. Handcock and Wallis [67] parametrize the Matérn class of covariance functions as:

$$R(z) = \frac{R(0)}{2^{\nu-1}\Gamma(\nu)} \left(2\sqrt{\nu} \frac{z}{l}\right)^{\nu} K_{\nu} \left(2\sqrt{\nu} \frac{z}{l}\right), \quad (3.26)$$

where K_{ν} is a modified Bessel function of the second kind and $\Gamma(z)$ is the Gamma function. The intercept at zero separation is the normalized variance of the penetration depth, $R(0) = \sigma_{\lambda}^2 / \langle \lambda \rangle^2 = (\langle \lambda^2 \rangle - \langle \lambda \rangle^2) / \langle \lambda \rangle^2$, and quantifies the magnitude of the inhomogeneity in $\lambda(\mathbf{r})$. The correlation length, l , controls the size of the fluctuations in $\lambda(\mathbf{r})$. The parameter ν controls the smoothness of $\lambda(\mathbf{r})$. Larger ν gives a smoother random field, since it is $\lceil \nu \rceil - 1$ times mean squared differentiable, where $\lceil \cdot \rceil$ is the ceiling function.[67]

Two members of the family deserve specific mention. When $\nu = 1/2$, Eq. 3.26 reduces to the exponential correlation function, $R(z) = R(0) \exp(-z\sqrt{2}/l)$, which is the correlation function of a Markov process in one dimension. The integrals for X in Eq. 3.25 diverge when $R(z) \propto e^{-z}$, making the case $\nu = 1/2$ invalid. In the limit where $\nu \rightarrow \infty$, $R(z) \rightarrow R(0) \exp(-z^2/l^2)$, which is labeled the squared exponential correlation function, to prevent confusion with the Gaussian probability distribution. This correlation function gives the smoothest possible $\lambda(\mathbf{r})$ that can be described within the Matérn covariance family.

3.3.3 Squared Exponential Correlations

We now consider the case of squared exponential correlations, $R(z) = R(0) e^{-z^2/l^2}$. In Fig. 3.1 we plot four realizations of a normally distributed penetration depth with squared




Figure 3.1: Sample realizations of a random penetration depth reveal the influence of the correlation length, l , and variance, $R(0)$. The variance, $R(0) = \langle \lambda^2 \rangle / \langle \lambda \rangle^2 - 1$, controls the width the penetration depth distribution, and the correlation length establishes the characteristic length over which $\lambda(\mathbf{r})$ changes.

exponential correlations, illustrating the effect of the two parameters l and $R(0)$ on $\lambda(\mathbf{r})$. Evaluating Eq. 3.25 gives:

$$\begin{aligned}
 X = & -R(0) \int_0^\infty dz e^{-z/\Lambda} e^{-z^2/l^2} \sinh\left(\frac{z}{\lambda_{\text{eff}}}\right) \frac{\lambda_{\text{eff}}}{\Lambda^2} \left[\left(1 + 2\frac{\Lambda^2}{l^2}\right) + 2z \frac{\Lambda}{l^2} \right] \left(1 + 2\frac{\lambda_{\text{eff}}^2}{l^2}\right) \\
 & + 2R(0) \int_0^\infty dz e^{-z/\Lambda} e^{-z^2/2l^2} \cosh\left(\frac{z}{\lambda_{\text{eff}}}\right) \frac{\lambda_{\text{eff}}^2}{\Lambda^2} \left[z \frac{1}{l^2} \left(1 + 2\frac{\Lambda^2}{l^2}\right) + 2z^2 \frac{\Lambda}{l^4} \right]. \quad (3.27)
 \end{aligned}$$

All of these integrals converge, so we evaluate X as:

$$\begin{aligned}
X = & R(0) \frac{2\lambda_{\text{eff}}^2}{l^2} \\
& + R(0) \frac{\lambda_{\text{eff}}\sqrt{\pi}}{4l^3\Lambda^2} \left\{ \left(l^4 - 2l^2\Lambda\lambda_{\text{eff}} + 4\Lambda^2\lambda_{\text{eff}}^2 \right) \exp \left[\frac{l^2}{4} \left(\frac{1}{\Lambda} + \frac{1}{\lambda_{\text{eff}}} \right)^2 \right] \text{erfc} \left[\frac{l}{2} \left(\frac{1}{\Lambda} + \frac{1}{\lambda_{\text{eff}}} \right) \right] \right. \\
& \left. - \left(l^4 + 2l^2\Lambda\lambda_{\text{eff}} + 4\Lambda^2\lambda_{\text{eff}}^2 \right) \exp \left[\frac{l^2}{4} \left(\frac{1}{\Lambda} - \frac{1}{\lambda_{\text{eff}}} \right)^2 \right] \text{erfc} \left[\frac{l}{2} \left(\frac{1}{\Lambda} - \frac{1}{\lambda_{\text{eff}}} \right) \right] \right\}. \quad (3.28)
\end{aligned}$$

After inserting Eq. 3.28 into Eq. 3.24, we solve for λ_{eff} over three decades in the correlation length, l , and in the disorder variance, $R(0)$ (Fig. 3.2). At large correlation length the effective penetration depth is larger than the average value, representing *suppressed* Meissner screening. Conversely, at small correlation length the effective penetration depth is smaller than the average, indicating *enhanced* screening. The separatrix, where $\lambda_{\text{eff}} = \Lambda$ for all values of $R(0)$, occurs near $l = 1.643\Lambda$. Note that the system is not symmetric about the separatrix, although it becomes more symmetric as $R(0) \rightarrow 1$. This is true for both linear and logarithmic spacing around the separatrix. In other words, neither $|\lambda_{\text{eff}}(l_s + \Delta l) - \langle \lambda \rangle| = |\lambda_{\text{eff}}(l_s - \Delta l) - \langle \lambda \rangle|$ nor $|\lambda_{\text{eff}}(al_s) - \langle \lambda \rangle| = |\lambda_{\text{eff}}(l_s/a) - \langle \lambda \rangle|$ are true, where l_s denotes the separatrix, and a is an arbitrary positive real number. As expected, $\lambda_{\text{eff}} \rightarrow \Lambda$ as $R(0) \rightarrow 0$. Yet even at small disorder, λ_{eff} has variations on the one percent scale, shown by the contours in Fig. 3.2. As we will discuss below, sub-percent variations of λ_{eff} could be significant in the context of a typical measurement of $\Delta\lambda(T)$.

The trends in λ_{eff} can also be seen in Fig. 3.3, where we plot $\lambda_{\text{eff}}/\Lambda$ vs. $R(0)$ at fixed correlation length. All three curves taper to $\lambda_{\text{eff}} = \Lambda$ as the magnitude of disorder decreases. At large correlation length, in this case $l = 10\Lambda$, λ_{eff} increases by ten percent when $R(0) = 0.02$. The effect at small correlation is more modest, but still reaches nearly ten percent by the time $R(0) = 0.1$ when $l = 0.1\Lambda$.

The penetration depth has a temperature dependence that it inherits from the underlying disordered superconducting state. It is natural to expect that $R(0)$ and l will have a temperature dependence of their own, which will create a temperature-induced change in λ_{eff} . This change contributes to any measurement of $\lambda(T)$, but is not related to the gap

structure in momentum space, because it arises from the spatial arrangement of the superconducting state. If we neglected the spatial variation of λ we would erroneously attribute the entire temperature dependence to the order parameter.




Figure 3.2: The effective penetration depth is a strong function of the parameters that characterize the distribution of local penetration depths. Here we show the value of λ_{eff} as the correlation length, l , and variance, $R(0)$, run across three orders of magnitude. This figure considers the case of squared exponential correlations in the penetration depth; a different case is shown in Fig. 3.4. The most important features of this color plot are the range of $\lambda_{\text{eff}}/\langle\lambda\rangle$ and the appearance of values both above and below 1. The calculation is valid when $R(0) \ll 1$, but we show the region with $R(0) > 0.1$ to emphasize the trends seen. Any temperature dependence in l or $R(0)$ will contribute to $\lambda(T)$. This temperature dependence is not accounted for by the superconducting gap.

3.3.4 General Matérn Correlations

To understand the impact of the smoothness of $\lambda(r)$ on the measured penetration depth, λ_{eff} , we now consider the general case of Matérn covariance. Recall that the parameter ν




Figure 3.3: The screening can either be enhanced ($\lambda_{\text{eff}} < \langle \lambda \rangle$) or suppressed ($\lambda_{\text{eff}} > \langle \lambda \rangle$), depending on the correlation length. The curves for $l = 0.1 \langle \lambda \rangle$ and $l = 0.01 \langle \lambda \rangle$ overlap.

controls the smoothness of the penetration depth. With the correlation function defined by Eq. 3.26, we evaluate Eq. 3.25:

$$\begin{aligned}
X = & -\frac{R(0)}{2^{\nu-1}\Gamma(\nu)} \int_0^\infty dz e^{-z/\Lambda} \sinh\left(\frac{z}{\lambda_{\text{eff}}}\right) \frac{\lambda_{\text{eff}}}{\Lambda^2 l^4} \left[l^4 \left(2\sqrt{\nu}\frac{z}{l}\right)^\nu K_\nu\left(2\sqrt{\nu}\frac{z}{l}\right) \right. \\
& + 4\nu l^2 (\Lambda^2 + \Lambda z + \lambda_{\text{eff}}^2) \left(2\sqrt{\nu}\frac{z}{l}\right)^{\nu-1} K_{\nu-1}\left(2\sqrt{\nu}\frac{z}{l}\right) \\
& \left. + 16\nu^2 \lambda_{\text{eff}}^2 \Lambda (\Lambda + z) \left(2\sqrt{\nu}\frac{z}{l}\right)^{\nu-2} K_{\nu-2}\left(2\sqrt{\nu}\frac{z}{l}\right) \right] \\
& + \frac{R(0)}{2^{\nu-1}\Gamma(\nu)} \int_0^\infty dz e^{-z/\Lambda} \cosh\left(\frac{z}{\lambda_{\text{eff}}}\right) \frac{4\nu \lambda_{\text{eff}}^2}{\Lambda^2 l^4} \left[l^2 z \left(2\sqrt{\nu}\frac{z}{l}\right)^{\nu-1} K_{\nu-1}\left(2\sqrt{\nu}\frac{z}{l}\right) \right. \\
& \left. + 4\nu \Lambda (\Lambda + z) z \left(2\sqrt{\nu}\frac{z}{l}\right)^{\nu-2} K_{\nu-2}\left(2\sqrt{\nu}\frac{z}{l}\right) \right]. \quad (3.29)
\end{aligned}$$

These integrals can be evaluated using equation 6.621.3 in Gradshteyn and Ryzhik:[68]

$$\int_0^\infty x^{\mu-1} e^{-\alpha x} K_\nu(\beta x) dx = \sqrt{\pi} \frac{(2\beta)^\nu}{(\alpha + \beta)^{\mu+\nu}} \frac{\Gamma(\mu + \nu)\Gamma(\mu - \nu)}{\Gamma(\mu + \frac{1}{2})} {}_2F_1\left(\mu + \nu, \nu + \frac{1}{2}; \mu + \frac{1}{2}; \frac{\alpha - \beta}{\alpha + \beta}\right), \quad (3.30)$$

which requires $\text{Re } \mu > |\text{Re } \nu|$ and $\text{Re } (\alpha + \beta) > 0$. The function ${}_2F_1(a, b; c; z)$ is Gauss' hypergeometric function. Using the integral in Eq. 3.30 to evaluate Eq. 3.29, we find the constraints

$$\nu > \frac{3}{2} \quad \text{and} \quad \frac{\lambda_{\text{eff}}}{\Lambda} > \frac{l}{l + 2\Lambda\sqrt{\nu}}. \quad (3.31)$$

The full solution for X is then:

$$\begin{aligned} X = & \frac{R(0) \sqrt{\pi}}{\Gamma(\nu)} \left(\frac{4\nu}{l^2}\right)^\nu \frac{\lambda_{\text{eff}}}{\Lambda} \left\{ \frac{\lambda_{\text{eff}}^2 + \Lambda^2}{2\Lambda} \frac{\Gamma(2\nu - 1)}{\Gamma(\nu + \frac{1}{2})} \left[a^{-(2\nu-1)} F1(\square) - b^{-(2\nu-1)} F1(\clubsuit) \right] \right. \\ & + \frac{\Gamma(2\nu)}{\Gamma(\nu + \frac{3}{2})} \left[\frac{\lambda_{\text{eff}} + \Lambda}{2\Lambda} a^{-2\nu} F2(\square) + \frac{\lambda_{\text{eff}} - \Lambda}{2\Lambda} b^{-2\nu} F2(\clubsuit) \right] \\ & + \frac{1}{\Lambda} \frac{\Gamma(2\nu + 1)}{\Gamma(\nu + \frac{3}{2})} \left[a^{-(2\nu+1)} F3(\square) - b^{-(2\nu+1)} F3(\clubsuit) \right] \\ & + \frac{\lambda_{\text{eff}}^2 \Lambda}{4} \frac{\Gamma(2\nu - 3)}{\Gamma(\nu - 1)} \left[a^{-(2\nu-3)} F4(\square) - b^{-(2\nu-3)} F4(\clubsuit) \right] \\ & + \frac{\lambda_{\text{eff}}}{4} \frac{\Gamma(2\nu - 2)}{\Gamma(\nu + \frac{1}{2})} \left[(\lambda_{\text{eff}} + \Lambda) a^{-(2\nu-2)} F5(\square) - (\lambda_{\text{eff}} - \Lambda) b^{-(2\nu-2)} F5(\clubsuit) \right] \\ & \left. + \frac{\lambda_{\text{eff}}}{2} \frac{\Gamma(2\nu - 1)}{\Gamma(\nu + \frac{3}{2})} \left[a^{-(2\nu-1)} F6(\square) + b^{-(2\nu-1)} F6(\clubsuit) \right] \right\}, \quad (3.32) \end{aligned}$$

where we have introduced the variables

$$\begin{aligned}
 a &= \frac{1}{\Lambda} + \frac{1}{\lambda_{\text{eff}}} + \frac{2\sqrt{\nu}}{l}, \\
 b &= \frac{1}{\Lambda} - \frac{1}{\lambda_{\text{eff}}} + \frac{2\sqrt{\nu}}{l}, \\
 \square &= \frac{l(\lambda_{\text{eff}} + \Lambda) - 2\lambda_{\text{eff}}\Lambda\sqrt{\nu}}{l(\lambda_{\text{eff}} + \Lambda) + 2\lambda_{\text{eff}}\Lambda\sqrt{\nu}}, \\
 \clubsuit &= \frac{l(\Lambda - \lambda_{\text{eff}}) + 2\lambda_{\text{eff}}\Lambda\sqrt{\nu}}{l(\Lambda - \lambda_{\text{eff}}) - 2\lambda_{\text{eff}}\Lambda\sqrt{\nu}},
 \end{aligned}$$

and functions

$$\begin{aligned}
 F1(\cdot) &= {}_2F_1(2\nu - 1, \nu - \frac{1}{2}; \nu + \frac{1}{2}; \cdot), \\
 F2(\cdot) &= {}_2F_1(2\nu, \nu - \frac{1}{2}; \nu + \frac{3}{2}; \cdot), \\
 F3(\cdot) &= {}_2F_1(2\nu + 1, \nu + \frac{1}{2}; \nu + \frac{3}{2}; \cdot), \\
 F4(\cdot) &= {}_2F_1(2\nu - 3, \nu - \frac{3}{2}; \nu - \frac{1}{2}; \cdot), \\
 F5(\cdot) &= {}_2F_1(2\nu - 2, \nu - \frac{3}{2}; \nu + \frac{1}{2}; \cdot), \quad \text{and} \\
 F6(\cdot) &= {}_2F_1(2\nu - 1, \nu - \frac{3}{2}; \nu + \frac{3}{2}; \cdot).
 \end{aligned}$$

Inserting this expression for X into Eq. 3.24, we can solve for λ_{eff} after choosing a value for the smoothness parameter ν . In Fig. 3.4, we have chosen $\nu = 2$, close to the lower bound of $\frac{3}{2}$ required for convergence of X . The results are almost identical to the case of squared exponential correlations (Fig. 3.2); evidently λ_{eff} is not much affected by changes in the smoothness of $\lambda(\mathbf{r})$ for the Matérn family of correlation functions. The qualitative features of interest to us are still present: there are regions of enhanced screening and regions of suppressed screening, the effect grows on increasing the variance of $\lambda(\mathbf{r})$, and changes in λ_{eff} at the one percent level persist down to small disorder. Quantitatively, the results in Figs. 3.2 and 3.4 differ by five percent in the region near $l = 1$ and $R(0) = 1$, where the difference is largest.




Figure 3.4: The effective penetration depth for Matérn correlations when $\nu = 2$ (shown here) has strong similarities to Fig. 3.2, which represents the limiting case where $\nu \rightarrow \infty$. These similarities imply that the smoothness of the random penetration depth does not strongly affect λ_{eff} .

3.4 Discussion

The measured $\lambda(T)$ in a non-uniform superconductor will be determined by both the momentum space gap structure and the real space variations of the penetration depth. We calculate the influence of spatial fluctuations in the penetration depth by solving the stochastic London equation in the limit of small fluctuations. This gives an equation (Eq. 3.24) for the disorder-averaged magnetic field in terms of the penetration depth correlation function. We then solve this equation for two example correlation functions to find λ_{eff} , the decay length of the disorder-averaged field, which we identify as the penetration depth measured experimentally. We find that λ_{eff} can be either smaller or larger than the average penetration depth, depending on the correlation length of λ . More importantly, the variance and

correlation length of λ will likely change with temperature, endowing the experimentally measured penetration depth with temperature dependence that is unrelated to the superconducting order parameter.

This work shows that there can be a disorder-induced change of the penetration depth that is not caused by the structure of the superconducting gap in momentum space. Rather, it reflects the real space variations of the order parameter. An interpretation that assumed a spatially uniform penetration depth would infer a larger modulation of $\Delta(\mathbf{k})$ than truly exists. Because $\Delta(\mathbf{k})$ is the starting point for investigations of the mechanism of the superconductor, this omission could lead us astray when we seek to determine the underlying mechanism.

How significant is the effect of disorder-induced change in the penetration depth, given that $\lambda_{\text{eff}}/\langle\lambda\rangle$ approaches 1 over a large segment of the $R(0)$ - l plane? Modern measurements can routinely resolve sub-nanometer changes in the penetration depth; [6, 10, 60] in cuprates and pnictides the penetration depth is approximately 200 nm, and a 1 nm change in λ yields $\frac{\Delta\lambda}{\lambda}$ of 0.5% — making even small changes in $\lambda_{\text{eff}}/\langle\lambda\rangle$ potentially significant.

Two issues are worth emphasizing. First, we have made no assumption about the distribution of $\lambda(\mathbf{r})$, i.e., whether it is normally distributed or follows a different probability distribution. However, the calculation presented here only extends to second order, and any non-normality only enters at third order and above. Second, λ_{eff} has a complicated dependence on the correlation function $R(z)$, and we know neither its functional form nor its temperature dependence. Hence we cannot make any tidy prediction for the low-temperature behavior of $\lambda(T)$; there is no power-law to be had.

Even without perfect knowledge of $R(z)$, it may be possible learn more about λ_{eff} by taking advantage of the general constraints that apply to all correlation functions.[66, 69] In particular, the strong similarities between the two cases presented here (Figs. 3.2 and 3.4) lead us to expect qualitatively similar behavior in λ_{eff} for most possible correlation functions.

