SOURCES OF NOISE IN NIOBIUM-BASED SUPERCONDUCTING QUANTUM CIRCUITS

BY

ALLISON DOVE

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2014

Urbana, Illinois

Doctoral Committee:

Associate Professor Nadya Mason, Chair
Professor James Eckstein, Director of Research
Professor Anthony Leggett
Assistant Professor Liang Yang

Abstract

Quantum computation is a fascinating field that combines novel physics with improvements in computation times and has been rapidly growing in the past few years. Using superconductors to form the qubits has the potential for large-scale computing, if the decoherence inherent in these devices can be understood and reduced. Two-level fluctuators due to defects in the materials are thought to cause changes to the Josephson critical current or the flux through the superconducting loop of a flux qubit, which leads to decoherence in the qubit. Alternatively, defects in the crystal lattice give rise to electron localization which in turn traps spins with random orientations at the substrate/metal interface, again producing decoherence.

My work studied the proposed noise mechanisms by using Molecular Beam Epitaxy to fabricate single-crystal niobium-based Josephson junctions. Using RHEED, AFM, and TEM/STEM, I studied the epitaxy of the niobium film at the substrate interface to reduce noise due to crystal defects. I then measured flux noise in these films; the noise in the epitaxial films is lower than the comparative polycrystalline films. Further measurements using *ex-situ* Josephson junctions and epitaxial niobium loops resulted in the lowest reported flux noise measurements to date. Additionally, EELS measurements made in the course of the STEM analysis of the crystal structure reveal oxygen depletion from the substrate at elevated growth temperatures. This depletion leads to oxygen vacancies in the aluminum oxide substrate, which can in turn lead to charge traps at the substrate/metal interface and hence to decoherence in the qubit.

Critical current noise was investigated by changing the oxidation dose of the alumina barrier in more than 50 separate Josephson junctions. Films which were grown with lower oxidation doses have a non-ideal critical current temperature dependence as well as higher critical current noise values, suggesting that insufficient oxygen is being incorporated into these junctions.

Finally, I studied the effect of surface cleaning after fabrication for resonators and transmons. Microscopy shows that an oxygen ash and BOE dip together removes all photoresist residue from the chip. Internal-Q measurements do not show a marked improvement due to this cleaning, but the coherence times do improve upon the cleaning step.

Acknowledgements

This work would not have been possible without assistance from many people. I would first like to thank my advisor, Jim Eckstein, for his advice and support during this project.

I would also like to thank my close colleagues Zachary Yoscovits and Gus Olson who were also studying noise in superconducting qubits. I am especially grateful for their assistance in the cleanroom while I was pregnant, as well as their invaluable assistance keeping System E operational. I would also like to thank my lab mates Brian Mulcahy, Carolyn Kan, Can Zhang, and Chi Xue for letting me brainstorm in their office. To my former lab mates Stephanie Law, Chandra Mohaptra, Mao Zheng, Mike Vissers, and Tim McCardle, thank you for teaching me and being patient while I learned.

I also want to thank Steven Anton (Clarke group, UCB), Steven Weber and Kater Murch (Siddiqi group, UCB), and Chris Nugroho and Vladimir Orlyanchik (Van Harlingen group, UIUC) for measuring devices that I fabricated, as well as their advisors John Clarke, Irfan Siddiqi, and Dale Van Harlingen for facilitating my measurements in addition to their own.

I would also like to thank the staff at the Frederick Seitz Materials Research Laboratory for training me to use equipment and always being willing to help troubleshoot. I would especially like to thank Scott MacLaren (AFM), J.G. Wen and Amish Shah (TEM), Tao Shang and Bharat Sankaran (Raith and MicroFab) and Doug Jeffers (MBE maintenance). I would have gotten nowhere without your expertise.

I would finally like to thank my husband, Noah, for suffering through the many years of frustrating PhD work. While it has been a bumpy road, I truly appreciate that you put up with it all – I could not have done this without your support. I would also like to thank my daughter Sylvia; you've only been here at the end, but your presence has been truly motivating! To my parents, thank you for always encouraging me even when you worried that I was striking out on a difficult path.

This material was based in part on work supported by IARPA. This work was carried out in part in the Frederick Seitz Materials Research Laboratory Central Facilities, University of Illinois.

Table of Contents

List of Abbreviations vii

Chapter 1 Introduction 1

Chapter 2 Superconducting Qubits 4

2.1 Introduction to Superconducting Qubits 4

2.1.1 The Josephson Junction 4

2.1.2 The Flux Qubit 6

2.1.3 The Transmon and 3D Transmon 8

2.2 Decoherence in Qubits 10

Chapter 3 Molecular Beam Epitaxy Growth Techniques 14

3.1 Vacuum Chamber Overview 14

3.2 Sapphire Substrates 16

3.3 Growth of the Niobium Base Electrode 16

3.4 Growth of the Aluminum Buffer Layer 17

3.4.1 Reflective High-Energy Electron Diffraction 18

3.5 Growth of the Alumina Layer 21

3.6 Post-Alumina Diffusion Process 22

3.7 Growth of the Niobium Top Electrode 23

3.8 Summary 23

Chapter 4 Materials Analysis 24

4.1 Transport Measurements of Niobium Thin Films 24

4.2 Thin Films of Single Crystal Niobium on A-Plane Sapphire 26

4.3 Reflection High-Energy Electron Diffraction and Atomic Force Microscopy Measurements 27

4.4 Transmission Electron Microscopy Measurements 31

4.4.1 TEM Techniques 31

4.4.2 STEM Analysis of the Niobium-Sapphire Substrate Interface 33
4.4.3 EELS Analysis of the Niobium-Sapphire Substrate Interface 36
4.4.4 EELS Analysis of the Aluminum Oxide Barrier 40
4.5 Summary 41
Chapter 5 Sample Fabrication 42
5.1 Lithography Techniques 42
5.1.1 Photoresist Application 42
5.1.2 Deep Ultraviolet Exposure 42
5.1.3 Electron-Beam Exposure 43
5.1.4 Hybrid Exposures 44
5.2 “Via”-Process Fabrication Techniques 44
5.2.1 Device Isolation 44
5.2.2 Mesa Etch 46
5.2.3 Passivation Deposition and Via-Hole Etch 47
5.2.4 Wire-Up Deposition 48
5.3 Self-Aligned Fabrication 50
5.3.1 Device Isolation 50
5.3.2 Wire-Up Deposition and Etch Back 52
5.4 Resonators and SQUIDs 54
Chapter 6 Transport Characteristics and Noise Measurements 55
6.1 Resonators 55
6.2 Josephson Junctions 59
6.2.1 Superconducting Order Parameter and Subgap Conductance 60
6.2.2 Josephson Junction Critical Current Measurements 67
6.3 Flux Noise Measurements of Epitaxial SQUIDs 72
Chapter 7 Sample Cleanliness 78

7.1 AFM of Unprocessed and Processed Films 78
7.2 Cleaning Etches and Resulting Microscopy 80
7.3 Internal-Q and Coherence Measurements of Cleaned Samples 88
Chapter 8 Summary and Future Directions 90
Appendix A The Shadow-Evaporation Technique 94
References 96

List of Abbreviations

ADF – Annular Dark Field

AFM – Atomic Force Microscopy

ALD – Atomic Layer Deposition

BCS – Bardeen Cooper Schrieffer

BF – Bright Field

BOE – Buffered Oxide Etchant

DUV – Deep UltraViolet

FCC – face centered cubic

EELS – Electron Energy Loss Spectroscopy

HAADF – High Angle Annular Dark Field

HRTEM – High Resolution Transmission Electron Microscopy

LEEM – Low Energy Electron Microscopy

NIST – National Institute of Standards and Technology

NMR – Nuclear Magnetic Resonance

MBE – Molecular Beam Epitaxy

MIGS – Metal Induced Gap States

MRL – Materials Research Laboratory

PECVD – Plasma Enhanced Chemical Vapor Deposition

PMMA – Polymethyl Methacrylate

PR – Photoresist

QCM – Quartz Crystal Monitor

RCA – Radio Corporation of America

RHEED – Reflection High Energy Electron Diffraction

RIE – Reactive Ion Etch

RMS – Root Mean Squared

RRR – Residual Resistance Ratio

SEM – Scanning Electron Microscope

SQUID – Superconducting QUantum Interference Device

STEM – Scanning Transmission Electron Microscopy

TCE -- Trichloroethylene

TEM – Transmission Electron Microscopy

UCB – University of California Berkeley

UHV – Ultra-High Vacuum

Chapter 1 Introduction

There are certain types of computations that are difficult to compute even on today's great supercomputers. Problems like the traveling salesman problem or integer factorization can be solved more efficiently with quantum algorithms than classical algorithms. Quantum algorithms such as those devised by Deutsch-Jozsa [1], Grover [2], or Shor [3] produce a significant computing power increase when factoring large numbers or conducting a search through an unstructured search space [4]. These developments have spurred interest in the development of quantum computers, both for security interests and for its intrinsic scientific value.

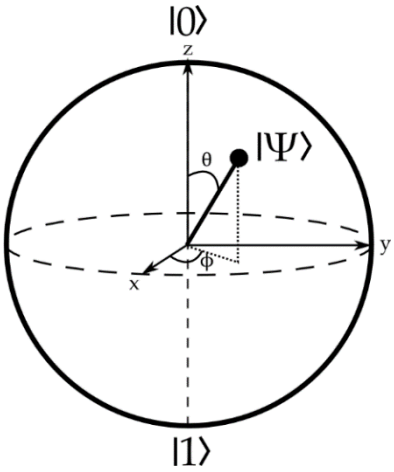

Figure 1 Bloch sphere representation of a superconducting qubit. The dot represents the state of the qubit at a given time, and the state can precess about the z-axis with frequency ω. Rotations about the x, y, and z axes represent different qubit operations.

Instead of traditional silicon-based transistors, quantum computers use quantum bits (qubits). A physical qubit can be made from nearly any two-level quantum system; this work will focus solely on qubits made using superconducting materials. At first glance, the two-level quantum states look like the zero and one states of a classical bit, but quantum mechanics allows for a superposition of states in a qubit. The general state of a single qubit is $a|0\rangle + b|1\rangle$ where a and b are complex numbers. Measurement of the qubit collapses the system to either the ground state or the excited state with probability $|a|^2$ or $|b|^2$ respectively and with the requirement that $|a|^2 + |b|^2 = 1$. Because the probabilities must sum to unity, the general state of a single qubit may also be written as $|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle$. This state can be represented on the Bloch sphere as shown in Figure 1. The state will precess in time about the z-axis; the precession rate ω is governed by the energy difference between the ground and excited states. Any deviation in the energy level spacing changes the time evolution of the qubits and the system decoheres.

There are several conditions that must be met to actually physically make a quantum computer; this set of five conditions is the DiVincenzo Criteria [5] and is described below.

DiVincenzo Criteria:

1. A scalable physical system with well characterized qubits.
2. The ability to initialize the state of the qubits to a simple fiducial state
3. Long coherence times
4. A "universal" set of quantum gates
5. A qubit specific measurement capability

The first item on the list dictates which quantum systems are appropriate for quantum computing. The individual qubits must have two distinct energy states that can be independently addressed from any other quantum states in the system. In Section 2.1.1 I shall demonstrate how Josephson junctions address this point. Furthermore, the choice for the qubit system should be scalable; superconducting qubits use the same fabrication techniques as today's integrated circuits so they are, in theory, easily scalable to thousands of chips.

The second requirement is that it must be possible to initialize the system to some ground state. With superconductors, this is achieved by cooling the qubits such that $k_B T$ is much lower than the difference between the ground and excited states. The fourth and fifth items (gates and readout, respectively) are realized by coupling the qubit to various electrical signals that can be incorporated onto the chip.

The third criteria, that the qubits have long coherence times, has proven to be the most challenging for solid-state qubits. There are a multitude of mechanisms by which superconducting qubits can interact destructively with their environment; this decoherence is one of the key hurdles that must be surmounted before a viable quantum computer can be made. I have studied the causes of decoherence in niobium-based superconducting devices from a materials perspective, using microscopy techniques to analyze the effect of temperature on niobium grown on sapphire, measuring the transport in Josephson junction to determine how best to grow the insulating barrier, measuring flux noise in epitaxial niobium SQUIDs, and using transmon qubits to evaluate the effectiveness of sample cleaning on reducing noise. This work used Molecular Beam Epitaxy (MBE) as the main tool by which the growth parameters for these devices could be controlled.

In this work, I will first outline the underlying physics governing superconducting qubits and their key component, the Josephson junction. I will then describe the growth of "trilayer" films using my MBE process. Even before devices are fabricated, these films can be analyzed using various microscopy techniques to better understand the effect of temperature or oxidation dosage. In Chapter 5 I discuss

the two fabrication processes that I used to create Josephson junctions, SQUIDs, resonators, and qubits. Transport measurements from these devices are shown in Chapter 6. Chapter 7 is a discussion of the effect of various wet and dry etches on the morphology of the films and reports data on the effect of cleaning on resonators and transmons. Finally, I summarize my findings and suggest future directions for understanding noise in niobium-based qubits.

Chapter 2 Superconducting Qubits

2.1 Introduction to Superconducting Qubits

There are three main types of superconducting qubits: the phase qubit [6], the charge qubit [7], and the flux qubit [8]. In recent years additional qubit geometries such as the transmon and quantronium have also been developed, but these newer structures are actually refinements of charge qubits [9, 10]. In this chapter I will describe the basic component of all superconducting qubits, the Josephson junction, and the physics dominating the flux qubit and the 3D transmon. I will also note the different ways in which decoherence can affect the different qubit systems.

2.1.1 The Josephson Junction

A Josephson junction is formed by separating two superconducting electrodes with a thin insulating layer. Each superconducting electrode is a condensate of Cooper pairs, which can be described as a single multi-particle wavefunction. When the insulating layer between the electrodes is thin enough that the wavefunctions can interact then Cooper pairs tunnel from one electrode to the other. The magnitude of the Cooper pair current is the Josephson supercurrent and is equal to $I = I_c \sin\phi$ where ϕ is the phase difference between the two superconducting wavefunctions and I_c is the maximum value of the critical current. This supercurrent flows without any corresponding voltage, one of the special properties of the Josephson junction. Nonzero voltages across the barrier are due to quasiparticle tunneling and are related to the phase of the superconducting electrodes by the fundamental voltage relation: $V = \frac{\hbar}{2e}\frac{d\phi}{dt}$.

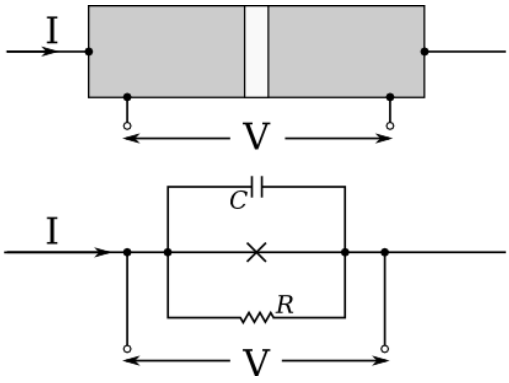

Figure 2 Schematic of a Josephson junction. (Top) The superconducting electrodes are gray, while the insulating barrier is shown as white. (Bottom) Circuit diagram for a junction. The cross is an ideal Josephson junction with intrinsic capacitance and resistance in parallel.