To make a stronger statement about $\lambda(T)$, we need to determine the local superconducting properties of a given chemically doped and intrinsically disordered material, which naturally depends on the microscopic details of the superconducting mechanism. Although it

should be possible to extract a local penetration depth or superfluid density from numerical methods such as solving the Bogoliubov–de Gennes equations on a lattice, to the best of our knowledge this has never been attempted. Several groups have calculated the *disorder-averaged* superfluid stiffness using this approach, for both s-wave[70] and d-wave[71, 72] models. The full temperature dependence of the disorder-averaged superfluid density can also be calculated,[73] but is incomplete, for we have shown that the real space inhomogeneity of the superconducting state also contributes to the temperature dependence.

The larger message is that some measured properties of disordered superconductors will not be determined by their disorder averages alone; inhomogeneities can affect the measured properties in an experiment-dependent manner. For example, the heat capacity will be given by the disorder average because it is additive, but we have seen that the penetration depth is non-trivially affected by the disorder. Nonetheless, these two experiments are both traditionally interpreted as measuring the same thing — the magnitude of the single-particle gap, $\Delta(\mathbf{k})$.

These results give a specific example of the potential impact of spatial variation on measurements of the penetration depth. With a full consideration of the impact of spatial variation on different measured quantities, as well as a complete understanding of how random chemical doping gives rise to a non-uniform superconducting state, we will be able to integrate a complete account of the effects of disorder into our understanding of unconventional superconductivity.

Acknowledgments

We thank John Kirtley, Steve Kivelson, Jim Sethna, and Jörg Schmalian for helpful discussions. We would also like to thank John Kirtley for checking some of these calculations. This work is supported by the Department of Energy, Office of Basic Energy Sciences, Division of Materials Sciences and Engineering, under contract DE-AC02-76SF00515

Chapter 4

Designing Disordered Dot Arrays

In this chapter I will describe how we make disordered dot arrays. Not the fabrication, which is covered in Chapter 5, but rather the process of determining target dot positions. We begin with an undisturbed two-dimensional Bravais lattice, with the position of each dot given by $\mathbf{r} = n\mathbf{a} + m\mathbf{b}$, where \mathbf{a} and \mathbf{b} are Bravais lattice vectors and n and m are two integers. All of the examples in this chapter are square lattices, where $\mathbf{a} = a\hat{x}$ and $\mathbf{b} = a\hat{y}$, but everything should be general. For a finite array I can define a single integer to index the entire array. So we restrict n and m to the sets $\{0 \dots N-1\}$ and $\{0 \dots M-1\}$, respectively. Then there are $N \times M$ dots in the array, and we can index all of them by defining $i = n + Nm$.

To add disorder to an array, I'll generate two sets of random variables, $\{X_i\}$ and $\{Y_i\}$, which are the x and y displacements of each dot from its undisturbed lattice position. The x and y displacements are independent from each other, although there certainly can be correlations within each set. The correlation structure of the displacements will control the lattice morphology and hence, the superconductivity.

4.1 Random Variables

It's a beautiful day outside, so we're in the lab, measuring two things as voltage signals. One is averaging to 1 V, with 0.3 V of noise. The other averages at -0.3 V, with 0.5 V of noise. Since our next experiment is to measure the sum of these two signals, what will the new signal be? One approach is to say you have two variables, x and y , with their own probability density

functions, $f(x)$ and $f(y)$. Then we need to find $f(x + y)$, which I think is the convolution of $f(x)$ with $f(y)$, provided x and y are independent. But then we need a function for $f(x)$ and $f(y)$. We could assume the noise is Gaussian and do the integral, although real noise is almost never Gaussian.

The other approach is to say that each signal is represented by a random variable, X or Y . We already know the mean and sample standard deviation for each one. The average value of the new signal is 0.7 V, but to derive that using random variables we use the expectation operator, which is linear. So we get $\langle X + Y \rangle = \langle X \rangle + \langle Y \rangle = 0.7$ V, as expected. Typically the expectation operator is denoted as $E[\cdot]$, but I've used $\langle \cdot \rangle$ for consistency with Chapter 3 and typical physics convention.

To talk about the noise on the new signal, I need to introduce two new operators, the variance and covariance. The covariance of two random variables is defined as:

$$\text{Cov}(X, Y) \equiv \langle (X - \langle X \rangle)(Y - \langle Y \rangle) \rangle. \quad (4.1)$$

Linearity of the expectation operator lets us simplify to:

$$\text{Cov}(X, Y) = \langle XY \rangle - \langle X \rangle \langle Y \rangle. \quad (4.2)$$

The variance of a random variable is just its covariance with itself, or:

$$\text{Var}(X) \equiv \text{Cov}(X, X), \quad (4.3)$$

which leads us to the familiar expression:

$$\text{Var}(X) = \langle X^2 \rangle - \langle X \rangle^2. \quad (4.4)$$

The variance is often denoted as σ^2 , which is to say that the standard deviation (σ) is the square root of the variance.

What I earlier described as just “the noise” was the sample standard deviation. Then $\text{Var}(X) = (0.3 \text{ V})^2$ and $\text{Var}(Y) = (0.5 \text{ V})^2$, or vice versa. We can easily calculate the variance

of the new signal:

$$\begin{aligned}
 \text{Var}(X + Y) &= \langle (X + Y)^2 \rangle - (\langle X + Y \rangle)^2 \\
 &= \langle X^2 + Y^2 + 2XY \rangle - (\langle X \rangle^2 + \langle Y \rangle^2 + 2\langle X \rangle \langle Y \rangle) \\
 &= \text{Var}(X) + \text{Var}(Y) + 2\text{Cov}(X, Y).
 \end{aligned} \tag{4.5}$$

If the two signals are uncorrelated, $\text{Cov}(X, Y) = 0$, and the variances add. Assuming that, then we find that the noise in the new sample is:

$$\begin{aligned}
 \sigma_{x+y} &= \sqrt{\text{Var}(X + Y)} \\
 &= \sqrt{\text{Var}(X) + \text{Var}(Y)} \\
 &= 0.6 \text{ V},
 \end{aligned} \tag{4.6}$$

to one significant figure. This is the reasoning behind the laboratory maxim that “noise sources add in quadrature”, since different noise sources are usually uncorrelated. Two things to remember are that we did not need to assume anything about the distribution of the noise in our two voltage measurements, and we also did not need to assume independence. We only had to assume the measurements are uncorrelated, which is a weaker statement.

4.2 Correlations

I’ve stated that the correlation structure of the displacements $\{X_i\}$ and $\{Y_i\}$ will determine what the disordered lattice looks like. This can be seen in Fig. 4.1, which shows three different disordered arrays, each with different correlation structure. To make “correlation structure” more precise, we need to define the correlation function. For a discrete set of random variables, like the set of x displacements $\{X_i\}$, the correlation function is defined as:

$$\begin{aligned}
 R(i, j) &= \text{Cov}(X_i, X_j) \\
 &= \left\langle \left(X_i - \langle X_i \rangle \right) \left(X_j - \langle X_j \rangle \right) \right\rangle.
 \end{aligned} \tag{4.7}$$




Figure 4.1: Three example disordered dot arrays. The obvious differences in morphology are due to different statistical correlations in the displacements used to create disorder.

The set $\{X_i\}$ is defined on a lattice, but we can also consider, like in Section 3.2, a random function of a continuously-variable position (a *random field*). In this case the correlation function is defined as:

$$\begin{aligned} R(\mathbf{r}, \mathbf{r}') &= \text{Cov}(\xi(\mathbf{r}), \xi(\mathbf{r}')) \\ &= \left\langle \left(\xi(\mathbf{r}) - \langle \xi(\mathbf{r}) \rangle \right) \left(\xi(\mathbf{r}') - \langle \xi(\mathbf{r}') \rangle \right) \right\rangle. \end{aligned} \quad (4.8)$$

A stationary random process has no special origin, so the correlation function depends only on the difference, $\mathbf{r} - \mathbf{r}'$. This also implies that the mean and variance are independent of position as well. In addition to stationarity, I'm going to require that the displacements also be isotropic, meaning the correlation function only depends on the distance between two points, denoted $z = |\mathbf{r} - \mathbf{r}'|$. With the two restrictions of stationary and isotropic, we can say generically that the correlation function has a maximum at $z = 0$, where it equals the variance of the process ($R(0) = \sigma^2$), and obeys the inequality $|R(z)| \leq R(0)$. [63]

The displacements are a discrete set of random variables defined on a lattice, so the correlation function has the form $R(i, j)$, as in Eq. 4.7. Because I've restricted this to stationary and isotropic displacements, we can use $R(z)$ by considering z to be the distance between lattice sites i and j .

If we want the displacements to be uncorrelated between sites, the correlation function

will simply be a delta function weighted by the variance. To pick a correlation function, I'm going to make the same choice as I did in Section 3.3.2, and use the Matérn family of covariance functions, defined in Eq. 3.26 as:

$$R(z) = \frac{R(0)}{2^{\nu-1}\Gamma(\nu)} \left(2\sqrt{\nu} \frac{z}{l}\right)^{\nu} K_{\nu} \left(2\sqrt{\nu} \frac{z}{l}\right), \quad (4.9)$$

where $\nu > 0$, K_{ν} is a modified Bessel function of the second kind and $\Gamma(z)$ is the Gamma function. The motivation for the choice is the same as before: the Matérn family has an additional parameter that controls the smoothness of the disorder. We can get some sense of the effect of this third parameter in Fig. 4.2, which plots $R(z)$ curves for several values of ν , as well as showing a realization of a random field with $\nu = 1/2$ and $\nu \rightarrow \infty$. (As a practical matter, a ν of 10 or more is tough to distinguish from this limiting case.) You can see that




Figure 4.2: Illustrating the effect of the Matérn smoothness parameter on a random field. (a) Plot of the covariance function, Eq. 4.9, for several values of the smoothness parameter ν . (b) An example random field in the limiting case $\nu \rightarrow \infty$. (c) An example random field when $\nu = 1/2$.

the two realizations are very different in appearance, showing the effect of ν . We can understand this by invoking the Wiener–Khinchin theorem, which states that for any stationary random process, the Fourier transform of its correlation function is the power spectrum of fluctuations.[64] In the plot, you can see that smaller values of ν develop a cusp in $R(z)$ near 0. This will create higher frequency components in the Fourier transform, leading to the spiky appearance. Loosely speaking, decreasing ν tunes from a very correlated to an uncorrelated random process. A random field must be a continuous function of position, so it cannot be truly uncorrelated. The dot displacements I’m constructing here live on a discrete lattice and become uncorrelated when ν goes to 0. This is why the difference between the continuous calculations and the discrete implementation becomes important at small ν , as in Section 4.5.2.

The form of Eq. 4.9 is a little strange, since ν appears in the argument to the Bessel function. The motivation for this parameterization of the Matérn family (it’s not the only one) is that the integral of $R(z)$ over the two–dimensional plane is $\pi l^2 R(0)$, regardless of ν . The importance of this is...something. As I mentioned, the Wiener–Khinchin theorem relates the correlation function to the spectral content of fluctuations. Because the normalization fixes the area under the curve of the correlation function, the dc component of the spectral density is always $\pi l^2 R(0)$. I’m not sure what to make of that. Perhaps it could define the correlation length l .

To recap, we choose a lattice geometry and periodicity. Then we choose a correlation function, either Matérn or uncorrelated. For Matérn we choose $R(0)$, l , and ν ; while for uncorrelated we only choose $R(0)$. Then we generate two independent sets of normally–distributed random numbers, $\{X_i\}$ and $\{Y_i\}$, both of which have the chosen correlation structure. These random numbers serve as the displacement of each dot position from its original location, creating a disordered lattice.

Three example arrays are shown in Fig. 4.3. All are based on a square lattice with 500 nm period; the section displayed is 30 μm by 15 μm . The magnitude of disorder, or $\text{std}(X)$ and $\text{std}(Y)$, is 150 nm for each example. Furthermore, the correlation length is 2 μm in each case. What differs is the smoothness parameter ν , which is 10, 2, and 1 from top to bottom.

The next two sections go into the details of how we can generate the displacements.




Figure 4.3: The effect of the Matérn smoothness parameter on the resulting dot arrays. At left are $30\ \mu\text{m}$ by $15\ \mu\text{m}$ sections of three dot arrays, $\nu = 10, 2,$ and 1 from top to bottom. At right is a histogram of the distance between nearest neighbor dots for each of the three arrays.

After that, I'll describe several different ways to characterize the resulting array, including the nearest neighbor distance shown in Fig. 4.3.

4.3 Generating Correlated Random Numbers

This section answers two related questions:

1. How do I generate random variables with a specified distribution and correlation function?
2. Given a set of random variables, what are the statistics of some function or weighted average of them?

Consider a set of independent and identically distributed random variables $\{G_i\}$. The index i is a position on a one-dimensional line, but the results generalize to higher dimensions. Define X_i as a weighted sum of the G_i :

$$X_i \equiv \sum_j a_{i-j} G_j. \quad (4.10)$$

It's easy to see that

$$\langle X_i \rangle = \langle G \rangle \sum_j a_j \quad (4.11)$$

for all i . The $\{X_i\}$ are stationary, because they are defined as a convolution of a set of independent and identically distributed random variables. Isotropy requires some symmetries among the a_i , but we'll take that for granted. To calculate the correlation function,

$$R(l) = \langle (X_{l+i} - \langle X_{l+i} \rangle) (X_i - \langle X_i \rangle) \rangle, \quad (4.12)$$

we insert Eq. 4.10 into Eq. 4.12 and simplify:

$$\begin{aligned} R(l) &= \left\langle \left(\sum_j a_{l+i-j} (G_j - \langle G_j \rangle) \right) \left(\sum_w a_{i-w} (G_w - \langle G_w \rangle) \right) \right\rangle \\ &= \sum_{j,w} a_{l+i-j} a_{i-w} \langle (G_j - \langle G_j \rangle) (G_w - \langle G_w \rangle) \rangle \\ &= \sigma_G^2 \sum_w a_{l+w} a_w. \end{aligned} \quad (4.13)$$

The last line follows from the independence of the $\{X_i\}$ and the definition of the variance, so that:

$$\langle (G_j - \langle G_j \rangle) (G_w - \langle G_w \rangle) \rangle = \sigma_G^2 \delta_{j,w}. \quad (4.14)$$

So we see that the correlation function of X_i is proportional to the self convolution of the filter kernel (a_j), provided it has inversion symmetry ($a_j = a_{-j}$). Then we can use the convolution theorem to construct a filter kernel which will give an X_i with the desired correlation function. For a continuous random variable with desired correlation function

given by $R(z)$, the required filter kernel is:

$$a = \mathcal{F}^{-1} \left[\sqrt{\mathcal{F}[R(z)]} \right], \quad (4.15)$$

where $\mathcal{F}, \mathcal{F}^{-1}$ is the Fourier transform and its inverse. To completely specify the probability distribution of X_i , we need to determine all of its statistical moments in terms of the corresponding moments of the G_j . This is easiest to do with the characteristic function.

4.3.1 The Characteristic Function

For any random variable G , the *characteristic function* of G is defined as:

$$\varphi_G(t) \equiv \langle e^{itG} \rangle. \quad (4.16)$$

The exponential of a random variable is defined as a power series:

$$e^{itG} = 1 + itG + \frac{(it)^2}{2} G^2 + \dots \quad (4.17)$$

We see that the power-series coefficients of $\varphi_G(t)$ are just the statistical moments of G . Specifically, the k -th moment of G is proportional to the k -th derivative of the characteristic function:

$$\langle G^k \rangle = (-i)^k \left. \frac{d^k}{dt^k} \varphi_G(t) \right|_{t=0}. \quad (4.18)$$

If G is a continuous random variable, the characteristic function turns out to be the Fourier transform of the probability distribution:

$$\varphi_G(t) = \langle e^{itG} \rangle = \int dg e^{itg} P_G(g), \quad (4.19)$$

where $P_G(g)$ is the probability density function of G .

From the definitions, Eq. 4.16 and Eq. 4.10, it's clear that the characteristic function of X is:

$$\varphi_X(t) = \prod_i \varphi_G(a_i t), \quad (4.20)$$

because the $\{G_i\}$ are independent and identically distributed. While this is complete and correct, it's pretty useless. To calculate the variance of X we would need to take a second derivative of a product of many power series, and so on. It helps to introduce an alternate set of statistical moments, the *cumulants*. Where the moments $\langle G^k \rangle$ are the power-series coefficients of $\varphi_G(t)$, the cumulants $\langle\langle G^k \rangle\rangle$ are the power-series coefficients of $\ln \varphi_G(t)$. Or:

$$\ln \varphi_G(t) \equiv \sum_{k=1}^{\infty} \langle\langle G^k \rangle\rangle \frac{(it)^k}{k!}. \quad (4.21)$$

The cumulants are linear combinations of the moments, which you can see by equating Eq. 4.16 and the exponential of Eq. 4.21, and manipulating the power series. The only important ones to know are that the first cumulant is the mean, and the second cumulant is the variance.

More importantly for us, we can combine Eq. 4.20 and Eq. 4.21 to see that:

$$\begin{aligned} \ln \varphi_X(t) &= \sum_j \ln \varphi_{G_j}(a_j t) \\ &= \sum_j \sum_k \langle\langle G_j^k \rangle\rangle \frac{(ia_j t)^k}{k!}. \end{aligned} \quad (4.22)$$

After changing the order of summation, we find:

$$\langle\langle X^k \rangle\rangle = \langle\langle G^k \rangle\rangle \sum_i a_i^k. \quad (4.23)$$

Which indirectly specifies the probability distribution of X . It tells us everything we could want to know about X , but getting back to a probability density function won't be easy. An important special case is when the G_i are normally distributed. The normal distribution has only two non-zero cumulants, the mean and variance. In this case, X is also normally distributed. This is an extension of the more familiar knowledge that the sum of two normal random variables is also normally distributed.

Now we know how to make the displacements. Generate independent normally distributed numbers for each lattice site, and filter with the appropriate filter kernel. I've chosen to normalize the filter kernel so that the variance of X matches that of G , Eq. 4.23 when

$k = 2$. This makes it easy to specify the standard deviation of the displacements, which is the knob we use to control the amount of disorder in the lattice. Although the standard deviation of the displacements may not be the best way to parametrize the “amount of disorder” in the lattice, see Section 4.5.

One final note is that if we generate a set of correlated displacements for a given lattice, the histogram of those displacements won’t actually be Gaussian. It will be missing events in the tails, which makes intuitive sense, because to get a filtered variable with a high z -score would require many of the original independent variables in that area to have large z -scores. The fancy way to say it is that this method of generating correlated random variables is non-ergodic. But if we look at any given lattice site over many realizations of disorder, that will have a Gaussian distribution.

4.4 Matérn Derivations

Back in Section 3.3.2 I mentioned two special cases of the Matérn covariance functions, $\nu = 1/2$, where $R(z) = R(0) \exp(-z\sqrt{2}/l)$ and $\nu \rightarrow \infty$, where $R(z) \rightarrow R(0) \exp(-z^2/l^2)$. To derive both of these results, we look up an integral representation of K_ν , equation 8.432.5 in [68]:

$$K_\nu(z) = \frac{\Gamma(\nu + \frac{1}{2})(2z)^\nu}{\sqrt{\pi}} \int_0^\infty \frac{\cos t dt}{(t^2 + z^2)^{\nu+1/2}}, \quad (4.24)$$

which requires $\nu \geq -1/2$ and $z > 0$ for $z, \nu \in \mathbb{R}$. The Bessel $K_\nu(z)$ diverge as $z^{-\nu}$ as $z \rightarrow 0$, except K_0 which diverges logarithmically. Plugging Eq. 4.24 into Eq. 4.9 and simplifying, we get:

$$\frac{R(z)}{R(0)} = \frac{\Gamma(\nu + \frac{1}{2})}{\Gamma(\nu)} \frac{2^{2\nu+1}}{\sqrt{\pi}} \nu^\nu \left(\frac{z}{l}\right)^{2\nu} \int_0^\infty \frac{\cos t dt}{\left(t^2 + 4\nu \frac{z^2}{l^2}\right)^{\nu+1/2}}. \quad (4.25)$$

We now make the perhaps-suggestive substitution $t = zk$ and simplify considerably to get:

$$\frac{R(z)}{R(0)} = \frac{\Gamma(\nu + \frac{1}{2})}{\Gamma(\nu)\sqrt{\nu}} \frac{l}{\sqrt{\pi}} \int_0^\infty \frac{\cos(zk) dk}{\left(1 + \frac{k^2 l^2}{4\nu}\right)^{\nu+1/2}}. \quad (4.26)$$

4.4.1 Exponential

When we set $\nu = 1/2$, the integral in Eq. 4.26 becomes:

$$\int_0^\infty \frac{\cos(zk) \mathbf{d}k}{1 + \frac{k^2 l^2}{2}} = \frac{\pi}{l\sqrt{2}} e^{-z\sqrt{2}/l}, \quad (4.27)$$

and we find that:

$$R(z) = R(0) e^{-z\sqrt{2}/l}. \quad (4.28)$$

4.4.2 Squared Exponential

In the limit that $\nu \rightarrow \infty$,

$$\lim_{\nu \rightarrow \infty} \left(1 + \frac{k^2 l^2}{4\nu}\right)^{-\nu-1/2} = e^{-k^2 l^2/4}, \quad (4.29)$$

which comes from a definition of the exponential function. Then the integral in Eq. 4.26 becomes:

$$\int_0^\infty e^{-k^2 l^2/4} \cos(zk) \mathbf{d}k = \frac{\sqrt{\pi}}{l} e^{-z^2/l^2}. \quad (4.30)$$

We also need to know that:

$$\lim_{\nu \rightarrow \infty} \frac{\Gamma(\nu + \frac{1}{2})}{\Gamma(\nu)\sqrt{\nu}} = 1, \quad (4.31)$$

which comes from the Stirling series. Putting these together, we find that:

$$\lim_{\nu \rightarrow \infty} R(z) = R(0) e^{-z^2/l^2}, \quad (4.32)$$

as promised.

4.4.3 Spectral Function

Recall from Section 4.2 that the spectral density of fluctuations is the Fourier transform of the correlation function. In this chapter we care about two dimensions. So let's calculate

the spectral density of fluctuations, cleverly denoted $R(k)$, for the Matérn family.

$$R(k) = \int \mathbf{d}z e^{-i\mathbf{k}\cdot\mathbf{z}} R(z). \quad (4.33)$$

Since the correlation function doesn't depend on the angle of \mathbf{r} , we can perform the angular integral:

$$\int_0^{2\pi} \mathbf{d}\theta e^{-ikz \cos \theta} = 2\pi J_0(kz), \quad (4.34)$$

to get:

$$\frac{R(k)}{R(0)} = \frac{4\pi}{\Gamma(\nu)} \left(\frac{\sqrt{\nu}}{l} \right)^\nu \int_0^\infty z^{\nu+1} K_\nu \left(2\sqrt{\nu} \frac{z}{l} \right) J_0(kz) \mathbf{d}z. \quad (4.35)$$

The integral can be done, and we end up with:

$$R(k) = R(0) \pi l^2 \left(1 + \frac{k^2 l^2}{4\nu} \right)^{-(\nu+1)}. \quad (4.36)$$

In the limit that $\nu \rightarrow \infty$, $R(k) \rightarrow R(0) \pi l^2 \exp(-k^2 l^2/4)$, which we can find from taking the limit in Eq. 4.36 or directly Fourier transforming Eq. 4.32.

4.4.4 Filter Kernel

In Section 4.3 we found that we can create a set of correlated random variables by filtering a set of uncorrelated variables. The filter kernel is determined by the desired correlation function according to Eq. 4.15, which in two dimensions takes the form:

$$a_\nu(\mathbf{r}) = \int \frac{\mathbf{d}\mathbf{k}}{4\pi^2} \sqrt{R(k)} e^{i\mathbf{k}\cdot\mathbf{z}}. \quad (4.37)$$

Of course, $R(k)$ only depends on the magnitude of \mathbf{k} , so we can perform the angular integral:

$$\int_0^{2\pi} \mathbf{d}\theta e^{ikz \cos \theta} = 2\pi J_0(kz), \quad (4.38)$$

so that:

$$a_\nu(z) = \int_0^\infty \frac{\mathbf{d}k}{2\pi} k \sqrt{R(k)} J_0(kz). \quad (4.39)$$

The limiting case of squared exponential correlations is more straightforward, so we start by evaluating:

$$\begin{aligned} a_{\nu \rightarrow \infty}(z) &= \frac{l}{2\sqrt{\pi}} \sqrt{R(0)} \int_0^\infty k J_0(kz) e^{-k^2 l^2 / 8} \mathbf{d}k \\ &= \frac{2}{l\sqrt{\pi}} \sqrt{R(0)} e^{-2z^2/l^2}. \end{aligned} \quad (4.40)$$

And we can easily check that

$$2\pi \int_0^\infty z a_{\nu \rightarrow \infty}^2(z) \mathbf{d}z = R(0). \quad (4.41)$$

In the general case, the integral is:

$$a_\nu(z) = \frac{l}{2\sqrt{\pi}} \sqrt{R(0)} \int_0^\infty k \left(1 + \frac{k^2 l^2}{4\nu}\right)^{-\frac{\nu+1}{2}} J_0(kz) \mathbf{d}k, \quad (4.42)$$

which evaluates to:

$$a_\nu(z) = \frac{2}{l\sqrt{\pi}} \frac{\nu^{(\nu+3)/4}}{\Gamma(\frac{\nu+1}{2})} \sqrt{R(0)} \left(\frac{z}{l}\right)^{-(1-\nu)/2} K_{\frac{1-\nu}{2}} \left(2\sqrt{\nu} \frac{z}{l}\right). \quad (4.43)$$

We can evaluate the $z = 0$ intercept separately:

$$a_\nu(0) = \frac{l}{2\sqrt{\pi}} \sqrt{R(0)} \int_0^\infty k \left(1 + \frac{k^2 l^2}{4\nu}\right)^{-\frac{\nu+1}{2}} \mathbf{d}k, \quad (4.44)$$

to get:

$$a_\nu(0) = \frac{2}{l\sqrt{\pi}} \frac{\nu}{\nu-1} \sqrt{R(0)}, \quad (4.45)$$

provided $\nu > 1$. For $0 < \nu \leq 1$, $a_\nu(z)$ diverges as z goes to 0 from the right.

4.5 Array Structure and Superconductivity

The problem with the preceding sections is that there is nothing interesting about the displacement of a dot from a fictive lattice position. The physically meaningful quantity is the

Josephson coupling between dots, which is some function of the distance between them. Because the physics is driven by the nearest neighbor couplings, it makes sense to split the analysis into the nearest neighbor distances and everything else. For this reason, the best measure of the degree of disorder in the array is not the variance of the displacements, $R(0)$, but probably the variance of the Josephson coupling between neighboring niobium islands. This quantity, in turn, will be related to the distribution of nearest neighbor distances.

This section will relate the displacements ($\{X_i\}$ and $\{Y_i\}$) to other quantities that characterize the disorder of the array, and are hopefully more relevant to the superconductivity. I'll start by determining the distance between pairs of dots when the two are nearest neighbors in the reference lattice, Section 4.5.1.

A key focus of future work will be addressing the effect of any long-range structure in the disorder on the inhomogeneous superconducting state. For example, by tuning the correlation length or by adjusting the smoothness parameter ν . This will require matching the local disorder across samples with different displacement correlations. To do this we can change the displacement magnitude ($R(0)$) and possibly the initial lattice constant in order to create arrays with matching nearest neighbor distance histograms, discussed in Section 4.5.2.

4.5.1 Nearest Neighbor Distance

Define the random variable D_i to be the distance between dots i and $i + 1$. The right edge cases, D_{N-1} and so on, are undefined. The situation is sketched in Fig. 4.4, showing two red dots and their respective lattice positions as tan crosses. The x and y displacements are also indicated. From the diagram, the bond distance is defined as:

$$D_i = \sqrt{(a + X_{i+1} - X_i)^2 + (Y_{i+1} - Y_i)^2}. \quad (4.46)$$

As defined, the $\{D_i\}$ only cover the “horizontal” bonds. The vertical bonds, for instance the distance between dots i and $i + N$, have the roles of X and Y interchanged. The statistics of the two sets of bonds will be identical, because the x and y displacements are independent and identically distributed.

Unfortunately, there's no straightforward way to evaluate the mean or variance of D_i ,

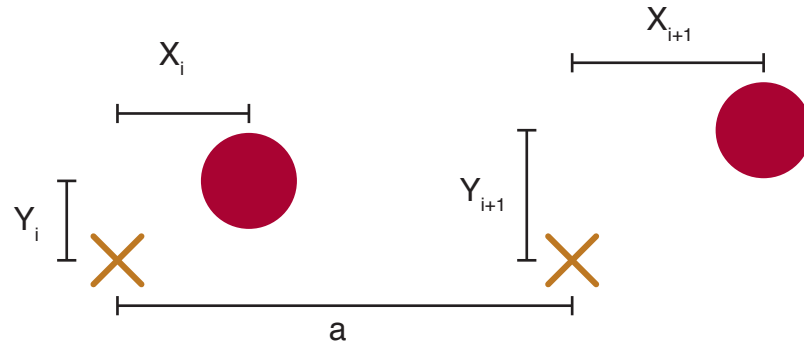


Figure 4.4: A diagram showing the position of two dots (red) and their original lattice positions (tan crosses). The displacements X and Y are indicated.

due to the square root. What I'd like to have is an expression for $\langle D_i \rangle$ in terms of $\text{Var}(X)$ and $\text{Cov}(X_{i+1}, X_i)$. Inverting this expression would allow me to match the local structure between samples with different long-range structure.

For instance, if I want two samples to have an identical spread of nearest-neighbor distances, the sample with correlations would require a larger $R(0)$ than the uncorrelated sample. I'm going to work around this three ways. First, I'll think about only the horizontal separation, because it's easy to calculate. Then we can set bounds on $\langle D_i \rangle$ using Jensen's inequality. Finally, I can calculate it numerically in a few cases by taking advantage of the independence of the x and y displacements, as well as their normal distribution.

Horizontal Separation

The horizontal distance between dots i and $i + 1$ is:

$$H_i = a + X_{i+1} - X_i; \quad (4.47)$$

since H seems like a good name for this variable. The average horizontal distance is the lattice period, $\langle H \rangle = a$, which I had expected to be true for the Euclidean distance (D) as

well. The variance is:

$$\begin{aligned}\text{Var}(H_i) &= 2 (\text{Var}(X) - \langle X_{i+1}X_i \rangle), \text{ or} \\ &= 2 (R(0) - R(a)).\end{aligned}\tag{4.48}$$

We saw in Section 4.2 that $R(a)$ is $R(0)$ times a number smaller than one. For a specific example, with squared exponential correlations we have $R(a) = R(0)e^{-a^2/l^2}$. If my lattice period is 500 nm and the correlation length is 5 μm , then $R(a) = 0.99R(0)$. If I want this sample to have the same variance of horizontal nearest neighbor separations as an uncorrelated sample (where $R(a) = 0$), I will have to increase $R(0)$ by a factor of one hundred.

Bounds on $\langle D \rangle$

Jensen's inequality relates the expectation value of a function of a random variable and the value of the function evaluated at the expectation value of the random variable.[68] The usual statement is that if X is a random variable and f is a convex function, then:

$$f(\langle X \rangle) \leq \langle f(X) \rangle.\tag{4.49}$$

The corollary, when f is a concave function, is:

$$\langle f(X) \rangle \leq f(\langle X \rangle).\tag{4.50}$$

To apply these inequalities, note that we can write $D = f(X_{i+1} - X_i, Y_{i+1} - Y_i)$, with $f(x, y) = \sqrt{(a+x)^2 + y^2}$, which is everywhere convex. Then we conclude that:

$$f(\langle X_{i+1} - X_i \rangle, \langle Y_{i+1} - Y_i \rangle) \leq \langle D \rangle,\tag{4.51}$$

or $\langle D \rangle$ is greater than or equal to the lattice period. Alternately, I can define $D = f(D^2)$, with $f(x) = \sqrt{x}$. This function is everywhere *concave*, so the inequality is:

$$\langle D \rangle \leq \sqrt{a^2 + 4 (\text{Var}(X) - \langle X_{i+1}X_i \rangle)},\tag{4.52}$$

where the quantity in the square root is simply $\langle D^2 \rangle$, assuming that X and Y are identically

distributed, so that $\text{Var}(X) = \text{Var}(Y)$ and $\langle X_{i+1}X_i \rangle = \langle Y_{i+1}Y_i \rangle$. Taking that all together, we have:

$$a \leq \langle D \rangle \leq \sqrt{a^2 + 4(\text{Var}(X) - \langle X_{i+1}X_i \rangle)}. \quad (4.53)$$

Since the variance of D is $\langle D^2 \rangle - \langle D \rangle^2$, it is bounded from below by 0, and from above by $4(\text{Var}(X) - \langle X_{i+1}X_i \rangle)$.

Direct Computation

While there is not a straightforward way to evaluate $\langle D_i \rangle$, there certainly is a tedious way. We can brute-force perform the integral to evaluate the average, provided we can construct the joint probability density function of X_{i+1}, X_i, Y_{i+1} , and Y_i . I've already chosen to make the x and y displacements normally distributed and independent of one another. So they each follow a bivariate normal distribution. If two random variables, X and Y , follow a bivariate normal distribution and both have mean zero and variance σ^2 , their joint probability density function is:

$$f(x, y) = \frac{1}{2\pi\sigma^2\sqrt{1-\rho^2}} \exp\left[\frac{-(x^2 + y^2 - 2\rho xy)}{2\sigma^2(1-\rho^2)}\right], \quad (4.54)$$

where $\rho = \text{Cov}(X, Y) / \sigma^2$ is the correlation coefficient. For our example, $\rho = R(a)/R(0)$ and $\sigma^2 = \text{Var}(X) / a^2$. With that defined, we can find the average nearest-neighbor distance by evaluating

$$\begin{aligned} \langle D_i \rangle = a \iiint \int dx_{i+1} dx_i dy_{i+1} dy_i & \sqrt{(1 + x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2} \\ & \times f(x_{i+1}, x_i) f(y_{i+1}, y_i). \end{aligned} \quad (4.55)$$

I've made the integral dimensionless by referring all lengths to the lattice constant. I should note that it's not really possible to directly evaluate the integral in Eq. 4.55. Instead we estimate $\langle D \rangle$ by Monte Carlo, by generating many bivariate normal random numbers and averaging the calculated D . In Table 4.1, I've done this for a few values of σ and ρ , as well as calculating the standard deviation of D . Both $\langle D \rangle$ and $\text{std}(D)$ increase with increasing σ , while increasing ρ decreases both. As required by the inequalities in Eq. 4.53, $\langle D \rangle$ is larger

than a . I've taken to calling this effect *dilation*, and it is less than ten percent for reasonable values of $\text{std}(D)$. Interestingly, $\text{std}(D)/a$ is larger than σ for small enough correlation.

$\langle D \rangle / a$	ρ					
	0	0.1	0.3	0.5	0.9	0.99
σ						
0.01	1.000	1.000	1.000	1.000	1.000	1.000
0.1	1.010	1.009	1.007	1.005	1.001	1.000
0.3	1.096	1.086	1.066	1.046	1.009	1.001
0.5	1.282	1.253	1.195	1.136	1.025	1.003
0.9	1.832	1.762	1.614	1.452	1.086	1.008
1	1.988	1.907	1.737	1.549	1.107	1.010
$\text{std}(D)/a$	ρ					
	0	0.1	0.3	0.5	0.9	0.99
σ						
0.01	0.014	0.013	0.012	0.010	0.005	0.002
0.1	0.141	0.133	0.118	0.100	0.045	0.014
0.3	0.399	0.381	0.341	0.292	0.134	0.042
0.5	0.597	0.574	0.522	0.457	0.221	0.071
0.9	0.940	0.900	0.814	0.716	0.381	0.127
1	1.025	0.981	0.885	0.776	0.417	0.141

Table 4.1: The mean (upper) and standard deviation (lower) of the nearest neighbor distance, normalized by the initial lattice constant a , for several values of the standard deviation (σ) and correlation coefficient (ρ) of the positional displacements. The mean of D is always larger than 1, increasing with growing σ or decreasing ρ . The same trend holds in the standard deviation of D .

4.5.2 Matching Local Properties

When creating the array I control the initial lattice constant, a , and the standard deviation of the displacements, $a\sigma$. (I've also referred to the latter as $\sqrt{R(0)}$ and $\sqrt{\text{Var}(X)}$.) The goal is to find one or more equations that determine a and σ from the desired mean and standard deviation of D . This depends on the choice of whether or not to correct for the

dilation.

First, if I choose to ignore the dilation the only equation needed is $\text{std}(D) = fa$, where f defines the standard deviation of D relative to the lattice constant. For convenience I will define the dilation function N as:

$$aN(\sigma, \rho) = \langle D \rangle, \quad (4.56)$$

which just labels the integral in Eq. 4.55 as N . The equation to solve is:

$$\begin{aligned} f^2 a^2 &= \langle D^2 \rangle - \langle D \rangle^2, \\ &= a^2 [1 + 4\sigma^2(1 - \rho)] - a^2 N^2, \\ f^2 + N^2 &= 1 + 4\sigma^2(1 - \rho). \end{aligned} \quad (4.57)$$

This is an implicit equation for σ , because the dilation function depends on it and ρ . With the lattice constant set, ρ is simply a number; it has no dependence on σ or anything else.

Correcting for the dilation requires a second equation, because $\rho = R(a)/R(0)$ depends on the lattice constant. I'm assuming that the properties of the correlation function, namely ν and l , are kept constant. Defining l in terms of a makes the dilation correction pointless. The system to solve is:

$$\alpha = \langle D \rangle \quad f\alpha = \text{std}(D). \quad (4.58)$$

Where α is the desired average nearest neighbor distance and f defines the standard deviation of D relative to α . Introducing N as before, the system works out to:

$$\alpha = aN \quad (1 + f^2)N^2 = 1 + 4\sigma^2(1 - \rho). \quad (4.59)$$

In the special case of uncorrelated disorder ($\rho = 0$) the two equations decouple and there is only one equation to solve, the other being a trivial choice of lengthscale.