From these equations it is possible to see why the Josephson junction is the fundamental building block of superconducting qubits. From an engineering point of view, any quantum harmonic oscillator could make a fine qubit, except that the energy levels are evenly spaced in a simple harmonic oscillator. The Josephson junction, in contrast, behaves like a nonlinear inductor so that the energy levels are not evenly spaced in Josephson junction-based qubits. The inductance of a junction can be found by considering the Ohm's law for inductors: $V = L\frac{dI}{dt}$. Plugging in the current relation for a Josephson junction produces $\frac{dI}{dt} =$

$I_c \frac{d\phi}{dt} \cos\phi$ and using the fundamental voltage relation shows that the inductance for a Josephson junction is

$$L = \frac{\hbar}{2eI_c \cos\phi}$$

The energy of a Josephson junction can also be calculated from the voltage and current relations. Integrating the power through the junction gives $E = constant - E_J \cos\phi$, where $E_J = \frac{\hbar I_c}{2e}$ is the Josephson coupling energy. Of course, a Josephson junction circuit also includes capacitance from the small separation of the two metal conductors and some shunt resistance across the barrier acting in parallel with the junction inductance. Summing the current through the different components produces the differential equation $C \frac{\hbar}{2e} \frac{d^2\phi}{dt^2} + \frac{1}{R} \frac{\hbar}{2e} \frac{d\phi}{dt} + I_c \sin\phi = 1$. The equation can be rewritten as $C \left(\frac{\hbar}{2e}\right)^2 \frac{d^2\phi}{dt^2} + \frac{1}{R} \left(\frac{\hbar}{2e}\right)^2 \frac{d\phi}{dt} + \frac{\partial U(\phi)}{\partial \phi} = 0$ with the potential $U(\phi) = -E_J \left(\cos\phi + \frac{I}{I_c}\phi\right)$. This is similar to the Josephson energy calculated earlier but with an additional linear bias. The energy potential $U(\phi)$can be understood with the "washboard analogy" where a particle with mass rolls down a hill as shown in Figure 3. When $I < I_c$, the potential has shallow wells which can trap the "particle" resulting in no change in the phase and correspondingly zero voltage. When $I > I_c$, the particle can roll down the potential and there is nonzero voltage in the system. There is no direct analog for the particle in a Josephon junction itself; rather the position of the particle becomes the wave function in the fully quantum mechanical system.

If the system is already in the nonzero voltage state and the current is reduced below the critical current then the particle still has some momentum and the junction stays in the nonzero voltage state until it loses enough momentum to be trapped in the well. This model helps explain the hysteresis observed in the voltage-current characteristics of Josephson junctions (see Section 6.2).

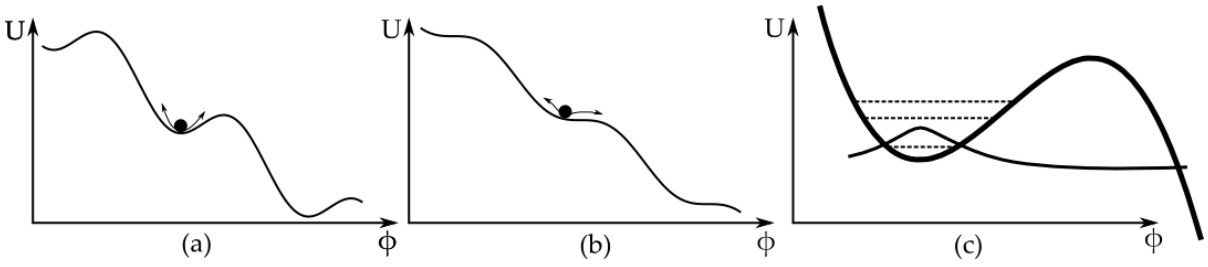

Figure 3 Energy potential for Josephson junctions with different current biases. Plot (a) shows the washboard potential when $I < I_c$. Plot (b) shows the washboard potential when $I > I_c$. Plot (c) shows the quantized states in a single shallow well. These states are used in the phase qubit and are included here to demonstrate the "quantum-ness" of the Josephson junction.

Additional features of the Josephson junction shall be discussed in greater detail in later chapters. Of particular interest will be the transport characteristics (Section 6.2.1) and temperature dependence (Section 6.2.2) of the junctions that I fabricated. The theory of transport and temperature dependence will be discussed in more detail there; in this chapter, I limit the discussion to the aspects of Josephson junctions that make them useful for quantum computation.

2.1.2 The Flux Qubit

To understand decoherence in superconducting qubits, it is first necessary to understand the energy landscape of the qubit, as this will determine how strongly different defects couple to the qubit. For this project I focused on causes of decoherence in flux qubit devices, so I will describe the energy of flux qubits in this section.

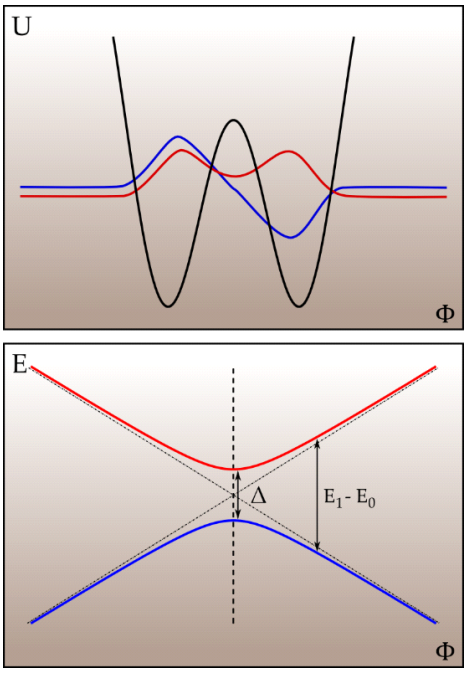

Figure 4 (Top) The black line shows the energy potential for the single junction flux qubit. The blue line shows the qubit ground state and the red line shows the first excited state. (Bottom) The energy splitting of the two lowest energy states. The vertical dashed line notes the degeneracy point $\Phi_0/2$. Reproduced from Reference [16].

The first system to consider is the single-junction flux qubit. When a Josephson junction is placed in a superconducting loop, it becomes a Superconducting QUantum Interference Device (SQUID). When the loop is placed in a magnetic field such that there is an external flux Φ_x through the loop, the flux through the loop is related to the phase across the junction by the relation $\phi = 2\pi n - 2\pi\Phi/\Phi_0$ where $\Phi_0 = h/2e$ is the flux quantum, the total flux is $\Phi = \Phi_x + LI$ and *n* is an integer. As with the single junction, summing the current around the loop produces the differential equation

$$C\frac{d^2\phi}{dt^2} + \frac{1}{R}\frac{d\phi}{dt} + I_c \sin\left(\frac{2\pi\Phi}{\Phi_0}\right) + \frac{\Phi - \Phi_x}{L} = 0$$

In the above equation, the phase-flux relation has been used to rewrite all instances of the phase. Again, as with the single junction, the last two terms can be written as $\partial U(\Phi)/\partial\Phi$. Doing so gives the potential for the single-junction flux qubit as

$$U(\Phi) = \frac{(\Phi - \Phi_x)^2}{2L} - E_J \cos\left(\frac{2\pi\Phi}{\Phi_0}\right)$$

When the external flux is one-half flux quantum, the potential forms two symmetric wells. Each well can have bound states located in it, but the qubit states are actually superpositions of the left and right wells. For appropriate choices of inductance and critical current, the wavefunctions of each well can

interfere forming symmetric and anti-symmetric states as shown in Figure 4. The energy separation of these wavefunctions can be defined as $\hbar\Omega$. Work by Van Harlingen *et al* [11] demonstrates that small shifts in flux will change the symmetry of the potential so that the energy separation then changes such that $\Delta\Omega \sim (\Delta\Phi)^2$. In comparison, changes in critical current produce an exponential shift in the energy spacing $\Delta\Omega \sim \exp(\Delta I_c)$. The comparison of these two fluctuations shows that the flux qubit is fairly insensitive to changes in flux while it is sensitive to changes in the critical current.

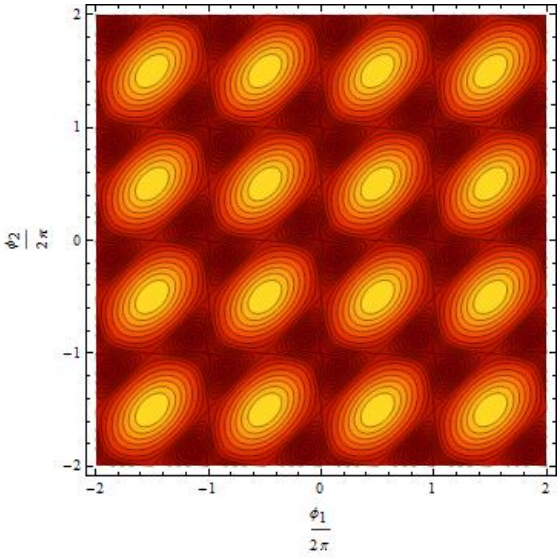

Figure 5 Potential landscape for the 3-junction flux qubit. The 3-dimensional double wells (dark red) such as those about (0,0) form the potential wells similar to the one-dimensional double-well for the single junction flux qubit pictured in Figure 4. Reproduced from Reference [15].

While the single-junction flux qubit is relatively easy to understand and model, in practice it has significant hurdles before it can be utilized in a quantum computer. Primarily, the one-junction flux qubit must have a large enough self-inductance such that the condition $2\pi LI_c/\Phi_0 \gg 1$ is met in order to form the shallow wells that form the basis of the qubit [12]. Unfortunately, having a large self-inductance requires loops diameters that are hundreds to thousands of microns across [13]. Even more problematic, these large loops are then more susceptible to external noise [14], by virtue of their size.

An alternative was found by using the Josephson junction again. As noted earlier, a Josephson junction has an inherent nonlinear inductance associated with it. Mooij *et al* [8] used this inductance to create a three-junction flux qubit that was considerably smaller than the one-junction qubits made to date. By making small loops the self-inductance of the loop was negligible.

In comparison with the one-junction qubit, the energy potential for the three-junction flux qubit is considerably more complex. Flux quantization requires that the flux-phase relation be written as $\phi_1 - \phi_2 + \phi_3 = -2\pi\Phi_x/\Phi_0$ where ϕ_i is the phase difference across each junction. In the three-junction flux qubit designed by Mooij *et al*, the junctions were designed such that junctions 1 and 2 had equal Josephson energy and capacitance while junction 3 had a Josephson energy and capacitance α times larger [15]. The energy of each junction is $E_{jn}(1 - \cos\phi_n)$ and when the energies are summed around the circuit, the potential for the system is

$$U = E_J[2 + \alpha - \cos\phi_1 - \cos\phi_2 - \alpha\cos(\phi_2 - \phi_1 + 2\pi\Phi_x/\Phi_0)]$$

For a specific external flux, this potential is dependent on both ϕ_1 and ϕ_2 and can be plotted as the contour plot shown in Figure 5. As with the one-junction flux qubit, the three-junction flux qubit has double-wells that can have localized states bound to them for appropriate choice of the tuning parameter α. And as with the one-junction flux qubit, Van Harlingen *et al* considered the critical current fluctuations for the three-junction flux qubit. These calculations showed that for the system described above, the small junction (junction 3) would provide the largest contribution to critical current noise. I shall now discuss how fluctuations in the energy potential for the flux qubit lead to decoherence and loss of information in quantum computing.

2.1.3 The Transmon and 3D Transmon

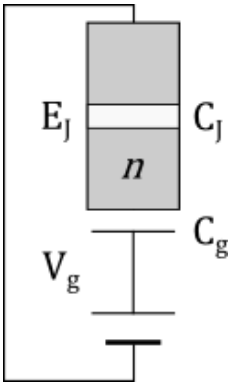

Figure 6 Schematic of a charge qubit. The color scheme is the same as in Figure 2 Reproduced from Reference [14].

In this section I will describe the physics of the transmon qubit, which was first developed by Robert Schoelkopf's group at Yale [9]. The transmon qubit is sensitive to different types of noise than the flux qubit, but is a relatively simple device to fabricate and measure, making it an attractive system to use to evaluate what materials to use in qubits of any type. The transmon is actually derived from the charge qubit, so I will begin by discussing the charge qubit.

Charge qubits work by having such a small superconducting island (denoted by the label *n* in Figure 6) that the charging energy of placing two electrons on the island is much larger than the thermal energy [16]. The charging energy is $E_c = e^2/2C_{total}$ where the sum of the junction and gate capacitances is $C_{total} = C_J + C_g$. The qubit states correspond to the number of Cooper pairs on the island, so *n* is the operator for the number of states on the island. The gate voltage is also accounted for by the number of charges on the gate, namely $n_g = C_g V_g/2e$. Neglecting any loss, the Hamiltonian for the charge qubit is $H = 4E_c(n - n_g)^2 - E_J \cos\phi$.

The first term represents the total electrostatic energy of the island and the second term is the familiar Josephson junction energy. The energy splitting for the lowest three states are shown in Figure 7, plotted for different values of the gate charge n_g. The energy splitting is highly sensitive to charge

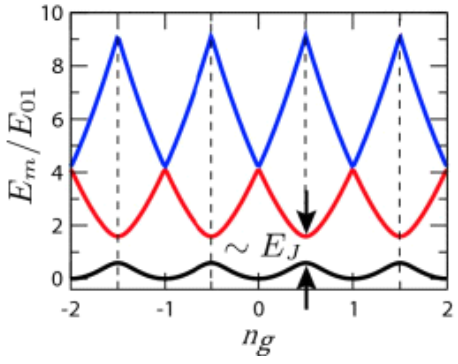

Figure 7 Energy level splitting for a charge qubit. Figure from Reference [9].

fluctuations, although the sensitivity is minimized at the bias point of $n_g = 1/2$.

The transmon and the 3D transmon are modifications to the charge qubit which reduce its charge sensitivity [9]. In the transmon, the superconducting island is shunted by a large capacitance C_B, which increases the total capacitance of the circuit $C_{total} = C_J + C_g + C_B$. This in turn lowers the charging energy of the circuit. The Hamiltonian for the transmon is identical to the original charge qubit, but the energy levels depend on the ratio of the Josephson energy to the charging energy E_J/E_c. When $E_J/E_c \gg 1$, the energies are nearly flat for all possible values of the gate charge n_g, leaving the device insensitive to charge fluctuations (see Figure 8). While increasing the ratio E_J/E_c insulates the transmon from charge noise, it also reduces the anharmonicity in the qubit, raising the concern that additional energy states can be excited during qubit operation. Functionally, however, this design takes advantage of the fact that the sensitivity to noise is reduced exponentially, while the anharmonicity is decreased algebraically with a slow power law dependence [9].