As an implementation note, this is one area where continuous vs. discrete will bite you. You expect $\rho = R(a)/R(0)$, but it doesn't quite work out. The true correlation between nearest neighbor displacements ends up being a bit smaller, depending on ν . The thing to

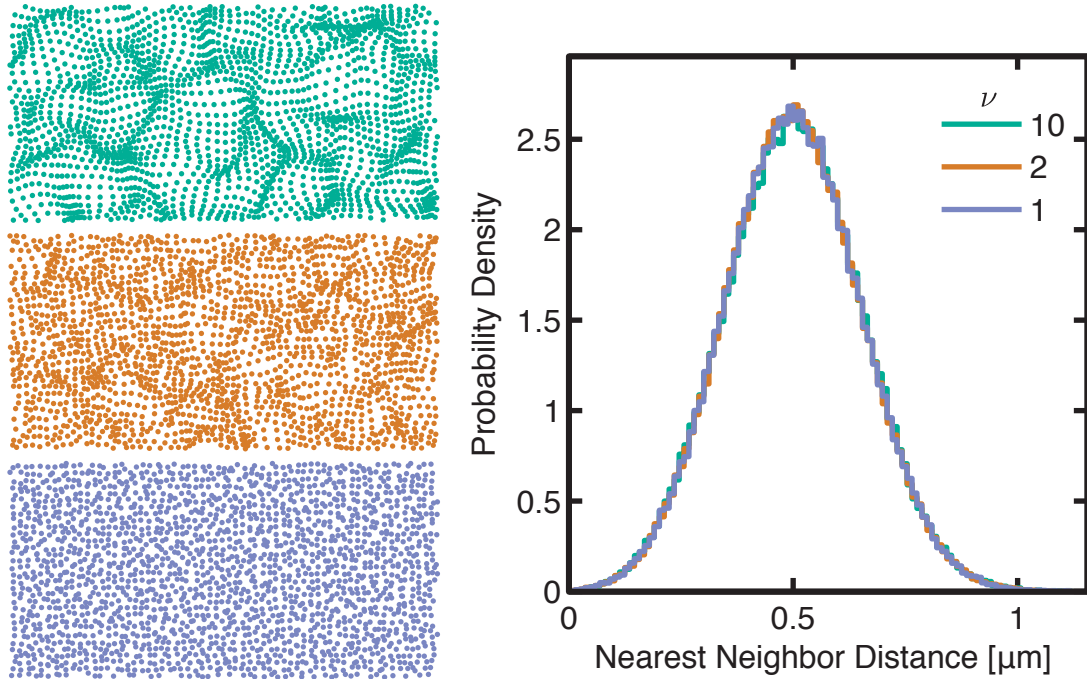


Figure 4.5: Three arrays similar to Fig. 4.3, after matching the local structure with $\text{std}(D)/a = 0.3$. Comparing this figure to Fig. 4.3, the arrays with $\nu = 10$ and 2 are more disordered here, while the $\nu = 1$ array is less disordered. This is evident both in the images at left and in the histograms at right. This trend is consistent with the data in Table 4.1, which shows that $\text{std}(D)/a$ is larger than σ for small enough correlation.

do is auto-correlate the filter kernel (a_ν) and pick off the value at one lattice constant.

As an example, I've made the local structure correction for the three arrays initially shown in Fig. 4.3, with $l = 2 \mu\text{m}$, $a = 500 \text{ nm}$, and $\nu = 10, 2$, and 1. The corrected arrays, set to have $f = 0.3$, are shown in Fig. 4.5. When I make the dilation correction by solving the system in Eq. 4.59, I get $\sigma = 0.95, 0.66$, and 0.24 for $\nu = 10, 2$, and 1 respectively. In each case the corrected lattice constant is 473 nm , so the dilation was six percent. If I skip the dilation correction by solving Eq. 4.57, the solutions are $\sigma = 0.85, 0.59$, and 0.22 in the same order as before.

Chapter 5

SQUID Measurements of Dot Arrays

5.1 Methods

The dot arrays are niobium islands placed on a patterned gold film by electron beam lithography. The gold film is 10 nm thick on top of a 0.4 nm titanium sticking layer, all on a silicon substrate. The metals are patterned into a four point structure for transport measurements, either by optical or electron beam lithography. Once the metals are deposited, resist spun on, exposed, and developed. Argon ion milling is used to clean the surface *in situ* before evaporating niobium onto the sample. The resulting niobium film is granular, with 50 to 100 grains per dot. The granularity has a significant impact on the physics of superconductivity in the arrays.[74]

I measured five of these samples. The first two samples were uniform triangular lattices with periods of 500 and 550 nm, pictured in Fig. 5.1(d) and (e). The dots are always 260 nm diameter, so that corresponds to spacings of 240 and 290 nm, respectively. These two samples were fabricated for transport experiments, so the gold pattern is long and skinny: 120 μm by 30 μm . Finally, there is some problem with the gold in these two samples, which is apparent both the optical image (Fig. 5.1(d,e)) and in the superconducting properties (Fig. 5.4(d,e)).

The next three samples were all disordered, with a 120 μm square active area. The gold pattern also has leads for a four point pattern transport measurement, but the the square active area is not ideal for determining the conductivity from the measured resistance. One




Figure 5.1: Optical photographs of the five arrays. (a) The uncorrelated sample, DS1, has dots placed randomly with a uniform density. (b) DS2 is a square lattice with correlated disorder, correlation length $3\sqrt{2} \mu\text{m}$. (c) DS3 is also a disordered square lattice, with correlation length $5\sqrt{2} \mu\text{m}$. (d) A uniform array, with dots on a triangular lattice of period 500 nm. (e) As (d), with period 550 nm. The arrays in (d) and (e) have patchy discolorations due to some schmutz in the gold layer.

possible fix to enable combined transport and SQUID measurements would be to use the Van der Pauw configuration.

One sample, DS1, is completely uncorrelated — dots were placed over a $150 \mu\text{m}$ square area with a uniform random distribution. The density is equivalent to a 500 nm period square lattice. The other two samples are disordered from a 500 nm period square lattice, following the procedure described in Chapter 4. Both have squared exponential correlations of the displacements. One, DS2, has a correlation length of $3\sqrt{2} \mu\text{m}$. The second, DS3, has a correlation length of $5\sqrt{2} \mu\text{m}$. The $\sqrt{2}$ factor is for consistency with Chapter 4; the displacements were made using a gaussian filter with equation given by $a(z) \propto \exp(-z^2/l^2)$, in contrast to Eq. 4.39, which has $a(z) \propto \exp(-2z^2/l^2)$. So the resulting correlation length

is larger by $\sqrt{2}$.

As I've mentioned, the uniform arrays have a triangular lattice, while the disordered arrays have a square lattice. Based on the previous results,[74] a triangular lattice should have stronger superconductivity —more screening and higher T_c — than a square lattice with the same spacing, because it has 6 nearest neighbors instead of 4. Future work (Section 5.3) will have uniform square lattice arrays to compare against the disordered structures.




Figure 5.2: Scanning electron microscopy images of two disordered arrays, showing that the lithography places the dots where intended. (a) The completely uncorrelated sample, DS1. (b) The sample with correlation length $5\sqrt{2}$ μm , DS3. Images courtesy Jared Schwede.

Two images of the disordered dot arrays taken by scanning electron microscopy (SEM) after the cooldown are shown in Fig. 5.2. These images confirm the lithography, that the dots are located where we wanted them. Even better, the SEM data shows that the capacitive touchdowns do not noticeably damage the arrays. I took several hundred touchdown curves on each sample, and expected to see isolated spots of damage in the SEM image corresponding to each touchdown location.

Some analysis of the three disordered structures is presented in Fig. 5.3. Panels a–c show 15 μm square sections of the structures. Each dot represents the center of a niobium island; the dot size is not to scale. Figure 5.3(g) shows the pair correlation function for each of the three disordered arrays, while panel h shows the histogram of nearest neighbor distances. Lastly, in panels d–f I show the local neighborhood of the array. For a random selection of

dots, plot the relative location of every other dot within a certain distance (here $2.5 \mu\text{m}$). This just the pair correlation function before radial averaging. I find it easier to understand; it looks like a fuzzed over square lattice and makes it very clear that DS2 and DS3 are weakly disordered.

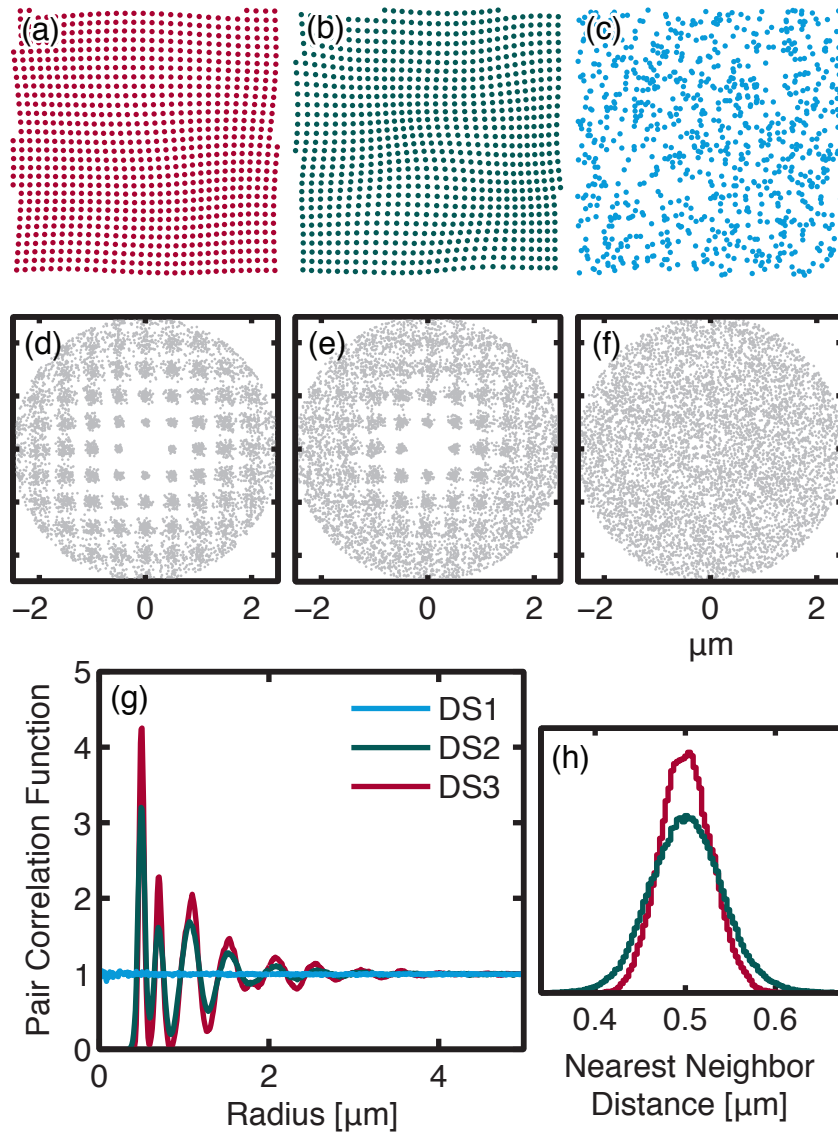


Figure 5.3: Structure analysis of the three disordered arrays. Top row: 15 μm square sections of each pattern. Each dot is the center of a niobium island, size not to scale. (a) DS3, with correlation length $5\sqrt{2}$ μm . (b) DS2, correlation length $3\sqrt{2}$ μm . (c) DS1, with a uniform random density of islands. Middle row: the local region of each array. Each gray dot represents a niobium island within 2.5 μm of 100 randomly selected islands. (d) DS3 (e) DS2 (f) DS1. (g) The pair correlation function for each structure. (h) Histogram of nearest neighbor distances, defined relative to the unperturbed lattice as in Section 4.5.1, for DS3 and DS2.

5.2 Results and Discussion

Base Temperature

In Fig. 5.4, I show SQUID susceptometry images of each of the five samples. All are taken at the base temperature of the ^3He microscope, roughly 0.4 K.

The uniform arrays don't have a uniform superconducting response, due to a defect in the gold layer that is clearly visible in Fig. 5.1(d) and (e). The dot array is undisturbed in those regions, which can be verified with a sufficiently powerful optical microscope. What is likely occurring is that contamination in the gold layer locally changes the normal coherence length. This would weaken the inter-island Josephson couplings, resulting in weaker superconducting screening.

The second observation in the uniform arrays is the comparison between the two samples. The 500 nm period sample has both a higher transition temperature (2.9 vs. 1.5 K) and a larger superconducting response (250 vs. 180 Φ_0) than the 550 nm sample. This is consistent with the transport data, at least for the transition temperature.[74]

The completely random sample (panel a) is barely superconducting at all, and no inhomogeneity is observed. The dominant feature in the image is a denim pattern due to correlated noise. From this data it is not possible to tell whether the sample has a uniform SQUID response, or if the inhomogeneity is simply under the noise floor. The transition temperature is 1.05 K.

The DS2 sample (panel b) has stronger superconducting response and a higher transition temperature (1.5 K). There is also some apparent inhomogeneity, the white lumps near the center of the square. In addition, notice the small tails extending from the top and bottom. Compare that to the optical image, Fig. 5.1(b), where you can see the niobium dots (a square of gray fuzz) extending beyond the gold layer visible underneath. What I take from this is that the data are not showing the diamagnetism of thousands of tiny niobium islands, but are imaging the screening currents flowing in the gold, because there is only response when there is both niobium and gold. This argument is strengthened by the data on the DS3 sample (panel c), where the tails are more pronounced. In addition, a lithography error resulted in a small patch of continuous niobium film at the corner of the dot array, visible in Fig. 5.1(c) as a white-yellow spot in the top-right corner. This patch of niobium was found




Figure 5.4: Susceptibility measurements of each array at base temperature, roughly 0.4 K. The lateral scale is identical across subfigures, but the colorscale is independent. White is stronger superconducting response, black is no superconductivity. (a) DS1 (b) DS2 (c) DS3. The dark square region in the middle was damaged by scanning in contact. (d) Uniform triangular array with period 500 nm. At bottom right is a sketch of the SQUID geometry to scale, showing the field coil (blue) and pickup loop (red). (e) Uniform triangular array with period 550 nm.

around 10 μm away from the top-right corner of the sample, with no measurable diamagnetic response found in between. Similar niobium patches are found near each of the other disordered samples, always on or near an alignment mark. My hunch is that imaging the alignment mark, which uses the electron beam as an SEM, exposed the resist there. This would result in a patch of continuous film near the four corners of the array. These patches were fortuitous; without them I would never have found the uncorrelated sample.

Now focusing on the DS3 sample shown in Fig. 5.4(c), the response is stronger yet, and the transition temperature higher still (1.8 K). The inhomogeneity is very apparent in this sample, for example the two bright white dots near the bottom each with a signal of roughly

$120 \Phi_0/A$. There is also substantial inhomogeneity in the region along the right edge. Of course, the most prominent feature in the data is the darker region in the middle. This was damaged by accidentally scanning in contact. For nine hours. Even so, the response in the damaged areas (roughly $80 \Phi_0/A$) is larger than the peak response over the DS2 sample.

A more detailed comparison of the DS2 and DS3 samples is in Fig. 5.5, which shows two identically sized regions of data over each sample, taken with the same nominal scan height of $4.9 \mu\text{m}$. We again see that the sample with a longer correlation length has both more obvious inhomogeneity and also a larger average magnitude. This is underlined by Fig. 5.5(c), which shows histograms of the data in each scan (the pixel values).

The data at base temperature suggest that the SQUID signal (and its variance) varies inversely with the correlation length. Perhaps at smaller correlation length the inhomogeneity is averaging out. In some sense, the correlation length can tune the system from one where the SQUID is a local probe to a regime where the SQUID is only sensitive to the disorder-averaged field.

Once caveat to the preceding discussion is the effect of the scan height. For each sample I defined the plane of the substrate with capacitive touchdown detection, as described in Appendix A. To avoid damaging the arrays, these touchdowns are done at a few points around the array, which means that I'm truly finding the location of the *substrate*. The very top of the niobium islands is some 100 nm higher than that. The data are taken by (open-loop) scanning in a plane a specified distance, the scan height, above the plane of the substrate. The data in Fig. 5.4 were taken with scan heights of $1.9, 0.8, 0.8, 1.1,$ and $1.1 \mu\text{m}$, respectively panels a–e. There is an additional offset height, corresponding to the height of the SQUID pickup loop above the surface when the sensor tip is touching the substrate. This offset height, of order $1 \mu\text{m}$, should be the same for samples on the same silicon chip in the same cooldown, but can change from cooldown to cooldown. Specifically, there should be one offset height for the two uniform arrays and another for the three disordered arrays.

From both Appendix A and the touchdown data in Fig. 5.8, we know that the scan height has a large effect on the observed magnitude of the SQUID signal. It will also affect the observed inhomogeneity in two ways. First, the height will change the spread of the data directly by the height dependence of the SQUID response. If we imagine touchdown curves taken on uniform infinite films with different Pearl lengths, the difference between

the raw SQUID signal in Φ_0/A is a function of height. So if the scanning data corresponds to a distribution of Pearl lengths, the width of the distribution of raw data will depend on the height. The second and more interesting (if problematic) reason is that the effective lengthscale being probed by the SQUID is itself is a function of height, as we saw in Section 2.3.2.




Figure 5.5: Comparison of the two arrays with correlated disorder: (a) DS2, and (b) DS3. Both images are the same size, with the same nominal scan height of $4.9 \mu\text{m}$, but independent colorscales including offset. (c) Histograms of the data in (a) and (b), showing that the sample with shorter correlation length (DS2, panel a) has both a smaller average SQUID signal and a smaller variance than the longer correlation length sample.

Temperature Dependence of DS3

In addition to the SQUID images at base temperature, I also took a lot of temperature series data. That is, you stabilize the temperature at some setpoint and take a scan, then change the temperature and take another scan with identical parameters. Here I'm going to focus on only one of these temperature series, taken on a 61 by $71 \mu\text{m}$ region over DS3. This is the very same damaged region mentioned earlier, and that's no coincidence. What happened was the ^3He pot ran dry while taking data at high temperature (4 K or so), which caused the SQUID to scan in contact. This happened in the middle of an overnight run, so it was not caught and stopped for quite some time. I was taking data at such high temperatures because I already had data showing the persistence of superconductivity in small islands.

All that aside, the data shown here were taken before the damage occurred.

The temperature series is summarized in Fig. 5.6, which shows the region at 0.42 K (panel a) and at 1.2 K (panel b). The data at 0.42 K are also used in Fig. 5.5(b). Figure 5.6(c) shows the data as a function of temperature; the gray region bounds the range of pixel values and the red line is the average. The diamagnetic response does not go to zero with finite slope, instead there is a tail that is probably indicative of phase fluctuations. The other feature to note is that the width of the gray region is larger at lower temperatures, and decreases with increasing temperature until it is set by the system noise. Taken together, this could suggest that the superconducting response uniformly scales down. That is, the only effect of the disorder is to introduce a distribution of local superfluid densities with identical temperature dependence. Of course, that is belied by Fig. 5.6(b), which does not look the same as Fig. 5.6(a).




Figure 5.6: Summary of the temperature dependence of the SQUID susceptibility over DS3. (a) Data at 0.42 K. (b) Data at 1.2 K. (c) The entire temperature series, showing the span of the data (gray region) and the average value (red line). Note that the vertical height of the gray region decreases along with the average value.

What's also cool and surprising is that the regions with the largest SQUID signal at base temperature don't have the highest transition temperature. This can be seen in Fig. 5.7, which again shows the region at 0.42 K (panel a). But now in panel b I show the highest temperature attained, 1.9 K, rather than 1.2 K as in Fig. 5.6(b). In panels a and b I have circled three regions of interest. Note how the region indicated by the green circle has the largest signal (is most white) in panel a, but is completely non-superconducting (black) in

panel b. In contrast, the two regions in the red and yellow circles are still superconducting at 1.9 K, and at 0.42 K are merely above average. This can be seen in panel c, which plots the difference between the average inside the region and the average of the entire scan as a function of temperature.

This is surprising because at the simplest level both the magnitude of the diamagnetic response and the transition temperature should be controlled by the local density of dots. Regions with an enhanced density have smaller distances between islands, hence they have stronger inter-dot couplings. The nearest neighbor coupling controls both the effective Pearl length and the transition temperature, see equations in chapter 4, from the LAT paper.

However, this neglects the fact that a region of higher density will necessarily be surrounded by a region of lower density. This is a consequence of how the arrays are constructed. The ring of weaker coupling can serve to suppress the transition temperature of the island inside, due to the fact that the superconducting phase in each dot is stabilized by its coupling to the neighborhood. At low temperatures, phase coherence is well-established throughout the array, so the island of enhanced density has a larger superfluid density, as expected. But at higher temperatures, the stabilizing effect of the rest of the array is weakened. The sparse ring surrounding the dense island has lower phase rigidity, allowing the island to decouple from the rest of the array and fluctuate itself out of the superconducting state.