The three-dimensional transmon is a further modification on the charge qubit. Where the transmon qubit uses an interdigitated capacitor to shunt the superconducting island, the 3D transmon uses two large paddles for the shunting capacitance (see Figure 9). The 3D transmon has the benefit of not requiring any bias circuitry [17], but the true advantage of this design is that it is operated in a three-dimensional cavity with the paddles simultaneously providing a shunting capacitance and acting as dipole antenna to the cavity. The large size of the paddles also serves to limit the qubit sensitivity to

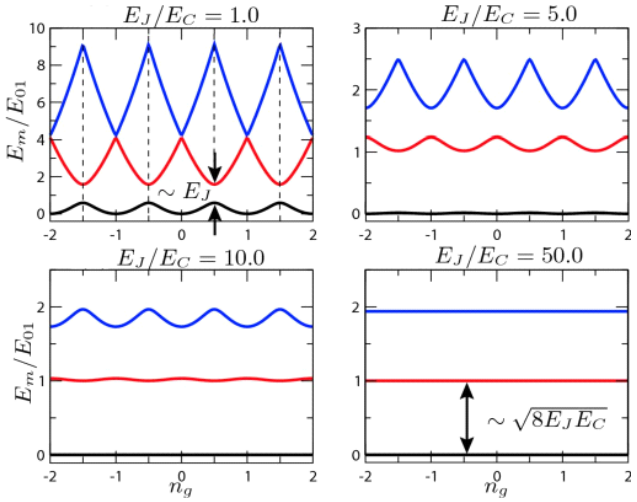

Figure 8 Energy levels for the transmon qubit for different values of the parameter E_J/E_c. As the ratio of Josephson energy to charging energy increases, there is less dispersion in the energy levels for different charges on the gate capacitor. Figure from Reference [9].

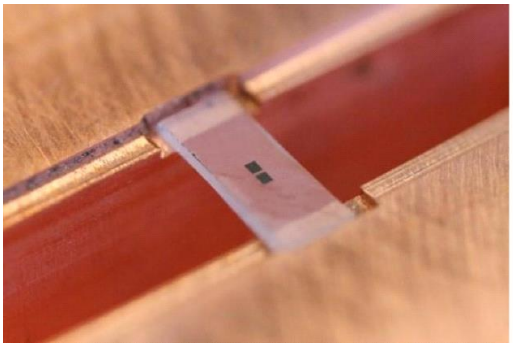

Figure 9 (Left) Photo of a 3D transmon mounted in its copper cavity. The paddles are visible -- the Josephson junction is between the two paddles. Photo courtesy Irfan Siddiqi, UCB. (Right) Schematic showing the Josephson junction (in white) between the paddles. Not to scale.

surface dielectric loss [17]. The simplicity of the 3D transmon qubit makes it an extremely attractive system for studying noise due to materials in superconducting qubit systems.

I will now move on to discuss the physical causes of noise in superconducting qubits.

2.2 Decoherence in Qubits

The primary problem preventing a large scale implementation of a quantum computer is the problem of decoherence. Since quantum computation operates on the entire ensemble of qubits, any deviation from any single qubit will lead to computational errors. In trapped-ions systems, with little environmental interaction, these deviations are relatively easy to contain. In solid-state systems like superconducting qubits, these deviations place severe limitations on current quantum computers.

Decoherence generally refers to any process by which the quantum system randomly fluctuates and loses information about the initial state. This can happen either by a relaxation process or a dephasing process. Relaxation occurs if the qubit couples to a random fluctuator and the system loses energy. These fluctuations can be two-level systems that occur due to material defects in the qubit, or these fluctuation can be due to external environmental effects. Any type of lossy material making up the qubit can cause relaxation of the superposition to the ground state. In contrast, dephasing refers to fluctuations that change the energy landscape and thus change the qubit frequency. Changes in the energy landscape of the qubit inherently change the qubit frequency, leading to small dephasing times. As in nuclear magnetic resonance (NMR) spectroscopy, the time associated with relaxation processes is T_1 while the time associated with dephasing processes is T_2 [16]. Qubit lifetime is limited by the T_1 time; the ability to allow two or more qubits to constructively interfere for a computation is limited by the T_2 time. The dephasing time depends on the relaxation time T_1 as $\frac{1}{T_2} = \frac{1}{T_1} + \frac{1}{T_\phi}$ where T_ϕ is the pure dephasing time absent any relaxation processes.

Direct measurement of the T_1 and T_2^* decoherence times can be done using standard NMR techniques such as Rabi oscillations, Ramsey fringes, or spin-echo experiments; T_2^* being the observed measurement of dephasing instead of T_2. Indirectly, sources of noise can be analyzed by performing

microscopy of materials, fabricating and measuring the quality-factor of resonators, and measuring critical current noise in fabricated Josephson junctions. This work will primarily draw on the latter measurements, but T_1 and T_2^* times will also be referenced in Chapter 7. The initial chapters focus on sources of noise that lead to dephasing, but the Chapter 7 does refer to noise mechanisms that can allow simple relaxation of the qubit.

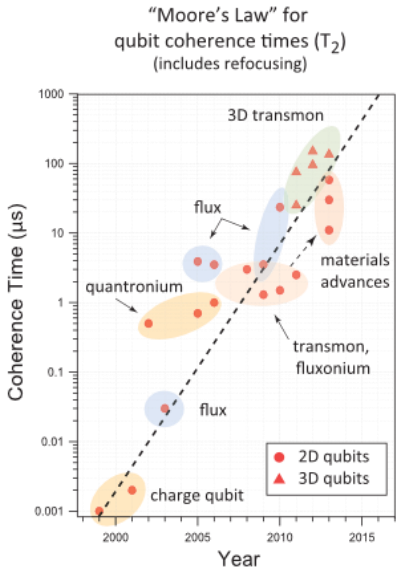

Figure 10 Coherence times for different types of superconducting qubits over a range of 1 nanosecond to >100 microseconds. The x-axis shows the publication year. Figure from [18].

Decoherence has been substantially reduced in superconducting qubits in recent years. In a review, Oliver and Welander [18] reported T_2 times for a myriad of different qubit devices as a function of the publication year (see Figure 10). Most of the increase in coherence times is due to new qubit designs which are less sensitive to environmental fluctuations, i.e. the transmon is less sensitive to charge noise than the charge qubit and the 3D transmon is less sensitive to dielectric loss than the transmon. Naturally, qubit designs which are inherently less sensitive to defects and fluctuators have greatly improved coherence times, but there is a limit at which defects will still couple with the electromagnetic fields required to operate the qubit. Further performance improvements likely depend on actual removal of defects and two-level fluctuators from the devices.

Decoherence in superconducting qubits is generally associated with flux noise, critical current noise, or charge noise, depending on the type of qubit in question. The phenomena of *1/f* flux noise in dc-SQUIDs and flux qubits has been well established [19, 20] although the mechanisms driving these fluctuations are still not completely understood. Measurements of the spectral noise density $S_\Phi^{1/2}(1\ Hz)$ for dc-

SQUIDs grown from niobium, lead, lead-indium, or aluminum reveal noise powers in the range of $1 - 20\ \mu\Phi_0 Hz^{-1/2}$ across 5 orders of magnitude in the size of the SQUID loops [21]. The size and materials insensitivity of this flux noise is part of the challenge in modeling the microscopic theory of flux noise.

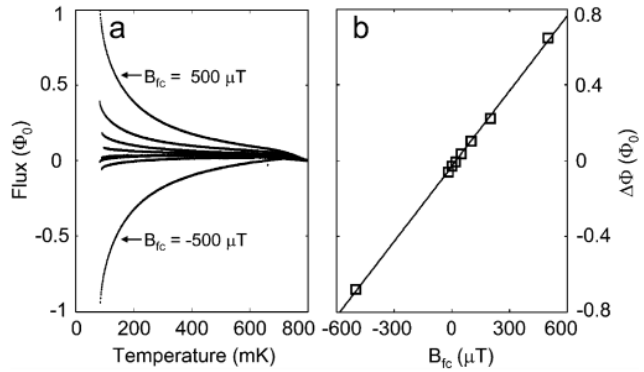

Figure 11 Temperature dependence of the flux in Nb-based SQUIDs showing trapped vortices due to paramagnetic impurities. Figure from Reference [22].

Recently, measurements have been able to provide an estimate for the areal density of spin defects which give rise to flux noise. Sendelbach *et al* [22] cooled a large Nb/AlO_x/Nb SQUID in a small magnetic field (-500μT – 500μT) trapping vortices in the superconducting loop. When they were well below T_c they turned off the magnetic field and observed the change in flux in the loop with respect to temperature. The inverse dependence of the data suggested that paramagnetic impurities were giving rise to the observed signal. They suggested that surface spins were strongly coupling to the pinned vortices with a defect density of $5 \times 10^{17}\ m^{-2}$. A similar defect density of $4 \times 10^{17}\ m^{-2}$ was obtained by Bluhm *et al* [23] for films made from gold and measured using a scanning SQUID microscope [24].

An alternative source of spins was proposed by Choi *et al* [25] in which the crystal defects at the metal-substrate interface can produce random fluctuations that couple to the qubit structure. When the Fermi-energy of the metal lies in the band-gap of the insulator, there are electron states which are allowed in the metal but which are evanescent in the insulator. This naturally occurs for any metal-insulator interface and by itself, this phenomenon does not lead to flux noise. Choi *et al* proposed that when there is sufficient disorder at the interface these metal-induced gap states (MIGS) can be localized with single occupancy due to Coulomb blockade and form local magnetic moments that can interact with the electromagnetic fields that operate the qubit.

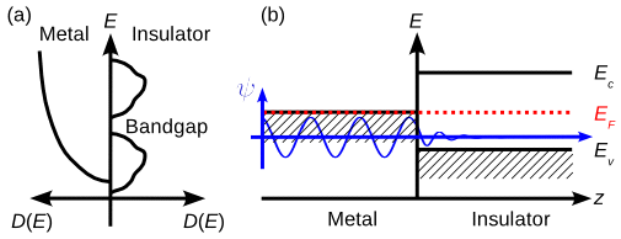

Figure 12 MIGS due to the disordered interface between the metal/superconductor and insulating substrate. Figure from Reference [25].

The Anderson localization model proposes a mechanism by which electron waves can become localized in a sufficiently disordered material [26]. In this case, disorder at the metal-insulator interface creates

enough scattering sites in the insulator so that some portion of incident electron waves become trapped. These trapping sites have only a single electron occupancy, but this leaves unpaired electron spins dotting the interface with random orientations. The spin of the electron is locked while the electron is localized, but thermal activation allows the electrons to hop on and off the defects [21]. As the electrons hop on and off defect sites they generate random telegraph signals that sum to produce the characteristic *1/f* noise power spectrum.

If flux noise is due to MIGS then reducing crystal defects at the substrate interface should reduce flux noise. I have grown single crystal niobium films on sapphire to reduce these crystal defects and hence reduce noise due to MIGS in these devices. In the following chapter, I present microscopy measurements showing the epitaxy of these samples and in Section 6.3 I present noise measurements made from these niobium films. I do a side-by-side comparison of epitaxial niobium films with polycrystalline films made by NIST and show that my films have lower flux noise, which is consistent with MIGS being a noise mechanism.

Electron localization may also occur due to stoichiometric imbalances in the insulating materials of the junction, resulting in critical current noise. Electron traps in the insulating barrier effectively change the barrier height, resulting in changes in the magnitude of the Josephson critical current [11]. Since the qubit frequency depends on the critical current, changes in the critical current result in changes in the qubit frequency leading to decoherence. I examine the effect of barrier growth on critical current noise in Section 6.2.2 by analyzing the transport characteristics of Josephson junctions. By comparing the effect of oxidation dose on the subgap conductance and the temperature dependence of the critical current, I show that films grown with low oxygen pressures and high aluminum flux rates have magnetic defects that would result in decoherence in actual qubits.

Chapter 3 Molecular Beam Epitaxy Growth Techniques

In this chapter I will describe the methods used in the growing the single-crystal niobium and aluminum films that were studied in this work. There are two types of samples described in this work: niobium-only films and full trilayer films. My "trilayer" films actually consist of 4 components: the niobium base, an aluminum buffer, the alumina barrier, and the niobium counter-electrode. The trilayer name is a reference to the creation of this type of device and is used for convenience when referencing the all-epitaxial Josephson junction. Trilayer films use all the techniques described in this chapter; niobium only films use the techniques through Section 3.3. I shall first give an introduction to the vacuum system used in this experimental work, then describe each subsequent step to produce a trilayer film.

3.1 Vacuum Chamber Overview

An ultra-high vacuum (UHV) MBE chambers was used to deposit the single crystal films studied here. This chamber was part of a multi-chamber growth facility and was connected by high vacuum tubing, allowing for *in-situ* transfer of films between growth chambers. In addition, a high vacuum oxidation chamber was also part of the growth facility, to facilitate clean oxygen dosing of samples without compromising the vacuum of the growth chambers.

The UHV deposition chamber (System E) had a base pressures of 5×10^{-11} Torr when warm, and was equipped with liquid nitrogen cryo-panels that both cooled the walls around the deposition sources and acted as a pump to further reduce the pressure inside the chambers to below 1×10^{-11} Torr. The chamber was pumped by both a cryo-pump and an ion pump, achieving a pumping speed of 1000L/s. Contamination monitoring was done by mass spectrometry with a residual gas analyzer. System E had a quartz crystal monitor that measured the atomic fluxes from the various sources. It also had a reflection high-energy electron diffraction (RHEED) system that will described in greater detail below.

Substrates were mounted on a sample heater capable of reaching temperatures up to 1350°C facing a source flange containing the thermal effusion cells and gas source. The effusion cells had the source material sitting at the bottom of a crucible and were heated by a resistive filament winding around the outside. The effusion cells were capable of heating material up to 1500°C. Refractory metal deposition, which required substantially higher temperatures, was done by electron-beam bombardment using an e-gun. In an e-gun the source material sits in a hearth where 10kV electrons can be directed by electromagnets. The heat from this bombardment is sufficient to sublimate the source material and

deposit a film on the substrate. In this work, aluminum was deposited by effusion cell, whereas niobium was deposited by electron-beam bombardment.

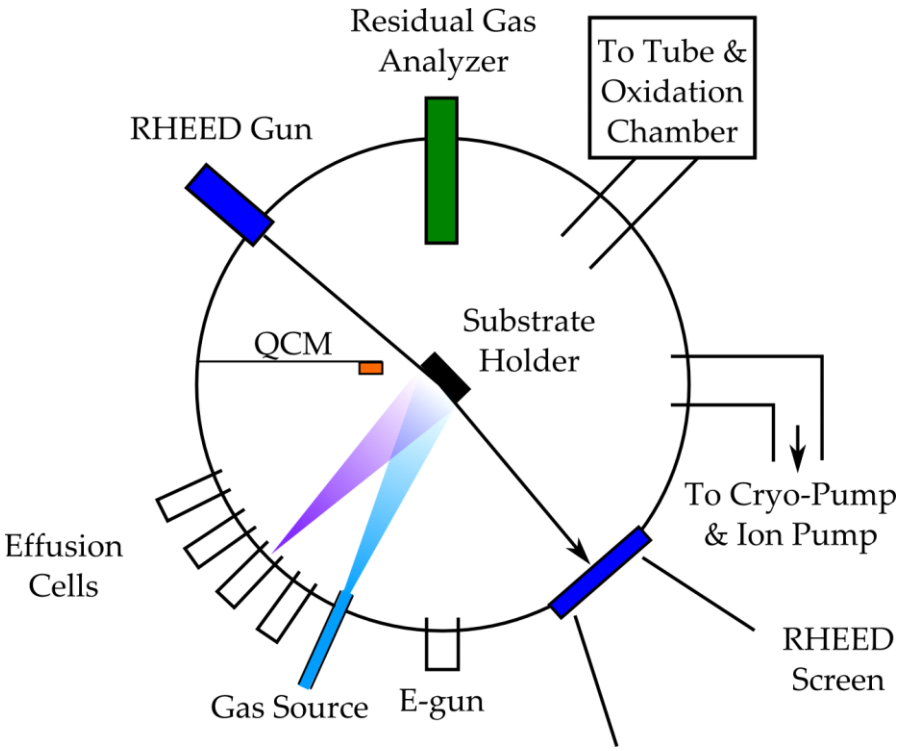

Figure 13 Schematic view of the MBE chamber used in this study. See text for chamber details.

System E had the ability to inject gas into the growth chamber to incorporate oxygen in the thin film. A bottle of ultra-high purity (99.995%) oxygen was mounted on a manifold near the chamber. A roots pump would evacuate the manifold and then reduce the inlet pressure once the gas was introduced to the chamber. A series of valves allowed for the control of the gas in the chamber at pressures as low as 10^{-9} Torr, although depositions were done at pressures of $10^{-7} - 10^{-6}$ Torr.

In addition to the gas injector system in the growth chamber, oxidation could also be done by diffusion. This was achieved by transferring the sample from the growth chamber to a separate chamber dedicated to this purpose. The pressure in the transfer tube and in the oxidation chamber during the transfer was typically around 1×10^{-9} Torr and the transfer could be done in under 1000 seconds, so there was less than 1 Langmuir of exposure, resulting in a very clean surface for diffusion. A bottle of either ultra-high purity oxygen or a mixture of 90% argon and 10% oxygen were connected to a manifold mounted on the chamber. As with system E, a roots pump would evacuate the inlet line and a precision needle valve would allow for fine control of pressure in the oxidation chamber. Pressures used were

typically between 1 Torr and 50 Torr, with diffusion times ranking between a few seconds and an hour or more.

3.2 Sapphire Substrates

Sapphire was used as the substrate in this work and offered several benefits. First, sapphire is an insulating substrate, which made device isolation trivial. Second, measurements of decoherence were done at microwave frequencies, and the substrate could potentially be a limiting factor in the measured decoherence times. Sapphire has a higher "Q" value than the next obvious candidate, silicon, and the measured decoherence times were expected to be longer with sapphire as the substrate. Third, and most important, epitaxial growth of niobium on sapphire is well understood [27, 28]. Since this work's aim was to study the relationship between single crystal films and decoherence, sapphire was the natural choice for the substrate. The epitaxy of niobium on sapphire is discussed in detail in Section 4.1.

Epitaxial growth of niobium on sapphire typically occurs at temperatures between 750°C and 900°C [28]. The heater in System E radiatively heats the samples, and as sapphire is a transparent material it requires a backside coating to absorb the supplied heat. The unpolished side of the substrate was coated with 3000Å of evaporated titanium to ensure uniform heating. The polished side was then scrubbed with trichloroethylene, and sonicated separately in baths of trichloroethylene, acetone, and isopropyl alcohol for a few minutes to remove any grease or dirt. Upon entering the vacuum, the substrate was heated above 100°C and left to out-gas water and other contaminants overnight.

3.3 Growth of the Niobium Base Electrode

Niobium was deposited by e-beam evaporation to create the base superconducting electrode in the Josephson junction trilayer. The temperature of the sapphire substrate was brought up to the desired growth point. The temperature could be measured by a thermocouple mounted behind the substrate, although this was later removed from the system and was not available for all films. The substrate temperature was also recorded by an optical pyrometer using $\varepsilon = 0.58$ as the value of optical emissivity for sapphire.

The base pressure in the chamber before growth was on the order of 10^{-11} Torr, when the e-gun source was hot, the pressure would rise to 10^{-8} Torr. Spectrometry of this environment using the residual gas analyzer showed that the primary species present in the chamber was hydrogen, likely given off by the tungsten filament. The e-gun source was operated such that the rate of niobium deposition was quite slow – the rate of growth was typically between 0.1 Å/s and 0.3 Å/s. As the electrons

impacted the niobium in the e-gun hearth, they would create a hole at the impact site. This has the effect of continuously reducing the growth rate, requiring frequent power increases to compensate for this effect. Using the QCM to monitor the thickness, 1000Å were deposited on the substrate.

Temperature readings were taken at 6 minute intervals during the growth (this time interval corresponded to approximately every 100Å). However, the temperature readings were found to be consistently higher during the growth than found immediately after the e-beam evaporator was powered off. This is either due to the amount of heat produced by the e-gun or is a function of the amount of light produced by the instrument – the hearth cannot be viewed directly but must be observed through welding glass. For consistency, the temperature of record was that measured immediately following the power-down of the e-gun but before the substrate heater was turned off. Using $\varepsilon = 0.20$ as the value of optical emissivity for niobium, the range of temperatures for films in this work was 460°C – 844°C.

Following the deposition of the niobium base electrode, the substrate heater was turned off. Prior to the growth of subsequent layers, RHEED was taken of the niobium surface. Analysis of these RHEED images is discussed in Section 4.2.2.

3.4 Growth of the Aluminum Buffer Layer

After depositing the base electrode, a buffer layer of metallic aluminum was deposited. This layer served two purposes: it allowed for careful calibration of the growth rate and thickness of the subsequent alumina layer via RHEED oscillations, and it provided a buffer so that at no point does the niobium base electrode "see" oxygen during the sample growth.

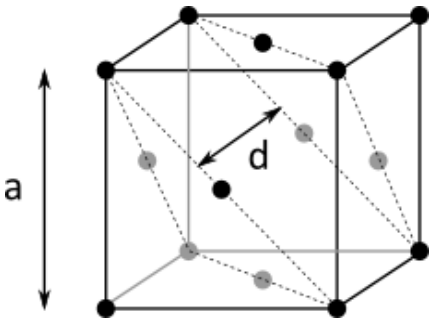

Figure 14 Face-centered cubic structure of metallic aluminum. The length a is the lattice constant and the distance d is the monolayer spacing.

While the niobium layer was deposited at temperatures upwards of 600°C, these temperatures were close to (if not exceeding) the melting point of aluminum. Work done by Flynn [29] indicates that the optimal growth temperature of metallic superlattices follows the "3/8 rule", that the growth temperature $T_g \approx 3/8\, T_m$ where T_m is the melting temperature (in Kelvin) of the metal being deposited. Given the melting temperature of aluminum is 660°C, the optimal substrate temperature would be around 76°C, or close to room temperature. The substrate was allowed to cool for 3 hours before beginning the deposition of the aluminum film. Thermocouple readings indicated that the

substrate temperature was typically around 40°C during these growths which varied in thickness between 20Å and 200Å.

Aluminum grown on (110) niobium has been shown to grow in the (111) orientation [30]. Given that aluminum has a face-centered cubic (FCC) lattice structure with lattice constant 4.05Å, the monolayer spacing in the (111) orientation is 2.34Å. In the following sections, reference to the number of monolayers will be related to thickness using 2.34Å per monolayer as the conversion factor.

Aluminum was deposited by use of a hot-filament effusion cell. The aluminum sat at the bottom of a ceramic crucible which was heated by a resistive winding around the outside of the cell. The heater was capable of reaching 1200°C and the temperature was monitored with a thermocouple. A feedback circuit controlled the power going to the cell to maintain a constant temperature during growths. Temperature was the largest factor determining the growth rate, but the growth rate was independently measured using both a quartz crystal microbalance (QCM) and by taking RHEED oscillations.

The QCM is an instrument that is mounted inside the vacuum chamber with line of sight to each of the sources. A quartz crystal is mounted at the far end of the instrument between two electrodes. An alternating current is applied to the crystal and its resonant frequency is measured. As material is deposited onto the crystal, the resonance changes, which can be converted into a thickness. The conversion factors depend on the atomic weight, as well as the geometry of the chamber. The geometry is accounted for by a tooling factor, which is found and corrected for by calibration films following chamber maintenance. The QCM was a very powerful tool, but I found that it was not a useful measurement tool once oxygen was introduced into the chamber; the oxygen would oxidize the fresh aluminum on the surface of the crystal and report a negative growth rate. RHEED oscillations were an equally powerful way to calibrate the growth rate, with the added benefit of being a measurement taken directly on the film surface with no geometrical effects.

3.4.1 Reflective High-Energy Electron Diffraction

RHEED uses a beam of high-energy (10kV) electrons to scatter off the sample surface and create a diffraction pattern on a detector screen. The electrons are incident at a very shallow angle (typically about 1 degree) and thus interact with only the topmost monolayer or two of the film. This makes RHEED very sensitive to surface roughness, and also enables RHEED to be used to monitor the rate of film growth during deposition.

The RHEED pattern is produced by constructive interference that occurs when the Laue condition is met, that is when $\vec{k}' - \vec{k} = \vec{q}$, where $\vec{k}'$ is the outgoing wave vector, $\vec{k}$ is the incoming wave vector, and $\vec{q}$ is a reciprocal space lattice vector belonging to the crystal. Because the electrons only interact with the uppermost monolayers of the film, the crystal is effectively two-dimensional. The reciprocal lattice for a two-dimensional film is a series of infinite rods perpendicular to the sample surface. When these rods intersect with the Ewald sphere, constructive interference occurs that produces the observed diffraction pattern. In the ideal case, with a perfectly flat film and monochromatic incident electrons, the diffracted pattern would be spots in a circular ring on the detector screen. In reality, the electrons have a range of energies and the films are not flat, leading to streaks instead of spots in the RHEED pattern. Broad, fuzzy streaks correspond to films with small flat regions, whereas narrow streaks indicate that the flat areas of the film are larger.

As atoms are deposited on the surface of the film, the roughness of the film varies with time. As RHEED is very sensitive to surface roughness, this can be measured by analyzing the intensity of the RHEED pattern. When half a monolayer has been deposited, the surface is maximally rough, resulting in a minimum in the pattern intensity. As more atoms are deposited, the surface levels out, leading to a maximum in the pattern intensity. The periodicity of these oscillations can be inverted to find the monolayer per second rate of deposition; if the crystal structure is known then the thickness is determined from RHEED oscillations.

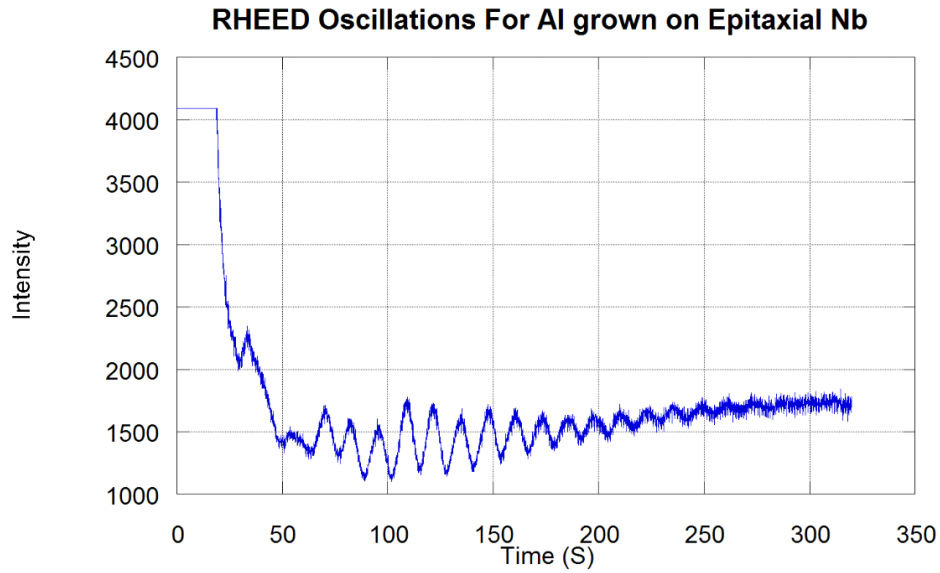

Figure 15 RHEED oscillations shown for aluminum grown on niobium.

In the case of the aluminum growth, the RHEED intensity pattern also provides some information about the growth interface on the niobium surface. Most aluminum growths began with a sudden decrease in the intensity of the RHEED pattern. For the first few monolayers of growth, no oscillations were observed, probably due to the small lattice mismatch between aluminum and niobium. After the initial few monolayers, oscillations began and frequently grew in intensity, sometimes to the point of saturating the camera. An example of a RHEED oscillation pattern following this general description is shown in Figure 15. The RHEED beam was turned off while the metallic layer was still being deposited, to avoid creating any defects due to the high energy nature of the RHEED beam. Growth of the aluminum buffer layer was terminated by masking the film from the cell.

In addition to diffraction from the crystal, Kikuchi lines can also be observed. In flat films there can be channels between rows of atoms. Electrons can travel further into the material along these channels and inelastically scatter. While only a small amount of energy is typically lost, the wavevectors for these electrons are in a random orientation compared to the original beam. The diffraction pattern for this process forms bands across the observation screen. Since the electrons must be able to penetrate into the bulk of the material for these features to form, observation of Kikuchi lines are an indicator of a highly crystalline film.

In the case where the electrons interact with a rough surface, the observed pattern can instead be caused by transmission through the crystal. These transmission patterns appear as spots which do not move as the substrate is rotated (as opposed to diffraction spots, which do move when the sample is

rotated). The difference between transmission patterns and diffraction patterns is another way that RHEED is used to evaluate the quality of the sample surface.

3.5 Growth of the Alumina Layer

Concurrently with the deposition of the aluminum buffer layer, I prepared the chamber for the injection of oxygen for the alumina layer that follows the growth of buffer layer. The aim of the alumina growth would be to introduce oxygen into the chamber quickly and seamlessly. This was done in part so that the calibrations from the buffer layer growth would hold true, and to prevent any defects or impurities from embedding themselves in the film. The insulating alumina layer was thought to be most susceptible to critical current noise, and cleanliness at this stage was of utmost importance. To that end, I will describe the oxygen gas lines and the cleaning procedures used to ensure clean gas was injected into the chamber.

A bottle of ultra-high purity (99.995%) oxygen was mounted near the chamber and connected by stainless steel piping to the gas injector. The stainless steel piping could be heated and purged to remove any possible water vapor from the lines. A pumping line was also connected to the gas injector, which was pumped by either a removable turbo-pumping station or, later, a dedicated roots pump. A series of on-off valves and pressure regulators were used to control the gas in the lines. Before oxygen was introduced into the chamber, the gas injector was pumped and purged with oxygen. The pressure in the lines was monitored using a thermocouple gauge and the purging pressure was several times the pressure present in the lines during growth.