Perhaps that is a plausible explanation for this observation. But it is also an argument for the impossibility of isolated superconducting islands, which are observed. So there's that.

Touchdown Curves

In Fig. 5.8, I show touchdown curves taken at base temperature over four of the five samples. The missing one is the uniform array with 550 nm periodicity, which was destroyed before I had the chance to take any touchdown data. At first glance, the data are consistent with the picture drawn by the scanning images in Fig. 5.4. The non-disordered array has substantially more diamagnetism than any of the disordered arrays, and among the disordered arrays the response tracks with the correlation length. But a closer look reveals




Figure 5.7: Isolated superconducting islands near T_c in DS3. (a) Data at 0.42 K, with three regions of interest indicated by colored circles. The region inside the green circle at middle-left has the largest response in the scan, over $35 \Phi_0/A$. (b) The same region at 1.9 K, with the three regions also indicated. Most of the sample is non-superconducting, except the regions in the yellow and red circles, which were not the regions with strongest superconductivity at base temperature. (c) Plot of the average value inside each region, minus the average of the entire scan, at each temperature.

discrepancies. The largest signal observed in a touchdown curve over the DS3 sample is around $90 \Phi_0/A$. But data larger than that (up to $120 \Phi_0/A$) are seen in Fig. 5.4(c). This could be explained by the sensor being closer to the sample surface during the scan than the touchdown. This is ruled out by the capacitance data taken simultaneously during the scan, which indicates that the sensor tip is not touching the sample. At touchdown the sensor tip is touching the sample, by definition. Another possible source of the discrepancy could be lateral drift. The intended touchdown point was the peak in the center bottom of Fig. 5.4(c), so missing this point at the beginning or slowly drifting off of it due to creep in the x and y benders would result in a smaller than expected signal. I estimate this drift would have to be 5 to 10 μm in order to explain the discrepancy. This is larger than I expect, but it is not impossibly large. Beyond that, the scan and the touchdown curve were taken with the same excitation magnitude and frequency, so an appeal to nonlinearity is moot.

The second major discrepancy is the data on the uncorrelated sample. The scans show very weak superconductivity, with SQUID signals roughly $2 \Phi_0/A$. But the touchdown response is up to $30 \Phi_0/A$. The scanning data were taken at a nominal scan height of 2 μm .

Based on the touchdown, the signal at that height should be roughly $17 \Phi_0/A$, almost ten times higher than observed. Turning that around, $2 \Phi_0/A$ of signal would imply a scan height of just over $10 \mu\text{m}$ based on the touchdown curve. I have no explanation for this.

The final upsetting discrepancy is a handful, perhaps five to ten, touchdowns on DS2 with much larger response than not just the scanning data, but also all of the other several hundred touchdowns. The anomalous data are upwards of $100 \Phi_0/A$, and can only be explained by the sensor getting $5 \mu\text{m}$ closer to the sample than usual. This suggests that the SQUID wasn't touching down at the tip in all of the other data. But I don't even know. I also have no explanation for this, but it calls into question every conclusion otherwise made about the differences between anything.

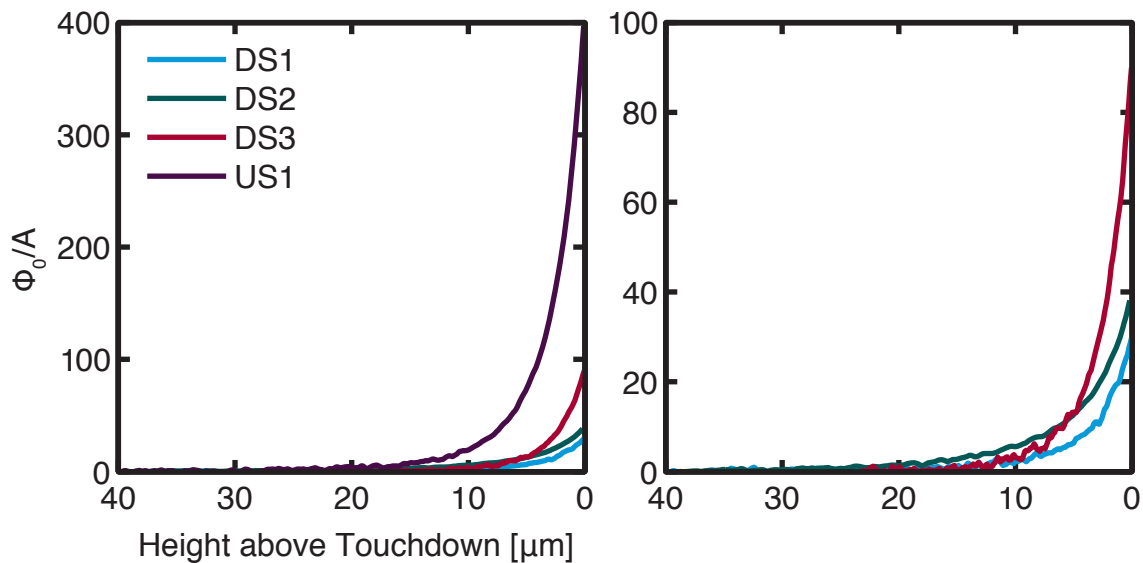


Figure 5.8: Touchdown curves at base temperature. Left: All four samples with touchdown data; US1 is the 500 nm period uniform array. Right: Only the three disordered arrays. Note how the data for DS3 cross over the other two curves. This is inconsistent with Fig. 2.3, indicating that an effective Pearl length does not adequately describe the response.

Considering the discrepancies just noted, it is clear to me that the offset height is not under control. It might not be consistent between the three disordered samples, and it may even change from touchdown to touchdown, a disheartening prospect indeed. Nevertheless, I went ahead and fit the touchdown curves in Fig. 5.8 to the two-coil model described

in Section 2.3. The results are summarized in Table 5.1, showing the fitted Pearl length and the offset height. The field coil radius (a) is fixed at $7.58 \mu\text{m}$, and the pickup loop radius (p) is fixed at $2.55 \mu\text{m}$.

The model fits the data on the uniform (US1) and uncorrelated disorder (DS1) samples reasonably well. But the model fails to describe the two samples with correlated disorder, DS2 and DS3. This can be seen from the unrealistic values of the fitted offset height and Pearl length. It can also be seen from the way the DS3 data crosses over the other two, which is inconsistent with simply having a smaller Pearl length based on Fig. 2.3. The failure of the model indicates that the effect of the disorder cannot be described by an effective Pearl length.

Sample	Pearl Length [μm]	Offset Height [μm]
DS1	102	1.89
DS2	2.65	8.5
DS3	64	0.16
US1	4.1	1.24

Table 5.1: The results of fitting the touchdown curves in Fig. 5.8 to the two-coil model as described in the text. The fitted offset heights for the two samples with correlated disorder (DS2 and DS3) stretch the bounds of plausibility, possibly indicating a failure of the model.

5.3 Future Possibilities

There is a large body of literature that studies the destruction of the superconducting state by disorder, but very little of it is spatially resolved.[59] The dot arrays provide a model system to study this question, because it is simple to create samples at any desired level of disorder. It is also possible to change more subtle properties of the disorder, including the lengthscale and smoothness of correlations in the disorder. This allows us to not only observe the destruction of the superconducting state by disorder, but also to investigate how the lengthscale of the disorder modifies the superconductivity.

A good place to start will be to fabricate two sets of arrays. Each set will be a disorder series, one set with uncorrelated disorder and the other with correlated disorder. Because I'm

interested in how the long-range structure of the disorder affects the superconductivity and its destruction, it makes sense to match the short-range structure between corresponding samples in each of the two sets. For our purposes “short-range structure” means the standard deviation of the distance between nearest neighbors, as described in Section 4.5.1. Matching this quantity between arrays with different correlation structure is non-trivial, see Section 4.5.2, but it can be done.

More specifically, each set will have four samples with the standard deviation of the nearest neighbor distance being 0, 5, 10, and 25 percent of the lattice constant. The first sample (0 percent) serves as a reference uniform array; for fabrication reasons each set will be on a separate chip, and probably measured in different cooldowns. Each array will be 80 μm square, with a four-point pattern for transport measurements, and based on a square lattice with 500 nm period. One set will have uncorrelated displacements of dot positions, the other will have squared exponential correlations (Section 4.4.2) with a correlation length of 5 μm .

Following the procedure and notation in Section 4.5.2, the samples in the correlated set can be generated with $\sigma = 0.36, 0.71, \text{ and } 1.84$ for $f = 0.05, 0.1, \text{ and } 0.25$ respectively. The uncorrelated set has $\sigma = 0.035, 0.071, \text{ and } 0.18$, in the same order. An example of the correlated set is shown in Fig. 5.9, while the uncorrelated set is in Fig. 5.10. Each figure shows 15 μm square sections of the arrays (a–c) along with the characterization plots described previously.

I’ve elected to ignore the dilation correction, which refers to the fact that the average nearest neighbor distance is larger than the lattice constant. If you wanted to account for that, the correlated set would have $\sigma = 0.36, 0.72, \text{ and } 2.00$ with $a = 499, 498, \text{ and } 481$ nm for $f = 0.05, 0.1, \text{ and } 0.25$ respectively. The uncorrelated sample has $\sigma = 0.036, 0.071, \text{ and } 0.19$ with $a = 499, 498, \text{ and } 483$ nm. In both cases the dilation effect is at most three percent, and ignoring it simplifies things. But it is worth keeping in mind that a ten percent change in the lattice constant is significant, as seen in the difference between the two uniform structures and in the previous transport data.[74]

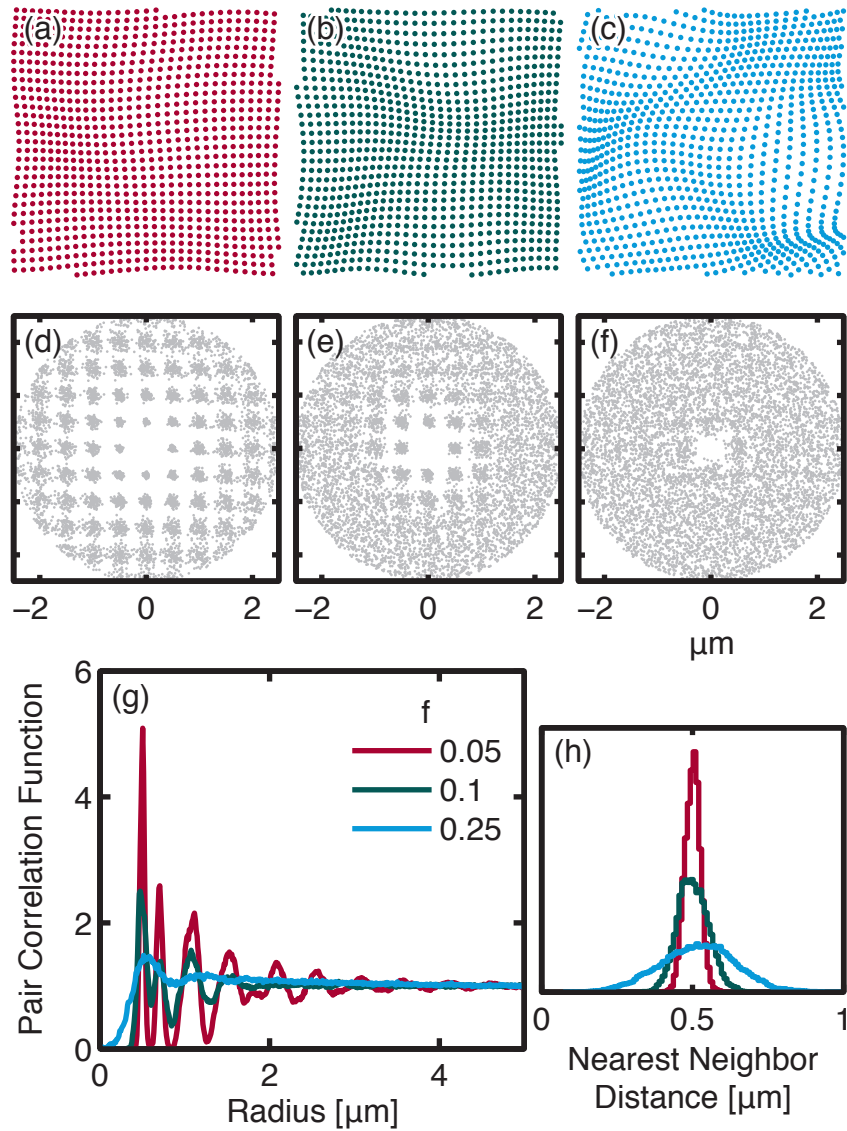


Figure 5.9: Structure analysis of the proposed disorder series with correlated disorder. Each panel is the same as Fig. 5.3, but now each array has squared exponential correlations with correlation length $5 \mu\text{m}$. What changes is the standard deviation of nearest neighbor distances, which are 5 (a,d), 10 (b,e), and 25 (c,f) percent of the lattice period of 500 nm .

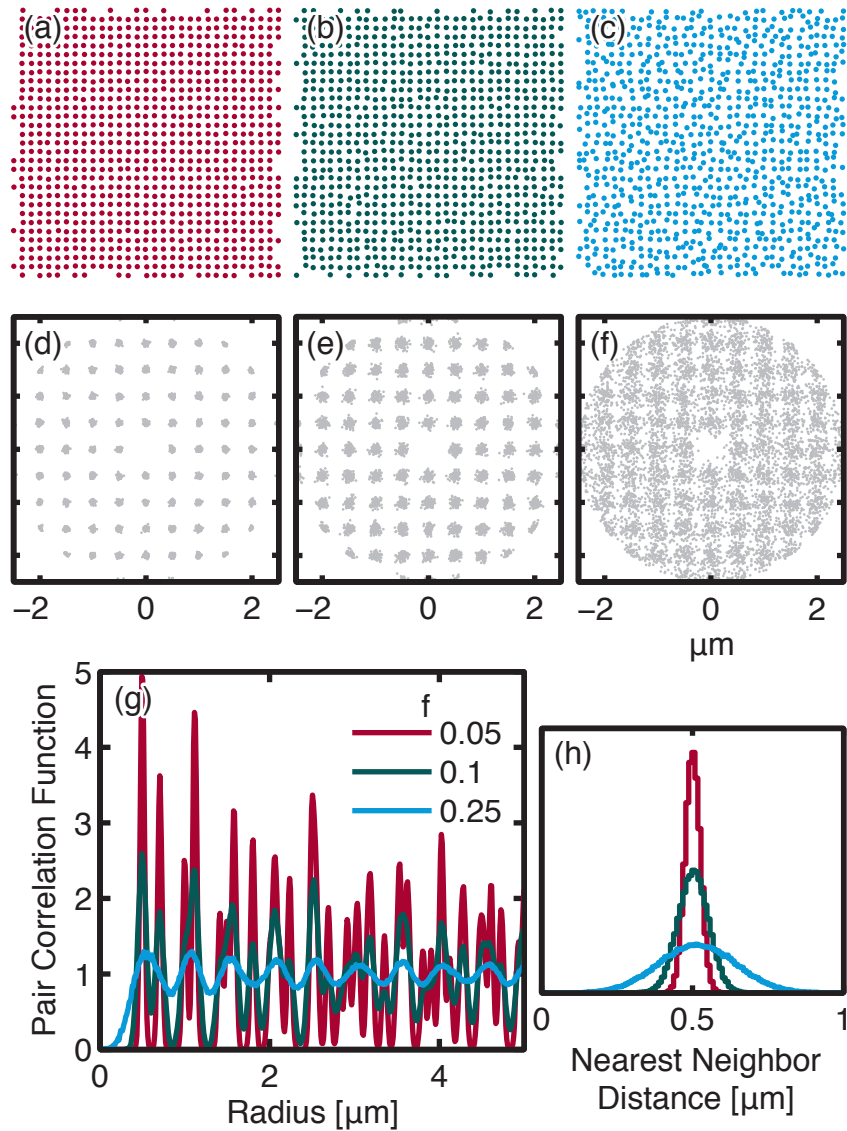


Figure 5.10: Structure analysis of the proposed disorder series with uncorrelated disorder. Each panel is the same as Fig. 5.3, but now each array has uncorrelated disorder. What changes is the standard deviation of nearest neighbor distances, which are 5 (a,d), 10 (b,e), and 25 (c,f) percent of the lattice period of 500 nm.

Appendix A

Capacitive Touchdown Detection

A.1 Introduction

Unlike other local probes, the scanning SQUID doesn't have a built-in way to detect the sample surface. For example, in STM measurements the tunnel current itself gives a very sensitive measure of the tip height above the sample surface. Force probe measurements have the Van der Waals force to tell them when they're too close to the sample. An incorporated height measurement allows true constant-height scanning; errors are canceled with negative feedback.

This lack of height measurement means that systematic uncertainty dominates in any attempt at precision, quantitative studies using the scanning SQUID. As we saw in Chapter 2, the sensor height above the sample is extremely important for determining the penetration depth, even in the limit where the penetration depth is much larger than the height. Height uncertainty is also a problem when determining the magnetic moment of dipoles[75] and ferromagnetic patches,[76] and for inferring the current flow in 2d structures.[77]

Now, perhaps all is not lost. We could add a parallel measurement of the height and use that for control. For example, we could glue a SQUID chip to a quartz tuning fork, incorporating a force measurement thereby. Another option would be to add an optical fiber to the SQUID tip to turn the space between the SQUID and sample into a Fabry-Pérot cavity. Now, there will be drawbacks to either of these examples, from increased system complexity to sample heating due to the laser light. Quite aside from that, a parallel measurement may

not give an absolute height measurement, since most distance measurements in the world are differential. And even if it did, it would be measuring the height to the tip or fiber, with some offset from that to the SQUID pickup loop itself. I am pessimistic about the possibility of calibrating that independently; can we ensure that the alignment angle, say, does not change at low temperature?

An interesting variation on this idea is to put the parallel measurement into the scanner itself. For example, there's the PicoCube™ from Physik Instrumente, which uses capacitors incorporated into the enclosed box of the scanning stage to offer three axis closed-loop scanning. Bear in mind this isn't really the same thing, since we're not measuring the distance to the sample, but rather the distance between two elements of the scanner. So really this is a method for eliminating drift and nonlinearity in the scan stage itself. This is nothing to sneeze at, but not a solution to the problem at hand. For example, we could not use it to maintain a fixed sample-probe separation while ramping the sample temperature, since the sample may wander off due to thermal expansion.

This seems like a good time to mention the current state of the art. With capacitive touchdowns, we can find the point where the tip of the squid first touches the sample surface. At that point, the sensor is some distance h_0 above the surface. This offset is defined by the SQUID tip geometry and alignment angle. We assume that this offset does not change across a sample, but may change between cooldowns and between samples in the same cooldown. Aside from that, we can get the penetration depth from touchdown curves, or from constant height measurements. They're both sensitive to the uncertainty in the bend calibration constant, and the latter is additionally affected by piezo drift. My experience has been that this can be a few nm over several hours, but is sensitive to the details of the recent history.

The lack of integrated height measurement is not all bad, since we are strongly encouraged, if not forced, to use open loop control. There are tradeoffs associated with the choice between open loop and closed loop scanning. Open loop control reduces system complexity and the need for a stiff scanner.