Before introducing the oxygen into System E, the gate valve to the ion pump at the base of the chamber would be closed, leaving only the cryo pump acting on the system. This still left a pumping capacity of approximately 1000 L/s acting on the chamber. All films had a pressure measurement made by an ion gauge installed at the base of the chamber; some films also had a pressure measurement made by an ion gauge which was installed behind the substrate holder for this purpose. This second ion gauge was a better measure of the “beam” pressure on the substrate and numbers based on this gauge will be referred to as “beam pressures”. A pressure of 5×10^{-7} Torr at the base of the chamber corresponded to a pressure of 1.6×10^{-7} Torr at the beam gauge. All of the measurements made with only the base gauge have been scaled to the beam pressure for consistency.

As discussed previously, the aluminum buffer layer would be terminated by raising the substrate mask between the film and the aluminum source. The aluminum source would remain untouched while the

oxygen valves were opened. An on/off valve capable of maintaining pressure differences of an atmosphere to 10^{-11} Torr was opened to begin injecting oxygen into the chamber. The pressure was controlled using a regulator valve and the chamber could be brought up to the required growth pressure within 60 seconds. Once at the growth pressure, the mask was lowered and the co-deposition growth was begun. Using the growth rate found during the aluminum buffer process, the requisite number of monolayers were deposited and the growth was terminated by closing the oxygen valve and shuttering the aluminum cell.

After the alumina layer was deposited, the aluminum cell was cooled to its standby temperature and the chamber was evacuated by the cryo pump. Once the pressure had been brought below 10^{-8} Torr, the gate valve to the ion pump would be opened, so that both pumps would work on evacuating the chamber. After approximately 20 minutes, the chamber pressure would be back to 1×10^{-9} Torr and the substrate could be transferred to the tube for further oxidation.

3.6 Post-Alumina Diffusion Process

The oxidation of the insulating alumina layer was typically completed by diffusion outside the growth chamber. Films grown at the beginning of my research were oxidized in the loadlock using the following procedure. A “tee” on the oxygen manifold described above directed gas to the loadlock. The film would be loaded into the loadlock, which had a base pressure of 1×10^{-7} Torr. The pumps on the loadlock would be isolated, and a pressure of 10 Torr of oxygen would be introduced into the chamber. The loadlock could maintain a pressure of this magnitude for upwards of an hour, whereupon it would be evacuated. Once the loadlock had again achieved its base pressure and the gate valve to the transfer tube could be safely opened, the film would be returned to the growth chamber for the deposition of the top electrode.

For films grown at the end of my research, I used a dedicated oxidation chamber that had a lower base pressure and could be baked out and left under vacuum indefinitely. The base pressure of the oxidation chamber was 1×10^{-10} Torr and the chamber had its own dedicated cryo pump and oxygen gas bottle. To calibrate the pressure in the oxidation chamber with the loadlock, a bellows was connected between the two chambers and the pressure was monitored on both chamber simultaneously. Oxygen was injected into the loadlock (and hence into the oxidation chamber) using the procedure described previously; at 10 Torr in the loadlock, the pressure in the oxidation chamber was recorded as 32.92 Torr. Oxygen was then injected into the oxidation chamber (and hence into the loadlock) using the oxidation chamber oxygen source and the pressure in the oxidation chamber was 31.96 Torr with 10 Torr in the

loadlock. The oxygen was injected from each direction to ensure that there was no pressure gradient between the two chambers. Thus, approximately 30 Torr of oxygen in the oxidation chamber is equivalent to the 10 Torr of oxygen in the loadlock. A majority of samples were grown using the loadlock procedure, so I have converted any oxidation chamber diffusions to loadlock equivalent pressures for clarity.

3.7 Growth of the Niobium Top Electrode

After the alumina layer was deposited, the film had lost its single-crystal structure. RHEED of the alumina surface shows rings, a signature of an amorphous surface. Attempts were made to deposit metallic aluminum between the aluminum oxide and top niobium, forming a symmetric film structure. Each attempt was thwarted by an unetchable eutectic between the aluminum and niobium films; it is likely the heat from the e-beam evaporator source which resulting in this layer. As a result, there was no possibility of growing single crystal niobium on the aluminum oxide surface.

Instead, polycrystalline niobium was deposited directly on the amorphous aluminum oxide. While this may result in increased noise from localized MIGS, heating up the sample to grow single crystal niobium posed the real danger of drawing oxygen out of the insulating alumina. On balance, noise from the disordered interface is likely less than the noise would be if oxygen was reduced from the insulator. After the film was returned to the growth chamber, the e-beam evaporator was powered on, brought up to the growth rate of $\approx 0.3\ \text{Å}/s$ and the required thickness (600Å - 1000Å) was deposited.

3.8 Summary

This chapter has detailed the steps used to grow the standard niobium/aluminum/aluminum-oxide/niobium “trilayer” that would then be fabricated into the desired electronic device (SQUID, transmon, Josephson junction, etc.). The use of MBE to control the growth allowed me to grow epitaxial niobium and to precisely control the oxygen dose for the aluminum oxide barrier. The following chapter will describe measurements made of just the films themselves and Chapter 5 will describe the techniques used to transform these films into testable device structures.

Chapter 4 Materials Analysis

This chapter will describe measurements made on thin niobium films grown as described in Section 3.3. I begin by discussing transport measurements of these films and then move on to discuss the microscopy of these samples, showing the ideal growth temperature to be between 750 – 850°C. I review the literature to better understand the epitaxy of niobium on A-plane sapphire, then describe my AFM and RHEED measurements of the surface structure, which indicate a slightly lower growth range (650 – 750°C). My final section will discuss transmission electron microscopy of both niobium and aluminum films, demonstrating that I have an epitaxial interface between the sapphire and niobium metals. I then use EELS to show that oxygen is being depleted from the sapphire substrate for temperatures greater than approximately 610°C. Combining the different measurements shows that the ideal growth temperature is in the range of 640 – 650°C.

4.1 Transport Measurements of Niobium Thin Films

Transport of niobium films can be used to reveal information about crystal structure and defect density in the metal sample. In particular, the temperature dependence of the resistivity indicates whether there are other sources of scattering in the metal beyond electron-phonon interactions in a perfect lattice [31]. The residual resistivity ratio ($RRR = \rho_{300}/\rho_{10}$) between the resistivity at 300 Kelvin and the resistivity at 10 Kelvin provides a measure of the temperature dependence of the films. Higher RRR values are found for samples that are largely free of impurities and crystal defects.

RRR measurements for superconductors are typically made at 300 Kelvin and just above the superconducting transition temperature, but many of my temperature measurements began below 295

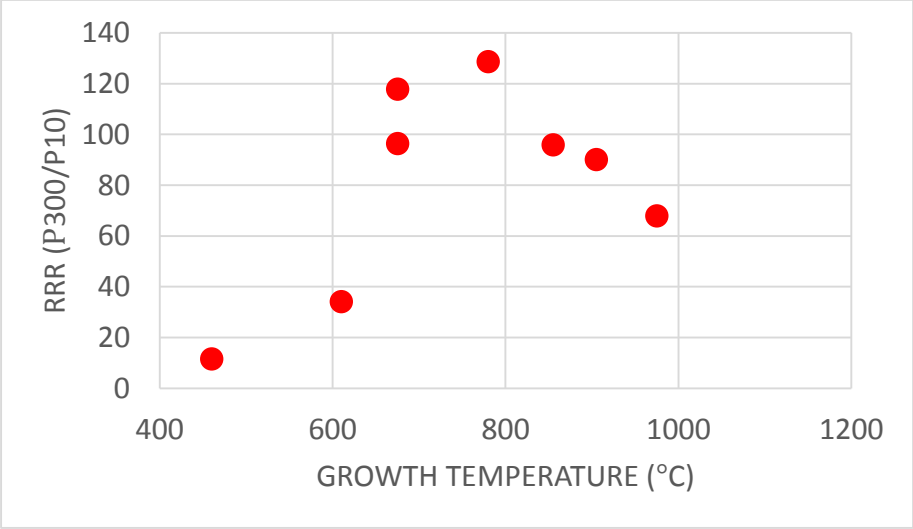

Figure 16 Residual resistivity ratio for 3000Å thick niobium samples.

Kelvin. The temperature dependence of resistivity is linear at room temperature, so all the resistivity measurements have been scaled to 300 Kelvin. The average shift was 2% and the largest adjustment scaled RRR up by 9%.

RRR for my films span the range of 10-130, with the maximum RRR occurring for a film grown at 780°C. These values agree with reported RRR measurements from other groups [32, 33]. The increase in RRR for samples at low-to-moderate temperatures (below 800°C) is likely due to improvements in the crystal structure of the film leading to longer mean-free-paths in the metal. The decrease in RRR for samples grown at high temperatures (above 800°C) is somewhat unexpected. Improvements in the crystal structure/surface morphology with higher growth temperatures have been reported by several groups [34, 35], which should lead to even higher values of RRR. The most likely explanation is that the films are being grown with increased oxygen in the bulk. A previous study by De Sorbo [36] showed that the resistivity of niobium wires increased as oxygen was dissolved into the metal. Additionally, RHEED and LEEM measurements of niobium films grown or annealed at high temperatures indicate the presence of oxygen at the sample surface. I will return to the topic of oxygen in the niobium film in Section 4.4.3.

In addition to measuring the RRR of my niobium films, I also measured the superconducting transition temperature (T_c) for each sample. The general increase of T_c seems to correlate with improvements in the crystal structure. There is a small decrease in T_c above 900°C for the 3000Å thick samples (difficult to see on the scale in Figure 17, which again may correspond to oxygen adsorption into the film.

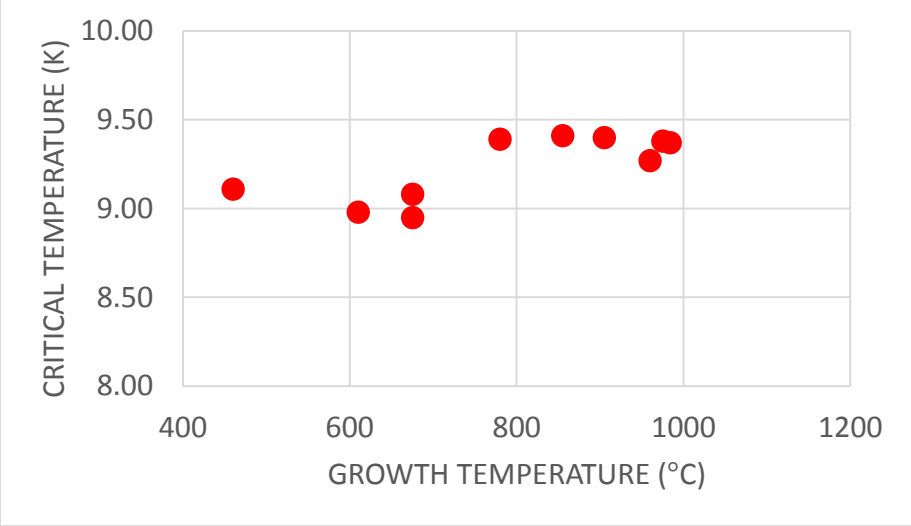

Figure 17 Superconducting transition temperature for 3000Å thick niobium films. A maximum transition temperature of 9.41K occurs for the sample grown at 855°C.

The highest RRR value occurs at 780°C and the highest T_c occurs for the sample grown at 855°C. Samples grown hotter than this appear to have too much mobility that leads to additional defects in the crystal lattice. Niobium grown at a substrate temperature of less than 700°C also has lower RRR and T_c measurements; based solely on these measurements the ideal range appears to be 750-850°C. I will now move on and discuss the crystal structure of the niobium itself, beginning with a discussion of the crystal structure of the underlying substrate.

4.2 Thin Films of Single Crystal Niobium on A-Plane Sapphire

The heteroepitaxy of niobium grown on sapphire has been well studied [28], and in this section I will review the growth of niobium on sapphire so that my own results are placed in a broader context.

In this work, niobium was grown exclusively on A-plane sapphire, and it is therefore useful to briefly discuss the nature of this crystal structure. Aluminum oxide (Al_2O_3) has a relatively complex crystal structure, with planes of oxygen atoms interlaced with planes of aluminum atoms. The oxygen anions are in a close-packed structure and form hexagonal close-packed layers ($ABAB$ …). The aluminum cations sit on two-thirds of the octahedral interstitial sites and the vacancies vary from plane to plane in the unit cell. The full structure for the cell is ($AC_1BC_2AC_3BC_1AC_2BC_3$ …) where C_i are the possible

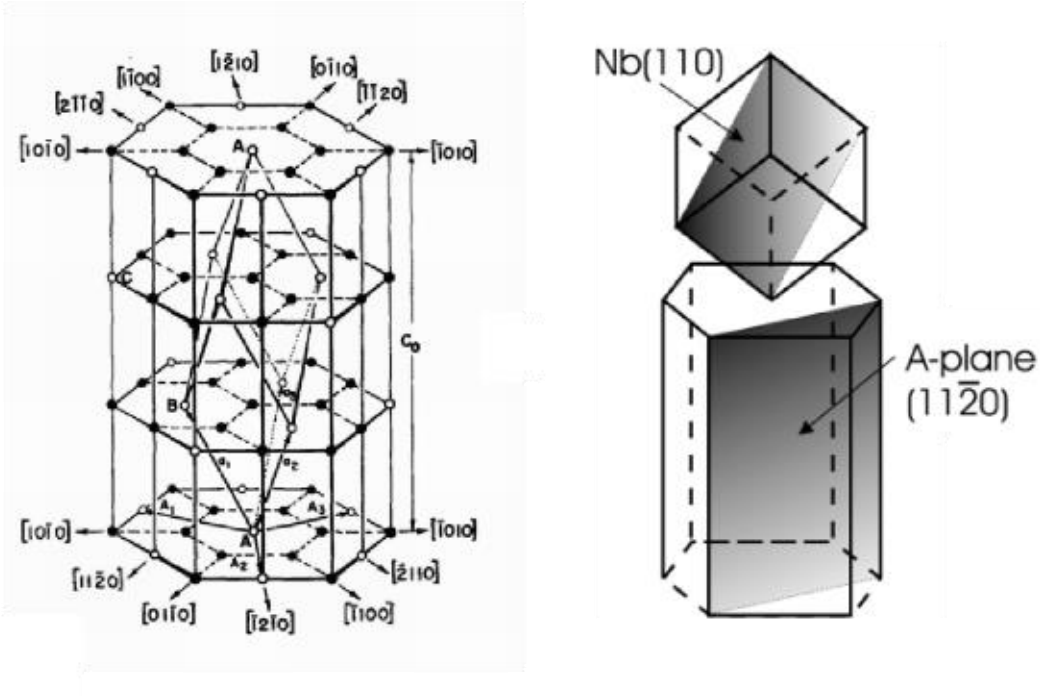

Figure 18 (Left) Morphological unit cell for sapphire showing aluminum anions. Empty circles are aluminum vacancies. Oxygen atoms are not shown. The structural unit cell of sapphire is twice the height along the c-axis. Figure from Reference [84]. (Right) Three-dimensional relationship between A-plane sapphire and epitaxial niobium. Figure from Reference [28].

aluminum vacancy orientations and A and B are the two oxide hexagonal close-packed planes. The lattice constants for this -six-layer hexagonal crystal structure are $a = 4.7592$ Å and $c = 12.9916$ Å at room temperature [37]. The rhombohedral unit cell has constants $\alpha = 85.7°$ and $a = 3.5$ Å [38, 28].