The scanner's mechanical rigidity sets its lowest resonant frequency, as well as affecting the achievable scan range. A stiffer scanner has a higher resonant frequency and smaller

scan range. Conversely, a less stiff scanner has a larger scan range and lower resonant frequencies. We consistently choose larger scan range at the expense of stiffness, due to the low spatial resolution of the SQUID. The scanner I used can scan an area $400 \mu\text{m}$ square. The resonance structure is unknown, sadly. The first resonance of a different scanner, with half the linear range, was 65 Hz.[78] So the first resonance of my current scanner is likely much lower. The usual way to measure this is by driving an axis with a heavily-attenuated function generator and looking for a noticeable vibration amplitude. Cliff described looking at the scanner under a microscope and watching for a speck to go blurry. This happens at room temperature, which is why we haven't done it for the current scanner. The silver electrodes are susceptible to diffusing through the piezo ceramic and shorting the bender. Typically a nickel underplating prevents this, but it also causes stray magnetic fields that were intolerable in a previous experiment.[79, 80]

Why is this important? The feedback electronics must have a bandwidth below the scanner's mechanical resonance frequency. This prevents the scanner from shaking itself apart with positive feedback. We imagine the sensor coming too close to the sample, so the controller pulls it away. But there is some overshoot and now the sensor is too far from the sample, etc. Now, that's just a description of a badly-tuned proportional controller, and doesn't have anything to do with the scanner resonance. But proper tuning of the controller depends on the scanner response. Crucially, on-resonance there is no proportional control tuning small enough to prevent this destructive over-oscillation. In electronics terms, the scanner resonance introduces a 180 degree phase shift into the feedback loop, turning negative feedback into positive feedback.

The answer is to add an integrator, to have a proper PI controller. But that amounts to limiting the bandwidth of the feedback electronics and we're back where we started. So what? The point is that the scanner resonance frequency limits the speed we can scan at. With resonance frequencies below 60 Hz, we might be looking at feedback bandwidths less than 1 Hz. This is intolerably slow.

Weighing all of these considerations, we have a *post-hoc* justification for using open loop control. But we still need to detect the sample surface, and can't always count on the magnetic signal directly. So we add a parallel measurement that is only useful for detecting when we've already touched the sample.

A.2 Capacitance

To find the sample surface, we use capacitive detection. The sensor is varnished to a thin metal shim, which is attached to the sensor mount with epoxy or more varnish. Attached at one end and free at the other, the shim is now a cantilever, which is what I usually call it. Below the shim is a copper pad, the base plate, that forms a small parallel-plate capacitor together with the shim. When the tip of the sensor touches the sample surface it deflects the cantilever, increasing the capacitance. The configuration is sketched in Fig. A.1(a). I've plotted an ideal touchdown (red line) in Fig. A.1(b). The capacitance is constant until touchdown, where it increases sharply. If only it were so easy. A typical touchdown is below it, in green.

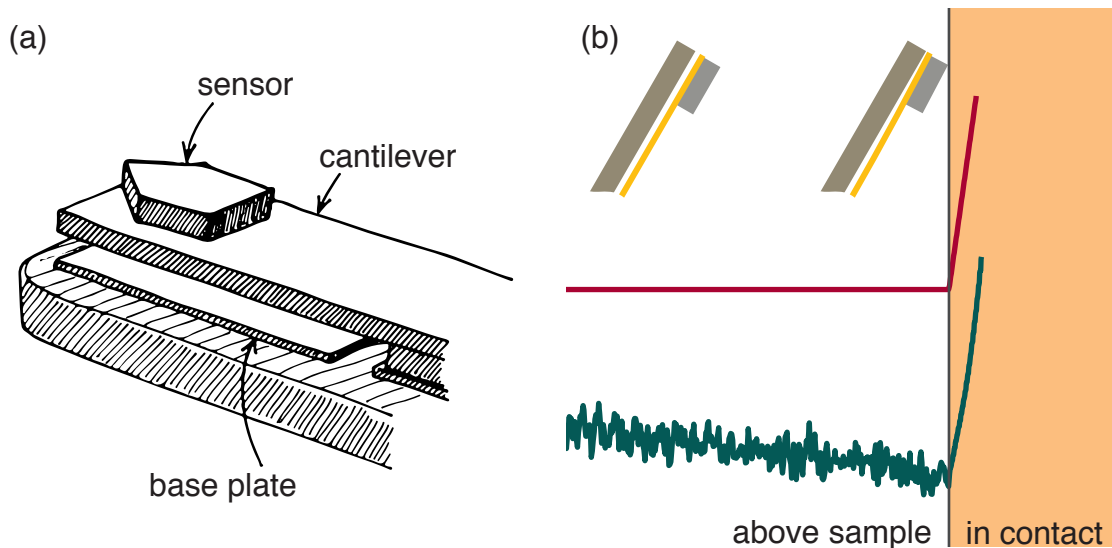


Figure A.1: Capacitive touchdown overview. (a) Sketch of the sensor mount, showing the sensor, cantilever, and base plate. Drawn by Clifford Hicks, reproduced with permission. (b) Concept of capacitive touchdown detection: the base plate-cantilever spacing, and hence capacitance, is constant as the scanner approaches the sample surface. Once the tip of the sensor touches the surface the capacitance increases sharply, signaling touchdown. Real data (green) doesn't match the ideal (red) due to noise and an incompletely-explained background slope.

The real data has noise, of course, and also a background slope. The capacitance decreases as we approach the sample. There is also an offset that changes with the scanner's position in x and y , as shown in Fig. A.2. I've always assumed that the background slope comes from stray coupling to the ground plane offered by the sample surface. However, the background slope is the same over grounded, floating, and insulating samples. Rafael Dinner observed his Hall probe sensors moving towards the sample prior to touchdown, perhaps due to electrostatic attraction between the sample and sensor.[81] This would account for a negative slope, but not its long range or linearity. The background slope is linear over perhaps $500\ \mu\text{m}$ of vertical range, which we observe during the initial approach to the sample.




Figure A.2: The ^3He system has two position-dependent backgrounds in the capacitance measurement. Shown are touchdowns at three locations, plotted against z -control voltage (a) and the height above the surface (b). The data show a decrease in capacitance as the scanner approaches the surface, as well as an offset that depends on the position in x and y . The offset can also be observed in constant-height scanning measurements, and is well-described by a plane.

We measure the capacitance in a bridge configuration, with excitation and detection provided by a lock-in amplifier. The specific bridge is an old General Radio bridge with an

internal standard and the ability to use a three-terminal configuration for good rejection of capacitance to ground. This is essential for a low-noise measurement, since the cantilever capacitance is at most a few pF and common lab coaxial cable (RG-58) has a capacitance of 24 pF per foot.

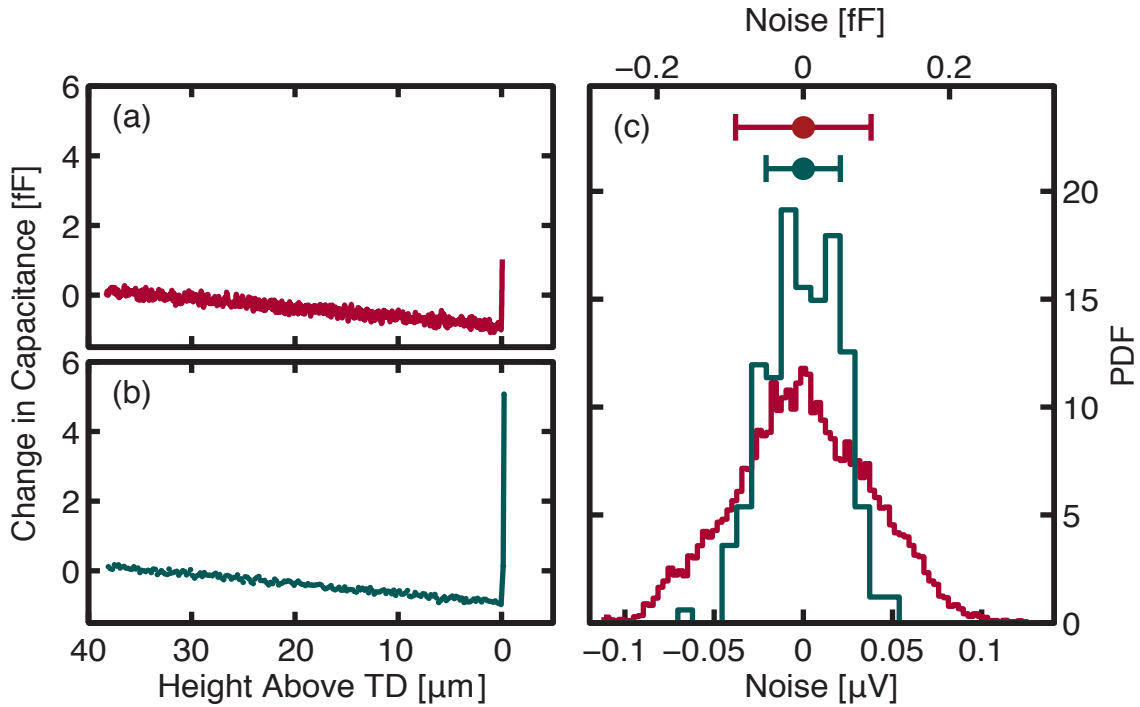


Figure A.3: Data taken in fast mode (a) and fallback mode (b) with otherwise identical parameters, at the same location in x and y . The two data modes are described in Appendix A.3.3. Subtracting a fitted line from the pre-touchdown data, we get the noise histograms in (c). The marker shows the mean and ± 1 standard deviation. Fast std = $0.038 \mu\text{V}$, 0.0926 fF . Fallback std = $0.0208 \mu\text{V}$, 0.0508 fF . Fitting the pre-touchdown data to a line gives a slope of $-0.0265 \text{ fF}/\mu\text{m}$ (fast) and $-0.0274 \text{ fF}/\mu\text{m}$ (fallback). Fitting the post-touchdown data to a line gives a slope of $22.73 \text{ fF}/\mu\text{m}$ (fast) and $19.14 \text{ fF}/\mu\text{m}$ (fallback). There is some curvature in the post-touchdown data likely due to the lockin timeconstant of 10 ms. With a 3 ms timeconstant, the post-touchdown slope goes up to $34.87 \text{ fF}/\mu\text{m}$.

The only intrinsic electronic noise source I've found is thermal (Johnson-Nyquist) noise across the capacitor. The capacitor is in an RC circuit with some resistor; the resistor has the usual Johnson noise, but the RC circuit itself sets the noise bandwidth. So the resistor

value drops out, and the rms voltage noise on the capacitor is just $\sqrt{k_B T/C}$. For a 3 pF capacitor at 4 K, this gives 4.3 μ V.

Another source of noise is mechanical vibration of the cantilever. The excitation voltage of the capacitance measurement can deflect the cantilever. Cliff estimated this deflection to be 8 nm for 1 V DC.[78] Current cantilever spacings are smaller (more on that later), which could increase this value. However, we run the bridge at 20 kHz or so, far above the resonant frequency of the cantilever (a few hundred Hz). So the motion of the cantilever due to our attempts to measure its capacitance is probably small, perhaps a few nm rms. The real vibration comes from the motion of the scanner itself. The scanner is not very stiff, so it has low resonant frequencies and is easily excited by ambient vibrations. The total vertical motion of the sensor is estimated to be 20 nm rms, based on rounding of the susceptometry response over a strong superconductor.[78] We could use the post-touchdown slope to convert this into a capacitance noise. Of course, 20 nm rms is the motion of the sensor, and we're only worried about motion of the cantilever and sensor relative to the base plate. Motion of the sensor and base plate together can only produce capacitance noise through some indirect mechanism involving the background slope.

A.2.1 Sensitivity

I want to distinguish between two distinct figures of merit for the capacitive touchdown detection. Either one could be described as “the sensitivity”, which is why I will belabor the point.

The first figure of merit is the *extent*, the distance the sensor pushes into the sample. This is important for measuring delicate samples, like the Nb dots described in Chapter 5, where the sensor can significantly damage the sample. Of course, the extent is equally important for delicate sensors, such as brittle GaAs Hall probes. Now, if we're concerned with damage, the truly important quantity is the *force* applied to the sample and sensor. The force applied will be proportional to the deflection of the cantilever. The maximum cantilever deflection is the distance the z-bender and base plate travel after the sensor tip first contacts the sample. So our simple definition of “distance the sensor pushes into the sample” is correct, but a bit imprecise. Importantly, we can calculate the cantilever deflection directly from the

z-bender control voltage, by taking the difference between the maximum control voltage and the touchdown position voltage, and converting to length.

The second figure of merit is the *touchdown position uncertainty*, or the error bar on the touchdown position. We could also describe this as the surface uncertainty, since a set of touchdown positions determine the sample surface. The touchdown position is determined by finding the intersection of two lines: the data before the touchdown, and the data after the touchdown. You can see right away that the two figures of merit are in conflict. To decrease the position uncertainty, we should take more data after the touchdown. This is increasing the extent.

So, which of these two things should we prioritize? That depends on your application, on the kind of samples being measured. If the sample and sensor are robust, then it might make sense to focus on the surface uncertainty. One thing to keep in mind is that vibrations, drift, and piezo creep and hysteresis all add to the uncertainty of the sample–sensor separation, but aren’t accounted for by the surface uncertainty determined from fitting touchdown data. It is relatively easy to take slow enough touchdown curves to have a 3 nm surface uncertainty,[78] which is quite impressive on a 20–40 μm z range. But this is in an environment where the sensor tip has an rms vibration amplitude of 20 nm, and creep and drift in the z bender can be tens of nanometers over the course of 12 hours. In some sense, the surface uncertainty is like the error on the mean and it isn’t meaningfully compared to the vibration amplitude, which is more like the standard deviation.

In the next section I’ll talk about efforts to reduce the extent, which was work I did to reduce my odds of damaging the delicate dot array samples in Chapter 5. I don’t have anything to say about the surface uncertainty, specifically how it changed as a side effect of the changes described below. Although I have wondered what bootstrapping the surface–finding fit would give you.

A.3 Reducing the Extent (of the damage)

A.3.1 Minimum Achievable Extent

The best we can hope to do, the minimum achievable extent (MAE), is set by the pre-touchdown noise and the post-touchdown slope. The noise post-touchdown is substantially smaller than pre-touchdown, and in many cases is not observed. It certainly is not a factor in minimizing the extent. Also irrelevant is the pre-touchdown downward mystery slope, since our touchdown algorithms can be designed to account for it. The simplest estimate of the MAE is just the ratio of the noise to the post-touchdown slope. It has the right units, and perhaps that's good enough. Figure A.3 shows a touchdown curve with a noise standard deviation of 0.093 fF and a post-touchdown slope of 23 fF/ μm . The ratio of these is 4 nm, so let's call the simple estimate of the MAE 12 nm.

We can motivate this estimate by thinking of the capacitance post-touchdown as the extent times the post-touchdown slope. Before the touchdown the measurement is random, perhaps normally distributed. Until the measurement exceeds the noise level, the data is indistinguishable from pre-touchdown. Now, "noise level" is a bit vague. If you'll grant that the noise is normally distributed and uncorrelated then we could choose to define the MAE as 3σ divided by the post-touchdown slope, where σ is an estimate of the noise standard deviation. We could quibble over the 3, whether we want a one-sided or two-sided test and what confidence level, etc. That's getting a bit too much into touchdown detection algorithms; see Appendix A.3.4.

The MAE can be improved by reducing the noise or increasing the post-touchdown slope. Reducing the noise requires knowing where it comes from, see Appendix A.2. Increasing the post-touchdown slope is easier than reducing system noise; we can just reduce the spacing between cantilever and base plate. The capacitor is roughly parallel-plate, so the capacitance is proportional to the area divided by the spacing. The post-touchdown slope is the derivative of that, so it is proportional to the area over spacing squared.

If the primary noise is extrinsic electronic noise, then changing the capacitance leaves the noise untouched, and the MAE is proportional to spacing squared. But if the noise is predominantly thermal noise across the capacitor, the noise decreases as we increase the

capacitance, and the MAE will be proportional to spacing to the five-halves power. So decreasing the spacing should have a large effect on the MAE. One complication is in the scenario where the noise is mostly set by cantilever vibration. In that case, the (differential) vibration amplitude is converted into capacitance noise via the post-touchdown slope, and there is no improvement from reducing the spacing.

A.3.2 Cantilever Improvements

I want to emphasize just how important the capacitor spacing is to every aspect of the capacitive touchdown detection. Smaller spacing gives a larger capacitance, which is easier to measure and has lower intrinsic (thermal) noise than a smaller capacitance. A reduced spacing also increases the post-touchdown slope, which brings several benefits. We've seen that it will decrease the extent, and it also allows us to take data faster. A larger post-touchdown slope also improves the topographic resolution when scanning in contact; when scanning the uniform arrays in Chapter 5 I found that the cantilever can detect a 10 nm thick gold film. The slope also makes it feasible to use a simple global threshold detection to guard against crashing into the sample. It could also reduce surface uncertainty, if that is your goal. Really, decreasing the capacitor spacing increases the signal-to-noise ratio of the touchdown measurement, and brings all of the benefits that implies.

I've just stressed the signal importance of the capacitor spacing for improving the touchdown detection. So what have I done to decrease the spacing? Well, it was sort of an accident. Early cantilevers in our lab were made using a piece of glass cover slip as a spacer.[81, 82] This gave a spacing of around 100 μm . I was taught to make cantilevers by Cliff, who used brass shim stock to make the cantilever. A second piece of the same shim stock was used as a spacer, but a spacer under the base plate. We would put silver epoxy on the PCB board and the spacer shim over the base plate. The cantilever is laid over both and a large piece of brass with a machined flat is put over the cantilever while the epoxy cures. Then the spacer shim is pulled out, and the final spacing is around 25 μm . [78]

While working on penetration depth in pnictides, I was having trouble keeping the SQUID behaving with the sample at high temperatures (20 K). Thinking that this was due to the sample heating the SQUID, I set about trying to improve the thermal grounding of

the sensor. (It was actually the increased ^3He boiloff heating the sorb, which is where the array amplifier is connected.) I wanted to make a good (low-resistance) connection from the cantilever base to the copper scanner frame, but that would ground one leg of the capacitor. The measurement would no longer be three-terminal, and how could we hope to measure 3 pF then? So I used cigarette paper and thin GE varnish to glue the cantilever to a large copper pad that was connected to the scanner frame. This gave me a decent thermal connection to the scanner frame, but kept the capacitor isolated from ground. (Wirebonds from a different copper pad to the cantilever provided electrical contact.)

So instead of epoxy, now I have cigarette paper and GE varnish connecting the cantilever to the PCB board. The rest of the procedure was unchanged. Spare shim stock over the base plate, big heavy brass piece, remove spacer shim. The capacitive touchdown detection still worked well, and perhaps the thermal grounding of the sensor was better. At some point after this, I noticed a bend in the cantilever. Apparently cigarette paper is thinner than the brass shim stock. Naturally, the next time I needed to make a new cantilever, I used cigarette paper as the spacer. Now there is no noticeable bend in the cantilever, and the post-touchdown slope is substantially larger. This happy accident is what made all of the improvements described in this appendix possible.

We can put some rough numbers on this historical improvement by looking at figures in theses. In Per's thesis, Fig. 2.6 shows a capacitive touchdown with a post-touchdown slope of $0.5 \text{ fF}/\mu\text{m}$ (that's femto-Farads) and an extent of $2.4 \mu\text{m}$.^[82] Rafael's thesis, Fig. 4.4(c), gives $1 \text{ fF}/\mu\text{m}$ with a $10 \mu\text{m}$ extent.^[81] Both of these figures clearly show that smaller extents were possible, but it wasn't a priority. Moving from glass cover slip to brass shim, $100 \mu\text{m}$ down to $25 \mu\text{m}$ spacing, Cliff estimated a post-touchdown slope of $50 \text{ fF}/\mu\text{m}$ from the geometry of the cantilever, then says that "In practice about half this is typical."^[78] Sure enough, from Fig. 1.5(b) I estimate a slope of $20 \text{ fF}/\mu\text{m}$ with an extent around $0.4 \mu\text{m}$.^[78] Using cigarette paper to set the spacing, I can achieve a slope of $30 \text{ fF}/\mu\text{m}$. The extent is typically $0.1 \mu\text{m}$. That's when taking fast, routine touchdowns; if I'm trying to not damage a delicate sample, the extent can be as small as $0.03 \mu\text{m}$.

A.3.3 Data-Taking Procedure

Like the cantilever spacing, I've also changed the method of taking touchdown data, with two variations. The traditional method was to position the scanner at a specified point, pause, measure the capacitance, then move to the next point, repeat. The user specified the spacing between the points, and the pause time was determined from the timeconstant of the lock-in. With this method the touchdown speed is specified by the point spacing, which will also affect the extent and the surface uncertainty. The time spent moving between points was always substantially larger than the time spent on the capacitance measurement.[78] The new method of taking touchdown data is to ramp the scanner continuously while retrieving data periodically to determine if a touchdown has occurred. The motivation for the change to continuous motion was to have SQUID susceptometry touchdowns simultaneously acquired with capacitive touchdown detection. Then the capacitance signal provides a consistency check on the SQUID data, which I felt was desirable when using susceptometry touchdowns to measure the penetration depth of superconductors.

In addition to allowing simultaneous data from multiple channels during a touchdown, the new method is also significantly faster. The touchdown speed can be controlled with a new parameter that sets the approach velocity. The sample rate of the analog output is chosen to give the desired velocity when subsequent samples are separated by one least significant bit of the 16-bit DAQ card. For typical touchdown speeds, between 0.1 and 0.5 (control) volts per second, the analog output rate ends up between 300 and 1600 Hz. I have chosen to match the rates of the analog input and output, reasoning that each data sample should have an exact corresponding position. It is possible to sample the analog input substantially faster, and use the average of many samples as the data at each position. This is how the scanning data is taken, and it seems to work well, although the input sample rate is almost never an integer multiple of the output rate.

The two variations on the new method differ in how often data is retrieved from the analog input. The first variation, a "fallback mode", uses the point spacing parameter to determine the number of samples that are averaged together into one data point. For typical settings this number is around 150. From the outside, fallback mode looks like the old method of taking touchdown data, with the additions of a velocity parameter and the ability to record multiple data channels simultaneously. The second variation, let's call it "fast

mode”, was created specifically for minimum extent. In this mode we retrieve data from the analog input as fast as it comes in. Since that is never faster than 2 kHz, this should be feasible. I’ll go into this a bit more in Appendix A.3.5.




Figure A.4: A touchdown curve in fast (left) and fallback (right) modes, showing the data (color) and the running threshold (black) which triggers the end of the approach. The tan square on the left indicates the point where the touchdown was triggered; the points after are due to latencies in the control loop.

A.3.4 Detection Algorithms

To reduce the touchdown extent, we need to examine and optimize the detection algorithms; meaning the calculations that determine whether or not to stop the approach. My goal was to reduce the touchdown extent, but that comes at the price of a higher false-positive rate. That’s the real tradeoff to be made in choosing detection algorithms; how many false alarms will you tolerate? The other tradeoff that comes to mind, between algorithm speed and sophistication, is really the false-positive tradeoff. Reducing the extent required streamlining the detection algorithms so that we could sample the data faster and

make stop/continue decisions faster. But a slower algorithm is (presumably) more sophisticated, and you would hope it can give you a lower false-positive rate without increasing the extent.

The detection algorithms do not determine the location of the surface. That is done after the fact by fitting the beginning and end of the data separately to two lines and finding the point of intersection. It might be worthwhile to use this post-processing step to perform a confidence test on whether the data describe a real touchdown or a false-positive. If the data are determined to be a potential false-positive we could re-take the touchdown. This could reduce the false-positive rate without increasing the touchdown extent unnecessarily.

Threshold

The simplest type of detection is threshold detection. If the signal exceeds some critical value, stop the approach. Then the choice of critical value will set the false-positive rate and minimum extent, based on the noise and so on. But our capacitance data is complicated by the offset and background slope illustrated in Fig. A.2. And so historically, we never used simple threshold detection. But the new cantilever construction has such a large post-touchdown slope that it was worth considering. In the end, the pre-touchdown slope and offset make it insufficient, but worth keeping as a final safety check.

The next step is to design a test that can account for the pre-touchdown slope. The algorithm I inherited was to fit the most of the collected data to a line. The fit gives an estimate of the noise as well as a prediction for the newest data points. Then we can do a simple threshold test against the predicted value plus the noise. This is a nice solution. It is elegant and can be self-calibrating because the noise and slope are estimated from the data. It is also simpler than it could be; it doesn't calculate the fit prediction interval with a t-test and so on. The primary downside is that it is slow, requiring a call to `polyfit` and `polyval` each iteration. The specific implementation in `scantdC` had some trouble with very small point spacing, just because some things were hard-coded in terms of data points (rather than distances).

We can also account for the pre-touchdown slope by calculating a weighted running average of the data. This is nice because it is very simple to calculate. If we can add an independent estimate of the noise, we can even keep the self-calibrating feature. This means

that the detection algorithms don't have to be tuned for every conceivable set of touchdown parameters, whether that's lockin gain and timeconstant or touchdown speed.

I implemented this by calculating a running cutoff as the exponentially-weighted moving average of the data. This is easy to calculate as

$$C_i = \alpha X_i + (1 - \alpha)C_{i-1}, \quad (\text{A.1})$$

where C is the running cutoff and X_i is the most recent data sample. Other details are that I use a different α for fast and fallback mode, and in fast mode we update the cutoff every 16 samples with the average of the last 16 samples. As an equation, we update the cutoff as:

$$C_i = \alpha \left(\frac{1}{16} \sum_{j=i-15}^i X_j \right) + (1 - \alpha)C_{i-1}, \quad (\text{A.2})$$

when i is a multiple of 16. To calibrate the noise and initialize the running cutoff, I sample the capacitance for half a second at the starting position of the touchdown. Then the detection is calculated as a threshold of the most recent data point against the running cutoff plus four times the noise estimate, which is the standard deviation of the starting sample. This is illustrated in Fig. A.4, which shows two example touchdown curves and their changing detection threshold.

Slope

Another possible test is to compare the slope of the data to a critical value. Data after touchdown has a large positive slope, while pre-touchdown data has a small negative slope, so a good choice of critical value is half of the final post-touchdown slope. We could choose to fit the most recent N data samples to a line and compare the fitted slope to the chosen threshold. But if you stare at the equation for the least-squares estimated slope of a line, you can show that absolute x offset drops out provided that the data spacing in x is uniform. This is the case here, so the slope can be estimated by calculating the dot product of the data with a vector of pre-computed constants.

The biggest drawback of the slope test is that it amplifies noise because it is a derivative estimator. This makes it sensitive to parameters like the noise level, lockin time constant,

and data point spacing. I found that I could get it to work nicely for a given parameter configuration with some tuning, but it would fail hopelessly for other configurations. The other drawback is that when the slope test was working well, it was redundant with the running cutoff threshold test. In the end, the slope test was at best useless, so I got rid of it.

A.3.5 Goddamnit, MATLAB

First, it's worth saying something about the guts of the data acquisition toolbox. All of the data acquisition (DAQ) code in the lab uses what is now called the "Legacy Interface". It isn't currently scheduled for deprecation (disappearance), although you can't use it on 64-bit MATLAB installations. In the legacy interface, the user interacts with the DAQ board through data acquisition objects. There are different classes for analog input, analog output, and digital i/o. The DAQ boards we use from National Instruments have all three, as well as digital clocks and counters. So we will have different objects in the MATLAB workspace that refer to different features of the same physical DAQ device.

The practical effect of this is that it is impossible to get the analog input and analog output to run at the same frequency. You can get close; for most of the touchdowns I've taken the two differ by 4 mHz out of about 1 kHz. It is also impossible to get them to start at the same time. Yes, yes, it is actually impossible to get two things truly synchronized. But the best I've been able to do is a 30 millisecond delay between the two devices starting. It doesn't matter whether I use external hardware triggering, where a digital output line sends a falling edge to the DAQ board's start triggers, or if I use the RTSI bus (Real-Time System Integration), which is built into the DAQ board and provided specifically for the purpose of synchronizing operations. Although it does look like the time difference is always the same using the RTSI bus, and somewhat random with external hardware triggering. Thirty milliseconds seem unnecessarily large when the computer processor is clocked above 1 GHz and the DAQ board's internal clock runs at 20 MHz. It is similarly impossible to get the analog input to sample at an integer multiple of the analog output rate.

The other thing to mention about the legacy interface is that it cannot interact with the clocks and counters. This is important for "real-time measurements", where an excitation such as the field coil current is slaved to the DAQ board to ensure a constant phase lag

between the excitation and measured response. In those situations we use C code to interact directly with the software drivers provided by National Instruments.

Well, if something's called the "legacy interface", there must be something newer, right? There is, and it is called the "session-based interface". The paradigm is completely different. Under the session-based interface, we interact with an object that represents the entire data acquisition *session*. (Hence the name.) All features of the session — analog input and output, plus counters, clocks, and digital i/o — belong to this object. So now it's trivial to keep the analog input and output running at the same frequency. In fact, we have to, there is no way to sample the data faster than the output changes. This is how the scanning data is taken, so the session-based interface is unsuited to our needs. The new interface does make it easier to synchronize input and output; the session object can have trigger properties that specify connections between various digital signals and the start triggers of different objects. In addition, the session object can mix and match signals between different physical data acquisition devices. It's really quite nice.

But there's always a catch, right? Sure enough, there is no simple way to *retrieve the data*. The new method uses callbacks and listeners to send data around. So now I need to create a data session, and give it callback functions to execute once it gets data, or runs out of data to output. It's completely bonkers. Especially considering the substantial overhead to calling functions in MATLAB. My opinion is that it is completely unsuited for real-time control. I do not see a path to migrating the Moler Lab DAQ to the session-based interface, even if someone was willing to do the complete re-write necessary.

Should the legacy interface be deprecated, there are a few options. Obviously, the session-based interface should be re-evaluated at that time, to see if the situation has improved. The next option is to keep the DAQ largely in MATLAB, but without the data acquisition tools. This would mean implementing the core routines using the National Instruments drivers (NI-DAQmx) in C, then use the MEX library to run it and pass the data back out to MATLAB. The final option is to abandon MATLAB entirely. There are few forms this can take. We could switch the DAQ to different software but continue to use MATLAB for analysis, or both DAQ and analysis can occur in the new environment. Naturally, LabVIEW has excellent support for National Instruments DAQ boards, if you're willing to program in Crayon. I don't know what analysis in LabVIEW is like, or if it even exists. Another option is Python,

which would actually require some things to be written in C. These core utilities can then be wrapped in Python, which can be used to provide a nice object-oriented interface. Analysis and plotting can be done in Python using a few different libraries.

Switching gears back to the present, I've learned a few things about how the legacy interface is implemented that may be useful. The generic advice is that the `struct` command can reveal the internals of any MATLAB object. For example, if `ai` is an analog input object, `struct(ai)` returns a struct with fields `uddobject`, `daqdevice`, `version`, and `info`. The first two fields refer to other classes, so we can call `struct` on them again to find their properties. The `daqdevice` object is boring - it only contains an empty cell array called `store`. The `daqdevice` class exists to provide common DAQ routines like `start`, `stop`, and `wait`. The `uddobject` contains every property of the analog input object. In fact, the `uddobject` is a reference to the physical hardware device. This allows us to have multiple objects in memory that refer to the same device. The *usefulness* of this knowledge is that we can speed up many interactions with the DAQ device by directly using the `uddobject`, since that's what happens under the hood anyways. For example, `start(uddobject)` executes faster than `start(ai)`.

Undocumented Buffers

There are two ways to read analog input objects using the legacy interface: `getdata` and `peekdata`. I lied, there's actually a third way, `getsample`, that immediately returns one sample. But it doesn't work with hardware triggers, so we never use it. We use `getdata` most of the time. It is a *blocking* function, which means that it doesn't return until the specified number of samples are available. It also removes the retrieved samples from the analog input object. So if we call `getdata` repeatedly during the approach, we will have to collect the data as we go. Usually we use `getdata` to collect data once acquisition has finished, such as at the end of a scan line. In contrast, `peekdata` is non-blocking; if you ask it for more samples than are available, it will give you what it has. Since it just peeks at the data, it also doesn't remove it from the DAQ engine. So we can call it repeatedly during the approach, and all the data will be waiting for us to retrieve at the end.

I hope you agree that `peekdata` seems like the obvious choice for rapidly previewing

the incoming data and making a touchdown decision. And according to the MATLAB documentation, “`data = peekdata(obj, samples)` returns the latest number of samples specified by `samples` to `data`.” So that’s what I used, and early versions of the fast mode used `peekdata` in a `while` loop. The loop condition was to check if the analog output is still running, and profiling the loop is what taught me about `uddobjects`. I also incremented a counter in the loop to have some idea how rapidly the code was peeking at the incoming data. After a lot of streamlining and optimization, I could get to about three times oversampled at the largest typical approach velocity; 0.5 control volts per second, about 2 $\mu\text{m/s}$. This works out to a loop frequency of just under 5 kHz. Since I’m trying to sample as fast as possible, there will be repeated data in the samples. This isn’t a problem; it only affects one of the touchdown tests, and there’s a simple workaround. (Use timing functions to control the update rate of the running cutoff)

So long as `peekdata` returns the most recently acquired data, this should work. But while testing, I kept getting data like that shown in Fig. A.5(a). We can see that there are several hundred samples taken after the data has clearly crossed the threshold defined by the running cutoff (black line). It should have stopped much sooner. At some point I started recording the most recent sample from `peekdata`, those samples are highlighted with blue dots. What gives? The test to do now was to record every response from `peekdata` and compare it to the complete dataset collected after the touchdown using `getdata`. What I expect, from the documentation, is for subsequent responses to either be identical or differ by two samples (the beginning and end). The mental picture is taking pictures of a window that always shows the most recent sixteen (in this case) samples. We’ll get the same snapshot until a new sample is available, when it will slide into the picture and the oldest one will slide out. So what is MATLAB really doing? A small section of that test is in Fig. A.5(b), which shows the data pulled by `peekdata` as blue dots clearly missing most of the red trace of the complete dataset. There must be some other data buffer that does not update fast enough. This won’t work; so I switched to `getdata(1)` and accumulating the collected data. This worked excellently for rapidly stopping the approach and minimizing the extent, but probably caused the problem described in the next section.




Figure A.5: The perils of believing the MATLAB documentation. Specifically, the DAQ function `peekdata` does not work as advertised. (a) A touchdown curve (red) with the running cutoff (black) and the sample from `peekdata` that triggered the touchdown tests (blue dots). The touchdown clearly should have stopped before it did. (b) The results of a test accumulating every sample from `peekdata` (blue dots) during a touchdown, then pulling the entire dataset (red) with `getdata`. Profiling indicated that `peekdata` was being called three times faster than the analog output rate, so you'd expect the blue dots to cover every data point three times. Instead the gaps between blue dots indicate that something undocumented in the internals is not updating fast enough.

Mysterious Timeout Error

Unfortunately, a rather mysterious error cropped up once the fast mode was working nicely. After around five hours taking repeated fast touchdowns, MATLAB would freeze up and require a system restart. I fought with it repeatedly, but it proved quite impossible to debug. I suspect it is due to a memory leak in a low-level DAQ routine, but that's all I can say.

Bibliography

1. W. B. Pearson, I. M. Templeton,
Superconducting Transition of Lead.
Physical Review **109**, 1094–1094 (1958). [Link](#)
2. Michael Tinkham,
Introduction to Superconductivity
(Dover Publications, Mineola, N.Y., 2004).
3. J. Bardeen, L. N. Cooper, J. R. Schrieffer,
Theory of Superconductivity.
Physical Review **108**, 1175 (1957). [Link](#)
4. J. File, R. G. Mills,
Observation of Persistent Current in a Superconducting Solenoid.
Physical Review Letters **10**, 93–96 (1963). [Link](#)
5. G.F. F. Sun, K.W. W. Wong, B.R. R. Xu, Y. Xin, D.F. F. Lu,
 T_c enhancement of $\text{HgBa}_2\text{Ca}_2\text{Cu}_3\text{O}_{8+\delta}$ by Tl substitution.
Physics Letters A **192**, 122–124 (1994). [Link](#)
6. W. N. Hardy, D. A. Bonn, D. C. Morgan, Ruixing Liang, Kuan Zhang,
Precision measurements of the temperature dependence of λ in $\text{YBa}_2\text{Cu}_3\text{O}_{6.95}$:
Strong evidence for nodes in the gap function.
Physical Review Letters **70**, 3999 (1993). [Link](#)

7. W. N. Hardy, S. Kamal, D. A. Bonn,
“Magnetic Penetration Depths in Cuprates: A short Review of Measurement Techniques and Results”,
The Gap Symmetry and Fluctuations in High-Tc Superconductors, vol. 371,
pp. 373–402 (Kluwer Academic Publishers, Boston, 2002),
[Link](#).
8. A. T. Fiory, A. F. Hebard,
Radio-frequency complex-impedance measurements on thin film two-dimensional superconductors.
AIP Conference Proceedings **58**, 293–298 (1980). [Link](#)
9. Stefan J. Turneaure, Aaron A. Pesetski, Thomas R. Lemberger,
Numerical modeling and experimental considerations for a two-coil apparatus to measure the complex conductivity of superconducting films.
Journal of Applied Physics **83**, 4334 (1998). [Link](#)
10. Ruslan Prozorov, Russell W. Giannetta,
Magnetic penetration depth in unconventional superconductors.
Superconductor Science and Technology **19**, R41–R67 (2006). [Link](#)
11. T. Pereg-Barnea, P. J. Turner, R. Harris, G. K. Mullins, J. S. Bobowski, M. Raudsepp, Ruixing Liang, D. A. Bonn, W. N. Hardy,
Absolute values of the London penetration depth in $\text{YBa}_2\text{Cu}_3\text{O}_{6+y}$ measured by zero field ESR spectroscopy on Gd doped single crystals.
Physical Review B **69**, 184513 (2004). [Link](#)
12. R. Prozorov, R. W. Giannetta, A. Carrington, F. M. M. Araujo-Moreira,
Meissner–London state in superconductors of rectangular cross section in a perpendicular magnetic field.
Physical Review B **62**, 115–118 (2000). [Link](#)

13. R. Prozorov, R. W. Giannetta, A. Carrington, P. Fournier, R. L. Greene, P. Guptasarma, D. G. Hinks, A. R. Banks, Measurements of the absolute value of the penetration depth in high- T_c superconductors using a low- T_c superconductive coating. *Applied Physics Letters* **77**, 4202–4204 (2000). [Link](#)
14. R. T. Gordon, H. Kim, N. Salovich, R. W. Giannetta, R. M. Fernandes, V. G. Kogan, T. Prozorov, S. L. Bud'ko, P. C. Canfield, M. A. Tanatar, R. Prozorov, Doping evolution of the absolute value of the London penetration depth and superfluid density in single crystals of $\text{Ba}(\text{Fe}_{1-x}\text{Co}_x)_2\text{As}_2$. *Physical Review B* **82**, 054507 (2010). [Link](#)
15. Hendrik Bluhm, Nicholas C. Koshnick, Martin E. Huber, Kathryn A. Moler, Magnetic Response of Mesoscopic Superconducting Rings with Two Order Parameters. *Physical Review Letters* **97**, 237002 (2006). [Link](#)
16. Jeff E. Sonier, Jess H. Brewer, Robert F. Kiefl, μSR studies of the vortex state in type-II superconductors. *Reviews of Modern Physics* **72**, 769–811 (2000). [Link](#)
17. S. F. J. Cox, Implanted muon studies in condensed matter science. *Journal of Physics C* **20**, 3187 (1987). [Link](#)
18. J. E. Sonier, R. F. Kiefl, J. H. Brewer, D. A. Bonn, J. F. Carolan, K. H. Chow, P. Dosanjh, W. N. Hardy, Ruixing Liang, W. A. MacFarlane, P. Mendels, G. D. Morris, T. M. Riseman, J. W. Schneider, New muon-spin-rotation measurement of the temperature dependence of the magnetic penetration depth in $\text{YBa}_2\text{Cu}_3\text{O}_{6.95}$. *Physical Review Letters* **72**, 744 (1994). [Link](#)