Niobium, which has a body-centered cubic lattice structure with lattice constant $a = 3.3$ Å, has a three-dimensional relationship with sapphire. For A-plane $(11\bar{2}0)$ sapphire, niobium grows in a (110) orientation [39]. Where the interface is coherent, i.e. when the niobium lattice is strained to maintain epitaxy, strain is estimated to be approximately 2% as along the $[\bar{1}1\bar{1}]$ axis [38, 40].

There have been a number of studies investigating epitaxial niobium grown on A-plane sapphire at temperatures between 700°C and 1000°C [39, 34, 32]. There has also been work studying the growth of niobium on sapphire at temperatures between 350°C and 530°C [40]. My work builds on previous films grown by Paul Welander using this MBE system [35] but extends substrate growth temperatures down to 550°C. I will now discuss the surface studies of these niobium films.

4.3 Reflection High-Energy Electron Diffraction and Atomic Force Microscopy Measurements

As discussed in Section 2.2, MIGS are thought to be a potential source of noise in flux qubits. As MIGS arise due to interfacial defects at the substrate/metal boundary, it was necessary to quantify the crystallinity of my niobium films. A non-invasive/non-destructive way of doing this was to measure the surface morphology using AFM and the surface crystal structure using RHEED. Both measurements are proxies for the crystal structure of interest, namely the crystal structure at the sapphire/niobium interface.

The basic physics of RHEED are discussed in Section 3.4.1, so this section will first discuss the technique of AFM to measure surface morphology. The results for samples that were grown at varying temperatures will then be discussed.

AFM is a type of scanning probe microscopy which produces high-resolution topographic maps of a surface. A cantilever with a very sharp tip is lowered to the surface under study. As the tip is brought into proximity with the surface, various forces repel the cantilever tip. A laser reflects off the back of the cantilever and is measured using a photodiode array. As the tip is moved along the sample, the deflection of the laser spot maps to changes in the surface topography. To ensure that the tip does not collide into the surface, a feedback mechanism is used to maintain a constant force between the tip and

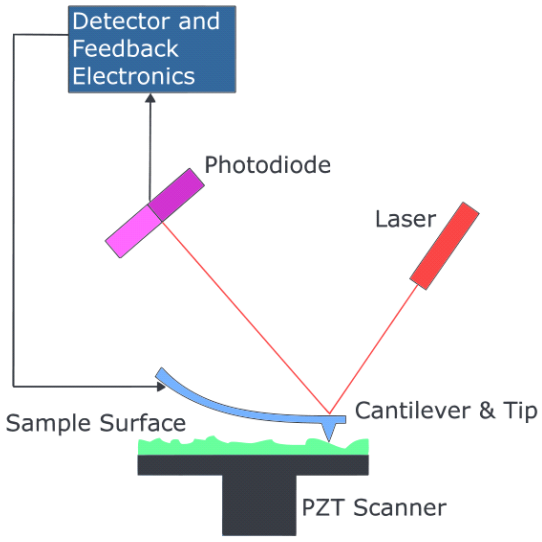

Figure 19 Schematic of an AFM. Figure from Reference [70].

the surface; piezoelectric motors position the sample in three dimensions. The Asylum Research MFP-3D microscopes at the Center for Microanalysis of Materials can achieve sub-nanometer resolution in the vertical dimension and sub-micron resolution in the lateral dimension.

The microscopy in this work was done using the tapping mode of cantilever control. When in tapping mode, the cantilever is oscillated near its resonance frequency and the tip makes intermittent contact with the surface. Furthermore, tapping mode helps to preserve the tip and allow for higher resolution imaging.

The samples that I will discuss here were grown at temperatures between 460°C and 984°C and were 1000-3000Å thick. RHEED measurements were made *in-situ* immediately following growth, while AFM measurements were made shortly after removing the sample from vacuum.

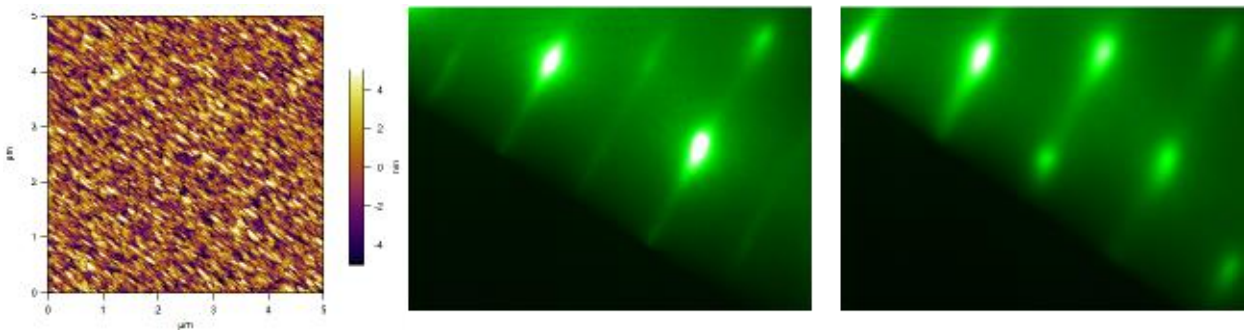

Figure 20 (Left) AFM image of a 3000Å thick sample grown at 610°C that shows a rough morphology. (Center) RHEED image along one crystal axis of this sample shows narrow lines. (Right) RHEED image of the same sample rotated 90° shows diffuse lines and 3D spots, indicating a rough surface.

AFM images of these samples generally show one of two distinct phases. The first phase has a rough surface morphology with an RMS roughness of ~20-30Å as in Figure 20 and were grown between 460°C and 620°C. The niobium surface has small grains of material that are aligned parallel to each other. This agrees with RHEED measurements which show well-defined, narrow patterns along one crystal axis, but diffuse spots when the sample is rotated 90° to the narrow axis. Narrow reflection lines are a qualitative measurement of surface roughness, and diffuse spots can also indicate surface roughness. The broad diffraction lines along one crystal axis indicate there is considerable surface roughness in that direction. However, the samples shown here do not have true 3D transmission spots, as the diffraction pattern of streaks and spots varies as the sample is rotated through the RHEED beam.

The second growth phase that I observed occurred for films that were grown at temperatures between 630°C and 750°C. For these samples, the RMS roughness was only about 2-3Å, and the surface morphology was noticeably different. For these samples, the niobium film had flattened and formed curvy "entrails-like" terraces with 2Å step-edges as shown in Figure 21. The RHEED patterns also show substantially different characteristics. The diffraction patterns have very narrow streaks, with thin

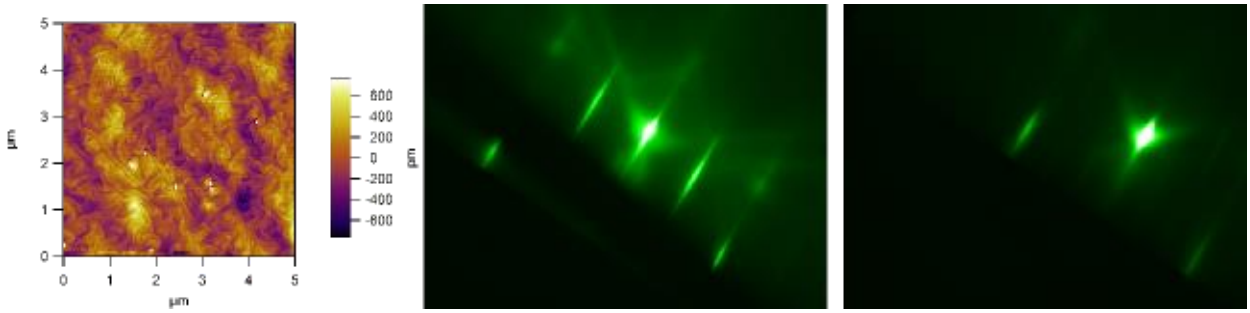

Figure 21 (Left) AFM image of a 3000Å thick sample grown at 675°C showing the "entrails-like" terrace discussed in the text. (Center) RHEED pattern along the $[001]$ *axis showing Kikuchi lines and narrow streaks. The dark region across the bottom portion of the image is a shadow from the sample mounting bracket. (Right) RHEED pattern along the* $[1\bar{1}1]$ *axis also shows Kikuchi lines and narrow streaks.*

specular spots, and Kikuchi lines are visible along all three possible crystal axes (these are $[\bar{1}11]$, $[1\bar{1}1]$ and $[001]$ for (110) niobium).

There were some samples that were grown in the 630°C - 750°C temperature range that were not quite as smooth and flat as the example shown above. These wafers typically had RMS roughness values between 5-10Å, but did not show the entrails-like terrace patterns seen in the flattest samples. The entrails pattern was achievable only under the best conditions – clean chamber, niobium source material in excellent condition – but this was the morphology that was used as the standard for all films. Samples that did not have this morphology but which were grown in the correct temperature range,

were suspected to have some source of contamination embedded in the niobium or to have been improperly prepared before the growth.

The lowest RMS roughness that I measured was 0.214 nm for a film grown at 675°C, but the next best film had an RMS roughness of 0.229 nm for a film grown at 750°C. As shown in Figure 22, the roughness then increases to a relative maximum at 855°C before falling again as the growth temperature increases. In the 850-900°C range, the niobium had faceted and did not have enough energy to reorder to a smoother structure. At higher temperatures, the lattice could relax from these facets to monatomic step-edges [35]. The smoothest films approximately correspond to the maximum RRR and T_c measured in Section 4.1, providing additional evidence that the crystal structure has few defects.

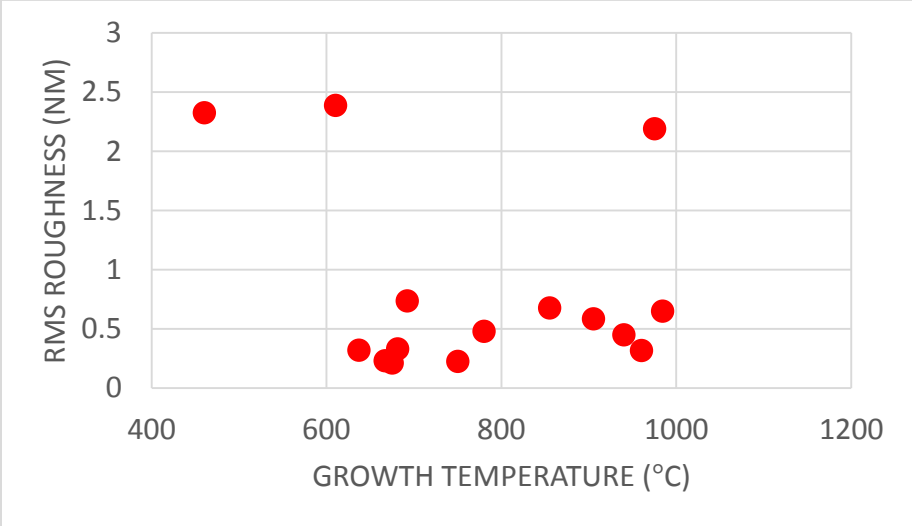

Figure 22 Root-mean squared surface roughness measurements of AFM scans of niobium films. The two distinct roughness regimes are clearly distinguishable here. A slight roughening is also seen for the samples grown above 800°C, but this is reduced again once the sample growth temperature is brought above 900°C.

The surface morphology shown in Figure 21 should be compared to the surface morphology measured by Zhou and Flynn [34], an example of which is shown in Figure 23. Zhou and Flynn also grew niobium on A-plane sapphire, but the "fingered" morphology that they observed is distinctly different from the terrace structure that I see. At 700°C, they observe a rough surface, while films at higher temperatures are either smooth or have wide faceted steps. The most likely cause of the difference in surface morphology is the miscut of the underlying sapphire substrate. Zhou and Flynn reported using substrates with a miscut of either .72° or .66° and the AFM features did change with the different miscuts. The substrates that I used were

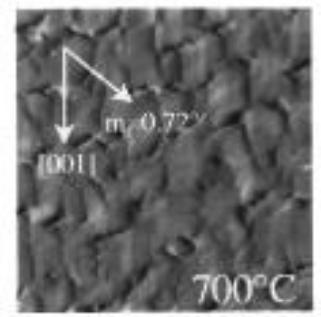

Figure 23 AFM of 1000Å niobium film grown at 700°C. Image scale is 2.1 µm. Image reproduced from [34].

obtained from Crystal GMBH and had a miscut of .1°, which may have allowed me to grow smoother films.

Both the RHEED diffraction patterns and the AFM roughness measurements indicate smooth growth between 650 – 750°C, a slightly lower temperature range than what was seen from the transport measurements. For growths around 600°C, I see a very rough surface morphology; TEM measurements shown in 4.4.2 demonstrate that this roughness is due to crystal defects deep within the niobium film. The correlation between the TEM measurements and AFM/RHEED allowed me to use the surface measurements as a proxy for the quality of the crystal interface and to then restrict device fabrication to those films with narrow streaks in RHEED and smooth surface morphology.

4.4 Transmission Electron Microscopy Measurements

While RHEED and AFM are extremely useful surface measurements, and transport reveals information about scattering sites for electrons, none of these techniques can actually reveal the crystal structure at the substrate/metal interface. Transmission Electron Microscopy (TEM) images the individual atoms in the structure and can directly measure crystal defects. Unfortunately, TEM is also a destructive technique, requiring a cross-sectional cut of the sample; for this reason my TEM measurements are limited to four samples.

There are several related techniques which all fall under the umbrella of TEM that will be described briefly here. This work makes use of High Resolution Transmission Electron Microscopy (HRTEM), Scanning Transmission Electron Microscopy (STEM), and Electron Energy Loss Spectroscopy (EELS). The instrument used to make these measurements was a JEOL 2200FS microscope located in the Materials Research Laboratory's Center for Microanalysis of Materials. I will begin by discussing the physics of TEM and EELS, then show STEM images of the niobium grown epitaxially on sapphire for samples grown at 610°C and above. I then show EELS measurements of these samples and calculate how much oxygen is depleted from the substrate. I finish with EELS measurements of oxygen content in the alumina barrier, demonstrating that oxygen is restricted only to the amorphous alumina barrier.

4.4.1 TEM Techniques

TEM uses energetic electrons (200 kV) to probe the internal structure of films. The wavelength of electrons accelerated to this voltage is nm which enables atomic resolution of the film under study. The electrons are focused by electromagnetic lenses before passing through the sample to create an

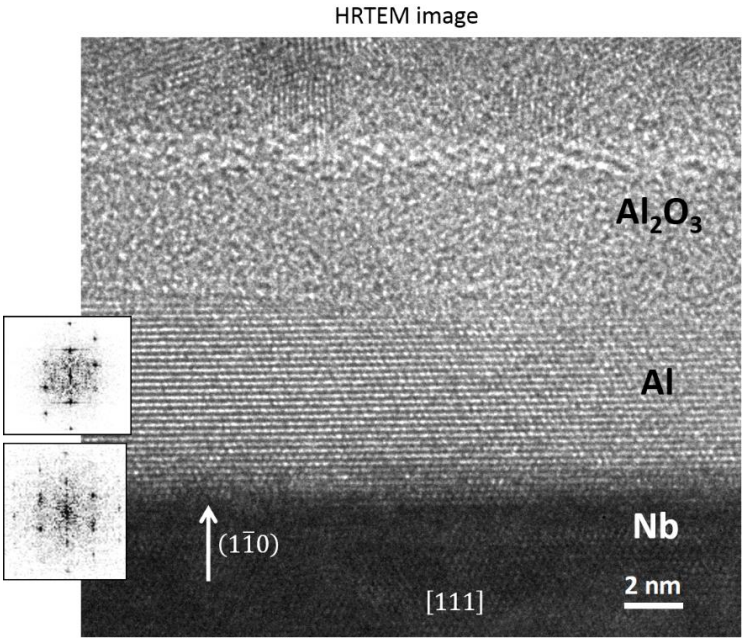

Figure 24 High-resolution TEM image showing the top of the base electrode, the metallic aluminum buffer layer, and the amorphous alumina barrier. The insets show the transmission patterns for both the niobium and aluminum regions which can be used to determine the crystal structure.

observed pattern. The resulting image can be observed using varying techniques that reveal different information about the sample.

HRTEM makes use of the crystalline structure of the sample to create a diffraction pattern. As the imaging electrons pass through the sample they interact with the electrons in the crystal. The resulting diffraction pattern can be analyzed to determine the crystal structure of the material and to measure the lattice constant of the grown films.

In contrast with HRTEM, STEM uses a rastering beam of electrons to generate an image. Electromagnetic lenses focus the beam to give atomic scale resolution. Several different types of detectors can be used (some simultaneously) to analyze the sample. A bright field (BF) detector measures electrons that pass through the sample at low angles. Surrounding the bright field detector is an annular dark field (ADF) detector that detects electrons that scatter at high angles relative to the incident beam. This latter type of imaging is also known as Z-contrast imaging, since the scattered electron intensity has an atomic number dependence.

Together with Z-contrast imaging, the energy of the transmitted electrons can be measured to provide chemical identification of the atoms under study. While the high-angle scattered electrons are collected by the ADF detector, the low-angle scattered electrons can be analyzed using a spectrometer. This technique, known as electron energy loss spectrometry (EELS), makes use of the fact that the

transmitted electrons have inelastically scattered with the sample. While there are many ways for electrons to scatter with the sample, chemical identification can be done by looking at specific (and known) inner-shell ionizations. Thus it is possible to spatially resolve different elements within the sample due to the presence of their energy spectra.

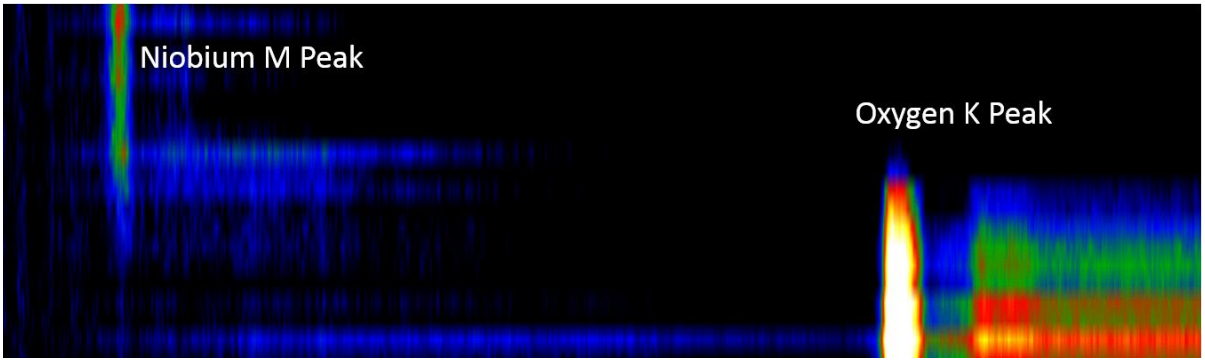

Figure 25 EELS scan of the niobium sapphire interface. The scan direction is along the y-axis starting at the top of the niobium film and ending at the bottom in the sapphire substrate. The x-axis is energy loss allowing separate ionization peaks to be characterized.

4.4.2 STEM Analysis of the Niobium-Sapphire Substrate Interface

Following the motivation provided in Section 2.2, I set out to examine the crystal structure of my niobium films at the interface between the niobium and the sapphire substrate. Although I have RHEED images of every film that I have grown, because RHEED measures reflected electrons it is solely a surface technique, and cannot reveal structures deep within the film. Furthermore, because of the amount of light emitted by the electron-gun during deposition, it is not possible to use RHEED to monitor the crystal structure during the niobium growth.

To measure the crystal structure, the samples were first diced and then polished so that the electron beam could be transmitted through them. As the dicing process was destructive, this was done for only four niobium films which were grown at room temperature, 610°C, 675°C, and 750°C, respectively. The

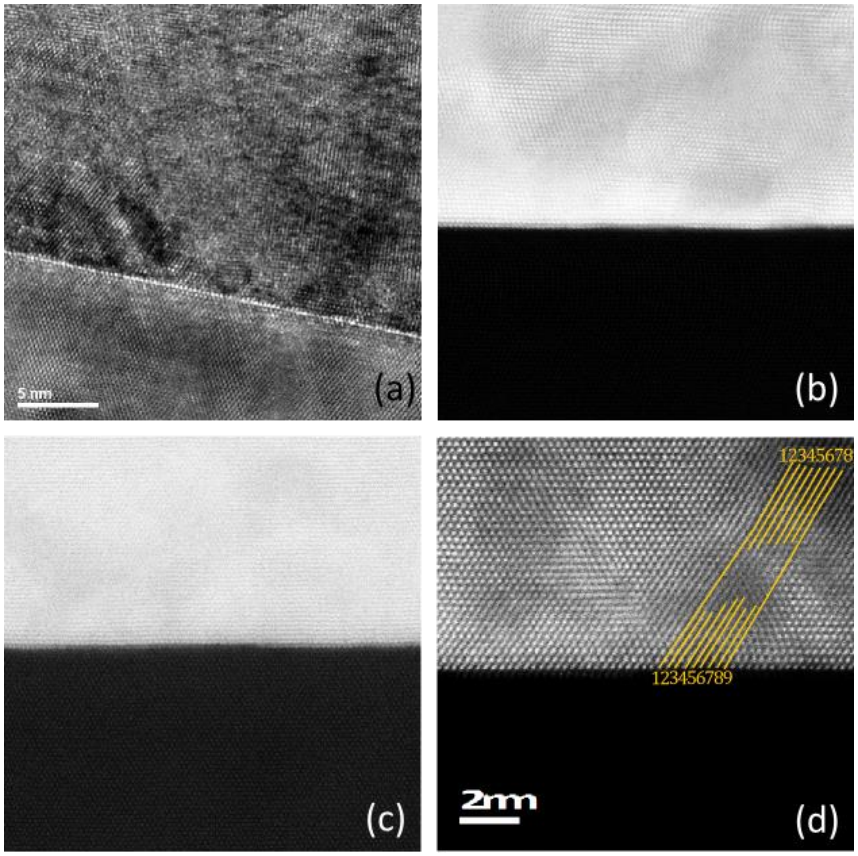

Figure 26 HR-TEM and STEM images of niobium grown on sapphire. In each image, the niobium film is at the top of the image. (a) HR-TEM image of the sample grown at room temperature. The material inside the circled areas is epitaxial, although the film is generally polycrystalline. (b) STEM image of the sample grown at 610°C. The curvature at the top of the image is an instrumentation artifact. (c) STEM image of the sample grown at 675°C. (d) STEM image of the sample grown at 750°C. The yellow lines help demonstrate a lattice relaxation point. There are nine lines at the base of the relaxation point, whereas only 8 lines emerge from the point. The relaxation point is approximately 10 monolayers away from the substrate.

film grown at room temperature is highly disordered and polycrystalline. In contrast, the film grown at 750°C shows perfect epitaxy at the substrate-metal interface. Careful inspection of the STEM image reveals that there is relaxation of the crystal lattice, but this occurs ~10 monolayers into the metal film.

A different magnification was used for the 610°C and 675°C samples. The crystal structure of these films resembles the structure of the 750°C (single crystal) sample, as opposed to showing grains like the room temperature film. There is no evidence of lattice relaxation in either of these two images. There is some bend/shift in the image, but this is an artificial shift of the electron beam, and does not reflect the true crystal structure of the film. The bending of the electron beam is likely due to charging effects because of the insulating nature of the sapphire substrate – similar effects are seen when imaging using scanning electron microscopy (SEM). While it is possible that lattice relaxation may be temperature-dependent, the small-scan size of my STEM images makes broad conclusions difficult.

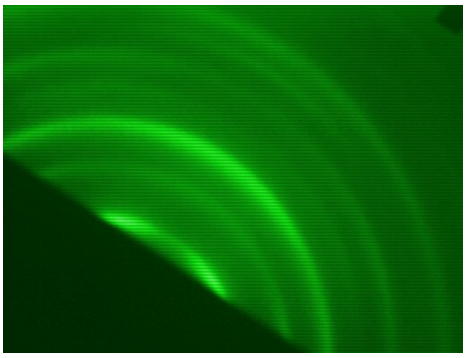

Figure 27 RHEED pattern showing the rings typical for polycrystalline films.

When the STEM scans are compared with each other, it is evident that single-crystal niobium can grow on A-plane sapphire at temperatures as low as 610°C. Measurements taken with RHEED support this conclusion; although the RHEED pattern for the 610°C has diffuse streaks, the diffraction pattern does not show the ring-pattern that is characteristic of polycrystalline films. Furthermore, RHEED indicates that single-crystal niobium-metal interface can be grown even on a 460°C substrate.

The STEM scans should be compared with the AFM images and RMS roughness measurements described in Section 4.3. In AFM there is an abrupt transition to rough surfaces for samples grown below 630°C, whereas STEM shows a single-crystal sapphire/niobium interface even at 610°C. Large-scale STEM images also show this behavior pattern. In Figure 28, defects are observed 200Å into the growth of the niobium film for the 610°C sample, whereas the 675°C film appears to be single-crystal throughout the entire film thickness (both films are 3000Å thick). The roughness measured in the AFM

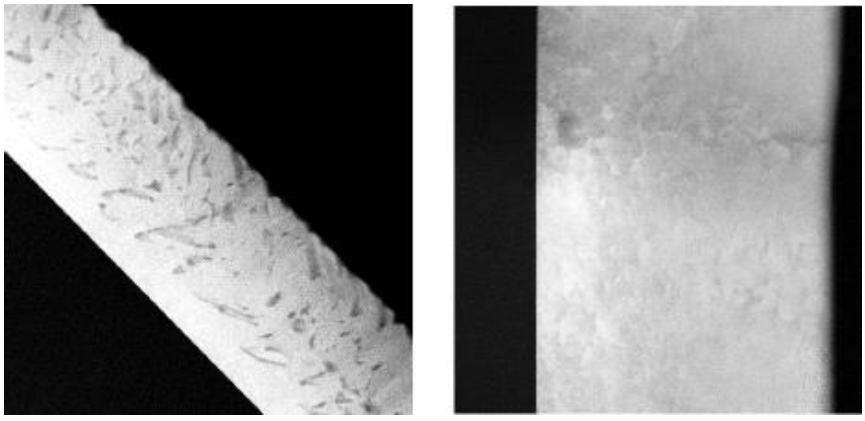

Figure 28 Low magnification TEM images of the 610°C and 675°C samples. The entire 3000Å thick films are shown. (Left) TEM of the 610°C sample, showing defects beginning at approximately 200 Å. (Right) TEM of the 675°C sample, showing a single crystal film throughout the entire thickness.

films is due to the defects as the thickness increases. However, since the RHEED pattern still shows diffraction patterns for single-crystal films (albeit it a rough surface), I conclude that the film is still mostly single-crystal with crystal defects or small areas of polycrystalline niobium at the surface.

The STEM measurements show that the AFM measurements are not perfectly correlated with the crystal structure at the base of the electrode. Based on AFM alone, it seemed unlikely that the 610°C

sample had any epitaxy at the sapphire/niobium interface. Not only does the STEM reveal why the niobium surface was so much rougher at a growth temperature of 610°C, but it also shows a single-crystal film for the first 200Å of the growth, revising the conclusion from the AFM micrographs.

STEM measurements show that niobium can be grown epitaxially on sapphire down to a temperature of at least 610°C. The lack of defects at the interfacial boundary should reduce flux noise in superconducting devices if the measured flux noise is due to MIGS. However, low growth temperatures do produce crystal defects away from the substrate, resulting in a rough surface morphology that may be problematic when trying to fabricate Josephson junction devices. The temperature range of 650-800°C, which has already shown good results from RRR, T_c, and AFM measurements, also shows good epitaxy throughout the film in TEM measurements.

4.4.3 EELS Analysis of the Niobium-Sapphire Substrate Interface

In addition to the crystal structure analysis at the substrate, the chemical make-up of the substrate and niobium film was also measured using EELS. These measurements were made simultaneously with the crystal structure measurements; the same four samples discussed in Section 4.4.2 were used here.

The motivation for this chemical analysis comes from the observation that film quality decreases with increased growth temperature (see RRR and T_c measurements in Section 4.1). Work done by Welander

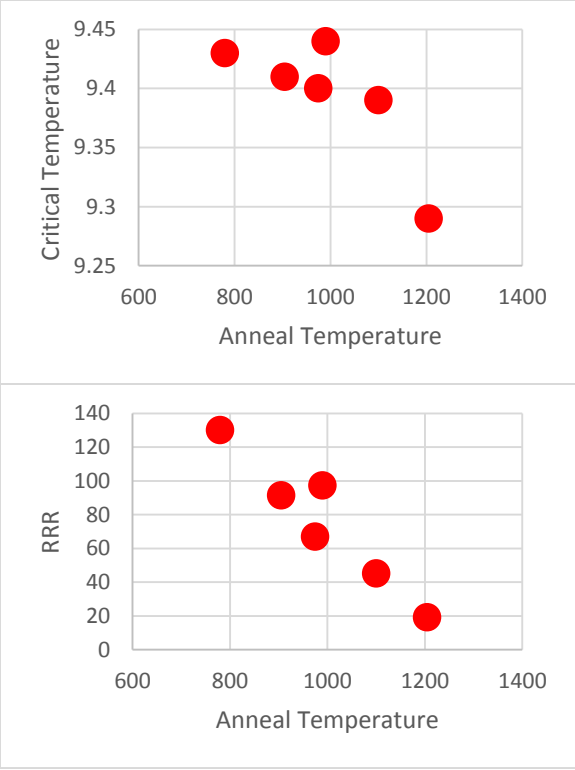

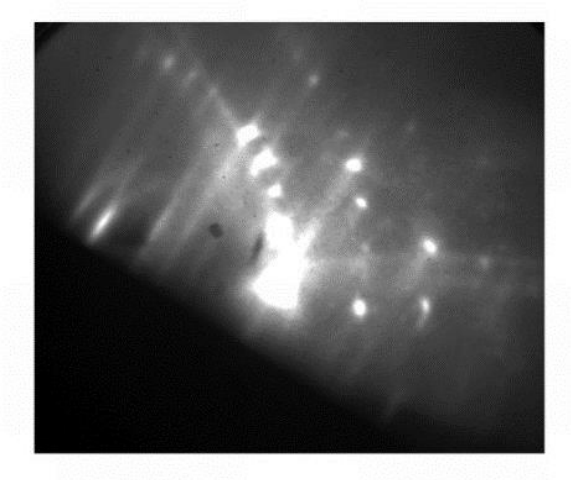

Figure 29 Data for niobium films which were annealed after the growth. (Upper left) Transition temperatures for the annealed films showing a decrease in T_c with increased anneal temperature. (Lower left) RRR measurements also show a decrease with the anneal temperature. (Above) RHEED image of a sample annealed at 1300-1370°C showing the RHEED reconstruction pattern discussed in the text.