19. D. R. Harshman, G. Aeppli, E. J. Ansaldo, B. Batlogg, J. H. Brewer, J. F. Carolan, R. J. Cava, M. Celio, A. C. D. Chaklader, W. N. Hardy, S. R. Kreitzman, G. M. Luke, D. R. Noakes, M. Senba,
Temperature dependence of the magnetic penetration depth in the high- T_c superconductor $\text{Ba}_2\text{YCu}_3\text{O}_{9-\delta}$: Evidence for conventional s-wave pairing.
Physical Review B **36**, 2386–2389 (1987). [Link](#)
20. Y. J. Uemura, V. J. Emery, A. R. Moodenbaugh, M. Suenaga, D. C. Johnston, A. J. Jacobson, J. T. Lewandowski, J. H. Brewer, R. F. Kiefl, S. R. Kreitzman, G. M. Luke, T. Riseman, C. E. Stronach, W. J. Kossler, J. R. Kempton, X. H. Yu, D. Opie, H. E. Schone,
Systematic variation of magnetic-field penetration depth in high- T_c superconductors studied by muon-spin relaxation.
Physical Review B **38**, 909–912 (1988). [Link](#)
21. D. R. Harshman, L. F. Schneemeyer, J. V. Waszczak, G. Aeppli, R. J. Cava, B. Batlogg, L. W. Rupp, E. J. Ansaldo, D. L. Williams,
Magnetic penetration depth in single-crystal $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$.
Physical Review B **39**, 851 (1989). [Link](#)
22. B. Pümpin, H. Keller, W. Kündig, W. Odermatt, I. M. Savić, J. W. Schneider, H. Simmler, P. Zimmermann, E. Kaldis, S. Rusiecki, Y. Maeno, C. Rossel,
Muon-spin-rotation measurements of the London penetration depths in $\text{YBa}_2\text{Cu}_3\text{O}_{6.97}$.
Physical Review B **42**, 8019 (1990). [Link](#)
23. E. H. Brandt,
Flux distribution and penetration depth measured by muon spin rotation in high- T_c superconductors.
Physical Review B **37**, 2349–2352 (1988). [Link](#)

24. H. Luetkens, H.-H. Klauss, R. Khasanov, A. Amato, R. Klingeler, I. Hellmann, N. Leps, A. Kondrat, C. Hess, A. Köhler, G. Behr, J. Werner, B. Büchner, Field and Temperature Dependence of the Superfluid Density in $\text{LaFeAsO}_{1-x}\text{F}_x$ Superconductors: A Muon Spin Relaxation Study. *Physical Review Letters* **101**, 097009 (2008). [Link](#)
25. Rustem Khasanov, Hubertus Luetkens, Alex Amato, Hans-Henning - H. Klauss, Zhi-An Ren, Jie Yang, Wei Lu, Zhong-Xian - X. Zhao, Muon spin rotation studies of $\text{SmFeAsO}_{0.85}$ and $\text{NdFeAsO}_{0.85}$ superconductors. *Physical Review B* **78**, 092506 (2008). [Link](#)
26. R. Khasanov, Takeshi Kondo, S. Strässle, D. O. G. Heron, A. Kaminski, H. Keller, S. L. Lee, Tsunehiro Takeuchi, Zero-field superfluid density in a d-wave superconductor evaluated from muon-spin-rotation experiments in the vortex state. *Physical Review B* **79**, 180507 (2009). [Link](#)
27. D. E. MacLaughlin, J. E. Sonier, R. H. Heffner, O. O. Bernal, Ben-Li Young, M. S. Rose, G. D. Morris, E. D. Bauer, T. D. Do, M. B. Maple, Muon Spin Relaxation and Isotropic Pairing in Superconducting $\text{PrOs}_4\text{Sb}_{12}$. *Physical Review Letters* **89**, 157001 (2002). [Link](#)
28. J. E. Sonier, W. Huang, C. V. Kaiser, C. Cochran, V. Pacradouni, S. A. Sabok-Sayr, M. D. Lumsden, B. C. Sales, M. A. McGuire, A. S. Sefat, D. Mandrus, Magnetism and Disorder Effects on Muon Spin Rotation Measurements of the Magnetic Penetration Depth in Iron-Arsenic Superconductors. *Physical Review Letters* **106**, 127002 (2011). [Link](#)
29. Ch. Niedermayer, E. M. Forgan, H. Glückler, A. Hofer, E. Morenzoni, M. Pleines, T. Prokscha, T. M. Riseman, M. Birke, T. J. Jackson, J. Litterst, M. W. Long, H. Luetkens, A. Schatz, G. Schatz, Direct Observation of a Flux Line Lattice Field Distribution across an $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ surface by Low Energy Muons. *Physical Review Letters* **83**, 3932–3935 (1999). [Link](#)

30. R. F. Kiefl, M. D. Hossain, B. M. Wojek, S. R. Dunsiger, G. D. Morris, T. Prokscha, Z. Salman, J. Baglo, D. A. Bonn, R. Liang, W. N. Hardy, A. Suter, E. Morenzoni, Direct measurement of the London penetration depth in $\text{YBa}_2\text{Cu}_3\text{O}_{6.92}$ using low-energy μSR .
Physical Review B **81**, 180502 (2010). [Link](#)
31. Oren Ofer, J. C. Baglo, M. D. Hossain, R. F. Kiefl, W. N. Hardy, A. Thaler, H. Kim, M. A. Tanatar, P. C. Canfield, R. Prozorov, G. M. Luke, E. Morenzoni, H. Saadaoui, A. Suter, T. Prokscha, B. M. Wojek, Z. Salman,
Absolute value and temperature dependence of the magnetic penetration depth in $\text{Ba}(\text{Co}_{0.074}\text{Fe}_{0.926})_2\text{As}_2$.
Physical Review B **85**, 060506 (2012). [Link](#)
32. T. J. Jackson, T. M. Riseman, E. M. Forgan, H. Glückler, T. Prokscha, E. Morenzoni, M. Pleines, Ch Niedermayer, G. Schatz, H. Luetkens, J. Litterst,
Depth-Resolved Profile of the Magnetic Field beneath the Surface of a Superconductor with a Few nm Resolution.
Physical Review Letters **84**, 4958–4961 (2000). [Link](#)
33. A. Suter, E. Morenzoni, R. Khasanov, H. Luetkens, T. Prokscha, N. Garifianov,
Direct Observation of Nonlocal Effects in a Superconductor.
Physical Review Letters **92**, 087001 (2004). [Link](#)
34. A. Suter, E. Morenzoni, N. Garifianov, R. Khasanov, E. Kirk, H. Luetkens, T. Prokscha, M. Horisberger,
Observation of nonexponential magnetic penetration profiles in the Meissner state: A manifestation of nonlocal effects in superconductors.
Physical Review B **72**, 024506 (2005). [Link](#)
35. C. W. Hicks, L. Luan, K. A. Moler, E. Zeldov, H. Shtrikman,
Noise characteristics of 100 nm scale $\text{GaAs}/\text{Al}_x\text{Ga}_{1-x}\text{As}$ scanning Hall probes.
Applied Physics Letters **90**, 133512 (2007). [Link](#)
36. Judea Pearl,
Structure of Superconductive Vortices near a Metal-Air Interface.
Journal of Applied Physics **37**, 4139 (1966). [Link](#)

37. V. G. Kogan, A. Yu. Simonov, M. Ledvij,
Magnetic field of vortices crossing a superconductor surface.
Physical Review B **48**, 392 (1993). [Link](#)
38. J. C. Wynn, D. A. Bonn, B. W. Gardner, Yu-Ju Lin, Ruixing Liang, W. N. Hardy,
J. R. Kirtley, K. A. Moler,
Limits on Spin-Charge Separation from $h/2e$ Fluxoids in Very Underdoped
 $\text{YBa}_2\text{Cu}_3\text{O}_{6+x}$.
Physical Review Letters **87**, 197002 (2001). [Link](#)
39. Hendrik Bluhm, Nicholas C. Koshnick, Julie A. Bert, Martin E. Huber,
Kathryn A. Moler,
Persistent Currents in Normal Metal Rings.
Physical Review Letters **102**, 136802 (2009). [Link](#)
40. John Clarke, Alex I. Braginski, Eds.,
The SQUID Handbook
(Wiley-VCH, Weinheim, 2004).
41. Martin E. Huber, Nicholas C. Koshnick, Hendrik Bluhm, Leonard J. Archuleta,
Tommy Azua, Per G. Björnsson, Brian W. Gardner, Sean T. Halloran, Erik A. Lucero,
Kathryn A. Moler,
Gradiometric micro-SQUID susceptometer for scanning measurements of
mesoscopic samples.
Rev Sci Instrum **79**, 053704 (2008). [Link](#)
42. J. Pearl,
Current distribution in superconducting films carrying quantized fluxoids.
Applied Physics Letters **5**, 65 (1964). [Link](#)
43. V. G. Kogan,
Meissner response of anisotropic superconductors.
Physical Review B **68**, 104511 (2003). [Link](#)

44. John R. Clem, Mark W. Coffey,
Vortex dynamics in a type-II superconducting film and complex linear-response functions.
Physical Review B **46**, 14662 (1992). [Link](#)
45. Stefan J. Turneaure, Eric R. Uml, Thomas R. Lemberger,
Numerical modeling of a two-coil apparatus for measuring the magnetic penetration depth in superconducting films and arrays.
Journal of Applied Physics **79**, 4221 (1996). [Link](#)
46. Mark W. Coffey,
Analyzing mutual inductance measurements to determine the London penetration depth.
Journal of Applied Physics **87**, 4344–4351 (2000). [Link](#)
47. John David Jackson,
Classical Electrodynamics
(Wiley, New York, 1999).
48. Shigehiko Okui,
Complete elliptic integrals resulting from infinite integrals of Bessel functions.
Journal of Research of the National Bureau of Standards **78B**, 113 (1974). [Link](#)
49. James Annett, Nigel Goldenfeld, S. R. Renn,
Interpretation of the temperature dependence of the electromagnetic penetration depth in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$.
Physical Review B **43**, 2778 (1991). [Link](#)
50. V. Mishra, G. Boyd, S. Graser, T. Maier, P. J. Hirschfeld, D. J. Scalapino,
Lifting of nodes by disorder in extended-s-state superconductors: Application to ferropnictides.
Physical Review B **79**, 094512–9 (2009). [Link](#)

51. Peter J. Hirschfeld, Nigel Goldenfeld,
Effect of strong scattering on the low-temperature penetration depth of a d-wave superconductor.
Physical Review B **48**, 4219 (1993). [Link](#)
52. A. B. Vorontsov, M. G. Vavilov, A. V. Chubukov,
Superfluid density and penetration depth in the iron pnictides.
Physical Review B **79**, 140507–4 (2009). [Link](#)
53. P. G. De Gennes,
Superconductivity Of Metals And Alloys
(Westview Press, Boulder, Co, 1999).
54. Øystein Fischer, Martin Kugler, Ivan Maggio-Aprile, Christophe Berthod,
Christoph Renner,
Scanning tunneling spectroscopy of high-temperature superconductors.
Reviews of Modern Physics **79**, 353 (2007). [Link](#)
55. Yi Yin, M. Zech, T. L. Williams, X. F. Wang, G. Wu, X. H. Chen, J. E. Hoffman,
Scanning Tunneling Spectroscopy and Vortex Imaging in the Iron Pnictide
Superconductor $\text{BaFe}_{1.8}\text{Co}_{0.2}\text{As}_2$.
Physical Review Letters **102**, 097002 (2009). [Link](#)
56. M. L. Teague, G. K. Drayna, G. P. Lockhart, P. Cheng, B. Shen, H.-H. Wen, N.-C. Yeh,
Measurement of a Sign-Changing Two-Gap Superconducting Phase in
Electron-Doped $\text{Ba}(\text{Fe}_{1-x}\text{Co}_x)_2\text{As}_2$ Single Crystals Using Scanning Tunneling
Spectroscopy.
Physical Review Letters **106**, 087004 (2011). [Link](#)
57. Y. Fasano, I. Maggio-Aprile, N. D. Zhigadlo, S. Katrych, J. Karpinski, Ø. Fischer,
Local Quasiparticle Density of States of Superconducting $\text{SmFeAsO}_{1-x}\text{F}_x$ Single
Crystals: Evidence for Spin-Mediated Pairing.
Physical Review Letters **105**, 167005 (2010). [Link](#)

58. Benjamin Sacépé, Thomas Dubouchet, Claude Chapelier, Marc Sanquer, Maoz Ovadia, Dan Shahar, Mikhail Feigel'man, Lev Ioffe, Localization of preformed Cooper pairs in disordered superconductors. *Nature Physics* **7**, 239–244 (2011). [Link](#)
59. B. Sacépé, C. Chapelier, T. I. Baturina, V. M. Vinokur, M. R. Baklanov, M. Sanquer, Disorder-Induced Inhomogeneities of the Superconducting State Close to the Superconductor-Insulator Transition. *Physical Review Letters* **101**, 157006 (2008). [Link](#)
60. Lan Luan, Thomas M. Lippman, Clifford W. Hicks, Julie A. Bert, Ophir M. Auslaender, Jiun-Haw Chu, James G. Analytis, Ian R. Fisher, Kathryn A. Moler, Local Measurement of the Superfluid Density in the Pnictide Superconductor $\text{Ba}(\text{Fe}_{1-x}\text{Co}_x)_2\text{As}_2$ across the Superconducting Dome. *Physical Review Letters* **106**, 067001 (2011). [Link](#)
61. Clifford W. Hicks, Thomas M. Lippman, Martin E. Huber, James G. Analytis, Jiun-Haw Chu, Ann S. Erickson, Ian R. Fisher, Kathryn A. Moler, Evidence for a Nodal Energy Gap in the Iron-Pnictide Superconductor LaFePO from Penetration Depth Measurements by Scanning SQUID Susceptometry. *Physical Review Letters* **103**, 127003 (2009). [Link](#)
62. J. R. Cave, J. E. Evetts, Critical temperature profile determination using a modified London equation for inhomogeneous superconductors. *Journal of Low Temperature Physics* **63**, 35–55 (1986). [Link](#)
63. Lawrence A. Mysak, Wave propagation in random media, with oceanic applications. *Reviews of Geophysics* **16**, PP. 233–261 (1978). [Link](#)
64. N. G. Van Kampen, Stochastic differential equations. *Physics Reports* **24**, 171–228 (1976). [Link](#)

65. Frank C. Karal, Joseph B. Keller,
Elastic, Electromagnetic, and Other Waves in a Random Medium.
Journal of Mathematical Physics **5**, 537 (1964). [Link](#)
66. Bertil Matérn,
Spatial Variation
(Springer-Verlag, Berlin, Second, 1986),
Vol. 36, p. 151.
67. Mark S. Handcock, James R. Wallis,
An Approach to Statistical Spatial-Temporal Modeling of Meteorological Fields.
Journal of the American Statistical Association **89**, 368–378 (1994). [Link](#)
68. I. S. Gradshteyn, I. M. Ryzhik,
Table of Integrals, Series, and Products
(Academic Press, Corrected and Enlarged Edition, 1980).
69. Noel A. C. Cressie,
Statistics for Spatial Data
(Wiley, New York, 1991),
P. 928.
70. Amit Ghosal, Mohit Randeria, Nandini Trivedi,
Inhomogeneous pairing in highly disordered s-wave superconductors.
Physical Review B **65**, 014501 (2001). [Link](#)
71. M. Franz, C. Kallin, A. J. Berlinsky, M. I. Salkola,
Critical temperature and superfluid density suppression in disordered high- T_c
cuprate superconductors.
Physical Review B **56**, 7882 (1997). [Link](#)
72. Amit Ghosal, Mohit Randeria, Nandini Trivedi,
Spatial inhomogeneities in disordered d-wave superconductors.
Physical Review B **63**, 020505 (2000). [Link](#)

73. Tanmoy Das, Jian-Xin Zhu, Matthias J. Graf,
Local suppression of the superfluid density of PuCoGa_5 by strong onsite disorder.
Phys. Rev. B **84**, 134510 (2011). [Link](#)
74. Serena Eley, Sarang Gopalakrishnan, Paul M. Goldbart, Nadya Mason,
Approaching zero-temperature metallic states in mesoscopic
superconductor–normal–superconductor arrays.
Nature Physics **8**, 59–62 (2011). [Link](#)
75. Beena Kalisky, Julie A. Bert, Brannon B. Klopfer, Christopher Bell, Hiroki K. Sato,
Masayuki Hosoda, Yasuyuki Hikita, Harold Y. Hwang, Kathryn A. Moler,
Critical thickness for ferromagnetism in $\text{LaAlO}_3/\text{SrTiO}_3$ heterostructures.
Nature Communications **3**, 922 (2012). [Link](#)
76. Julie A. Bert, Beena Kalisky, Christopher Bell, Minu Kim, Yasuyuki Hikita,
Harold Y. Hwang, Kathryn A. Moler,
Direct imaging of the coexistence of ferromagnetism and superconductivity
at the $\text{LaAlO}_3/\text{SrTiO}_3$ interface.
Nature Physics **7**, 767–771 (2011). [Link](#)
77. Katja C. Nowack, Eric M. Spanton, Matthias Baenninger, Markus König,
John R. Kirtley, Beena Kalisky, C. Ames, Philipp Leubner, Christoph Brüne,
Hartmut Buhmann, Laurens W. Molenkamp, David Goldhaber-Gordon,
Kathryn A. Moler,
Imaging currents in HgTe quantum wells in the quantum spin Hall regime.
Nature Materials (2013). [Link](#)
78. Clifford William Hicks,
“Scanning SQUID microscopy on iron pnictide and time–reversal
symmetry–breaking superconductors”,
PhD thesis, Stanford University, 2009. [Link](#)
79. J. R. Kirtley, C. Kallin, C. W. Hicks, E.-A. Kim, Y. Liu, K. A. Moler, Y. Maeno,
K. D. Nelson,
Upper limit on spontaneous supercurrents in Sr_2RuO_4 .
Physical Review B **76**, 014526–8 (2007). [Link](#)

80. Clifford W. Hicks, John R. Kirtley, Thomas M. Lippman, Nicholas C. Koshnick, Martin E. Huber, Yoshiteru Maeno, William M. Yuhasz, M. Brian Maple, Kathryn A. Moler,
Limits on superconductivity-related magnetization in Sr_2RuO_4 and $\text{PrOs}_4\text{Sb}_{12}$ from scanning SQUID microscopy.
Physical Review B **81**, 214501 (2010). [Link](#)
81. Rafael Baruch Dinner,
“Scanning Hall probe microscopy of supercurrents in YBCO films”,
PhD thesis, Stanford University, 2006. [Link](#)
82. Per G. Bjornsson,
“Low-temperature scanning magnetic probe microscopy of exotic superconductors”,
PhD thesis, Stanford University, 2005. [Link](#)