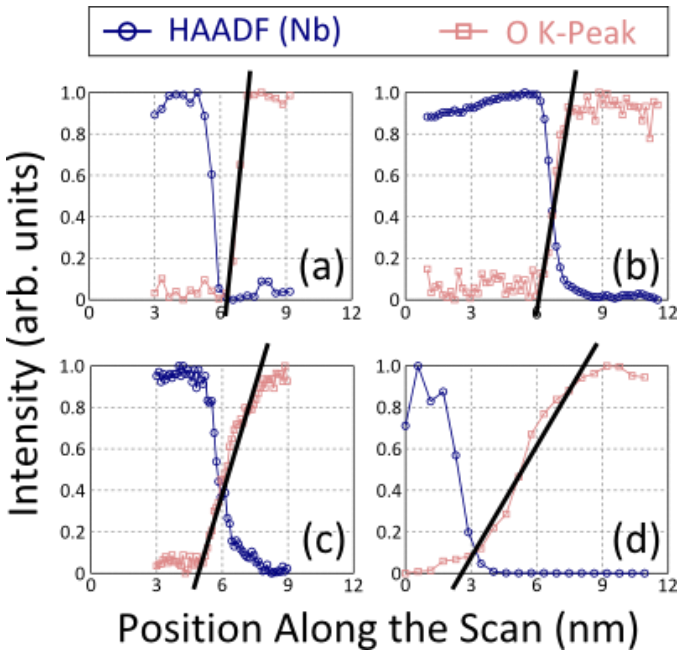

Figure 30 EELS scans of niobium films showing depletion of oxygen from the sapphire substrate. Blue is the High-Angle Annular Dark Field detector intensity and represents the presence of niobium. Pink is the oxygen K-peak signal. Black lines show the slope of the oxygen depletion. (a) Room temperature sample. (b) 610°C sample. (c) 675°C sample. (d) 750°C sample.

[35] measured properties of niobium films grown between 750-950°C and annealed at temperatures up to 1200°C and focused on the effect of annealing temperature on film properties. His films show a deterioration in film quality with an increase in the annealing temperature (Figure 29). He attributes the decrease in both T_c and RRR to oxygen adsorption from the substrate. This interpretation is supported by RHEED images of samples annealed at 1300-1370°C. After annealing, the diffraction pattern had 1/4-order streaks with spot rings whereas before annealing the diffraction pattern only had 1/3-order streaks. Other groups used Auger electron spectroscopy of the niobium surface grown at high temperatures to concur that oxygen is adsorbed to the surface [41, 42, 43]. While the presence of niobium oxide at the surface of annealed samples is well-established, the origin of this contamination has been unknown. My work, studying the chemical structure at the substrate-metal interface, suggests that adsorption from the substrate and diffusion through the film is the cause of the oxide at the surface.

Measurement of oxygen adsorption from the sapphire substrate into the niobium film was done by EELS. Oxygen was identified by its K-peak at 556 eV, while the niobium/substrate interface was determined using the z-contrast data. The z-contrast data measures the brightness of the image being measured by the dark field detector and is primarily sensitive to electrons which have interacted with heavy atoms like niobium. The niobium M-peak at 385 eV was not used because not enough niobium

atoms were deflected into the EELS detector to make a reliable comparison across samples. Because the niobium and oxygen peaks were not measured using the same detector, the measurement of oxygen adsorption is made by only measuring the change in the oxygen K-peak signal.

For all the measured samples, the niobium signal decays sharply, but for some samples there is an oxygen signal present at low levels in the niobium metal. The measurement of the presence of oxygen in the niobium is the oxygen depletion depth and is the distance over which oxygen is present above the noise floor to the point of maximum saturation. The values measured are shown in Table 1.

SAMPLE GROWTH TEMPERATURE	OXYGEN DEPLETION DEPTH (NM)
ROOM TEMPERATURE	1.0
610°C	1.1
675°C	2.6
750°C	3.5

Table 1 Oxygen depletion depths measured from the noise floor to the point of maximum saturation.

Another way to compare the quantity of oxygen depletion from the substrate is to examine the slope of the oxygen signal across samples. In Figure 30, the oxygen K-peak signal form EELS is shown with the same x-axis for all four samples. The decrease in the slope with growth temperature confirms that additional oxygen is being adsorbed from the samples grown at higher temperatures.

The measurements of the oxygen depletion depth for the samples grown at room temperature and at 610°C show approximately the same depletion depths although the samples were grown at vastly different temperatures. Assuming there is no diffusion of oxygen out of the substrate at room temperature, then this number – 1.0 nm – should be interpreted as the sensitivity limit of the spectrometer. Taking the oxygen depletion depth at 610°C to be at the sensitivity limit, I can find an approximate value for the activation energy of oxygen adsorption from the sapphire substrate. The slope of the fit provides the activation energy, multiplied by the Boltzmann constant, so the activation energy for oxygen to be adsorbed from sapphire is 0.65 eV.

Using this activation energy, I can estimate the point at which adsorption from the sapphire is frozen. As discussed in 4.2, sapphire has a lattice constant $c = 12.9916\text{Å}$. There are 6 oxygen-aluminum planes in

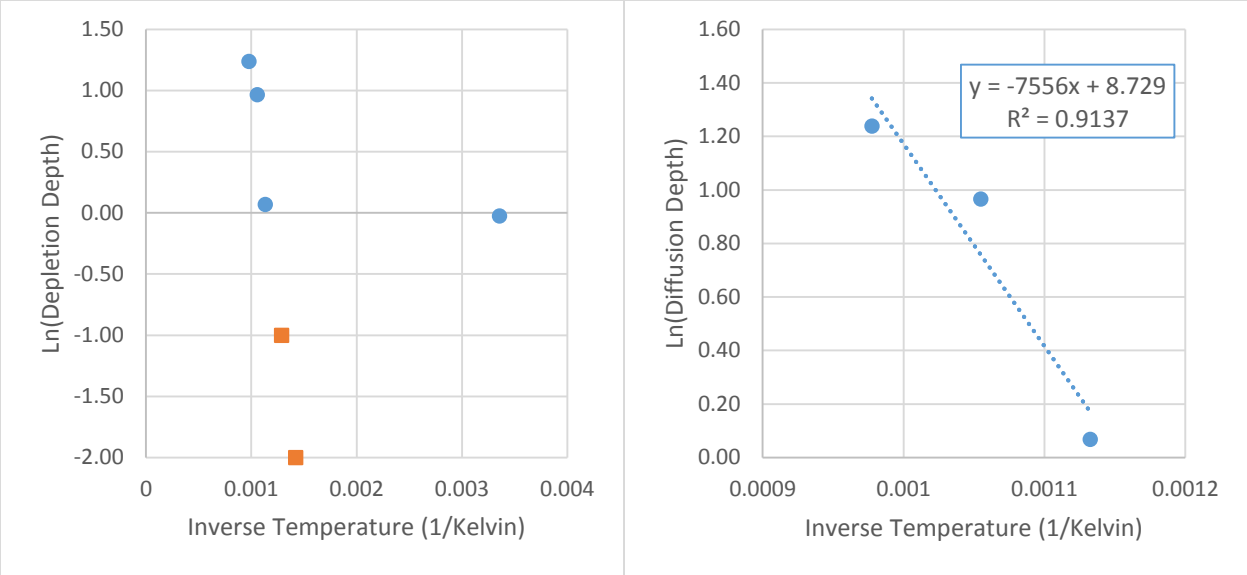

Figure 31 Plot of oxygen depletion depths listed in Table 1. The diffusion process is governed by the Arrhenius equation $d = Ae^{-E_A/k_BT}$ so to find the activation energy, I have plotted inverse temperature and the natural logarithm of the diffusion depth. (Left) The blue dots show all four data points taken from Table 1. The measurement saturation point is clear here. The orange squares show two calculated diffusion depths based on the slope calculated at the right. (Right) Linear fit for the three data points above the saturation point. The slope here is used to estimate the limit of diffusion.

this vertical structure, so any single plane is 2.16 Å tall. Using this as an upper limit on when oxygen can no longer depleted from the sapphire, I find that this corresponds to a sample growth temperature of 465°C. As a lower limit, I use the size of an oxygen atom in sapphire: 1.40 Å [44]. This height corresponds to a growth temperature of 430°C. Based on these limits, I estimate that adsorption no longer occurs for growths below 430°C.

Whereas the RRR, T_c, AFM roughness, and STEM measurements all indicated an ideal niobium growth temperature in the range of 650 – 750°C, the EELS measurements reveal that oxygen is depleted from the substrate in this range. While the epitaxial growth of niobium on sapphire may reduce flux noise due to MIGS, it leaves behind oxygen vacancies in the aluminum oxide substrate that can, in turn, produce their own noise. The stoichiometric imbalance will leave charge traps in the substrate which become electron traps. As discussed in Section 2.2, these traps can flux-lock an electron, generating flux noise due to the stoichiometric imbalance rather than any lattice mismatch.

The ideal growth temperature therefore is less than 650°C. STEM shows good epitaxy for the sample grown at 610°C, so flux noise due to MIGS should be alleviated at this lower temperature. RRR and T_c are not maximal, but are both sufficiently high to allow for useful measurements. The surface roughness is the one problem with growing this low – a 2nm rough surface may cause pinhole defects in a Josephson junction structure, effectively shorting the device. However, a smooth film was observed at

637°C; this growth temperature might be low enough to have a low density of oxygen vacancies and subsequently low flux noise. Moreover, the usual practice of coating the niobium base electrode with an overlayer of aluminum is known to provide smoothing (i.e. as in Figure 15).

4.4.4 EELS Analysis of the Aluminum Oxide Barrier

In this last section, I move away from the niobium/sapphire substrate to analyze the chemical structure of the aluminum oxide layer grown as the Josephson junction insulating barrier. Just as the chemistry of the niobium-sapphire substrate can be interrogated with EELS, the chemistry of the Josephson junction barrier can also be studied with TEM imaging. As described in Section 3.4, a buffer layer of metallic aluminum is deposited before the aluminum oxide is grown. Part of the purpose of this buffer layer is to prevent oxygen from diffusing into the niobium base electrode. However, if oxygen diffuses into the buffer layer, a stoichiometrically imbalanced layer can form in the aluminum instead, possibly leading to critical current noise in the Josephson barrier.

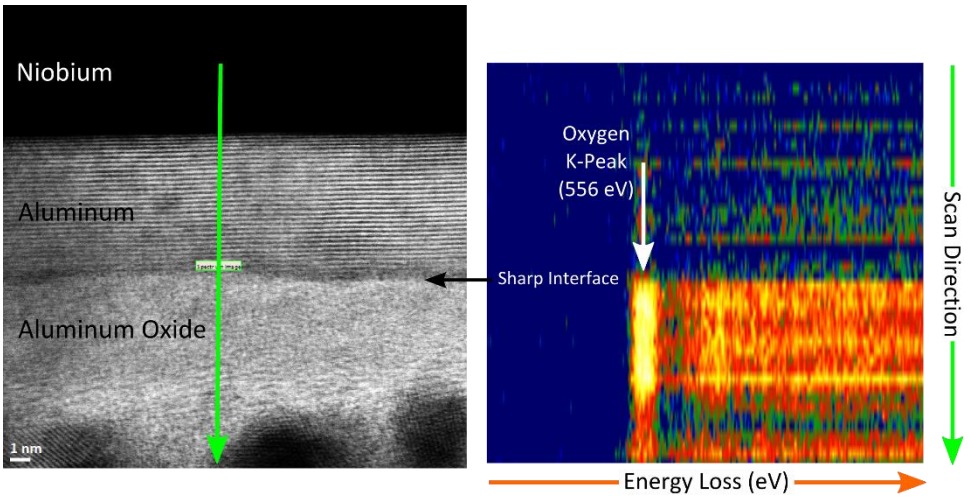

Figure 32 (Left) STEM image of the niobium, aluminum, and aluminum oxide barrier of the film under study. Green arrow shows the EELS scan line. (Right) Spectroscopy along the scan line of the left image. Bright colors correspond to species at the given energy. Bright yellow and white is due to oxygen K-peak at 556 eV.

The EELS scan shown in Figure 32 demonstrates that the oxygen is confined to the grown aluminum oxide and has not diffused to the underlying metallic aluminum. The sample studied was the same 750°C niobium film imaged for the substrate analysis in Section 4.4.2; it had also been used for calibration purposes and as such did not have a top niobium electrode deposited. The termination point of the aluminum oxide in the TEM images is thus somewhat obscured as the aluminum oxide has similar contrast to the background scattering. However, the interface between the metallic aluminum and aluminum oxide is both clear and sharply defined. EELS of this interface yields a similar sharp boundary

in the oxygen K-peak data. If oxygen were present in the metallic aluminum, the oxygen peak would not turn on sharply and would be present outside the aluminum oxide.

4.5 Summary

This chapter has characterized niobium and alumina films to determine the optimal film parameters. Transport measurements have shown maximal RRR and T_c values for films grown at 780°C and 855°C, respectively, suggesting the ideal growth temperature to be in the range of 750 – 850°C. In contrast, AFM roughness measurements show an increase in surface roughness for films grown above 750°C, with the flattest sample grown at 675°C. Furthermore, EELS has shown that oxygen is being depleted from the sapphire substrate for films grown at 610°C and above. The oxygen vacancies in the sapphire substrate may cause charge traps that interact with the electromagnetic fields required for qubit operation, leading to dephasing of the qubit. The critical temperature is still quite high (~9K) for all the measured films, leading me to conclude that the optimal growth temperature is 640 – 650°C for niobium films designed for qubit fabrication. In this temperature range the films still show good epitaxy at the substrate interface and flat surfaces where the aluminum will be deposited to form the Josephson junction barrier while limiting how much oxygen is depleted from the sapphire substrate.

In the final section, I also showed an EELS measurement of oxygen from the alumina barrier. The oxygen signal is confined to the amorphous alumina, and has not diffused into the aluminum buffer layer. This suggests that the growth pressure is sufficiently low so as not to force oxygen into the underlying buffer layer.

Chapter 5 Sample Fabrication

There were two main types of sample fabrication reported in this project. The first type of sample fabrication used four lithographic steps to create the Josephson junction and leads. This complex process had several failure modes as well as a size limitation issue. To overcome these issues, I created a new sample fabrication process with only two lithographic steps that allowed me to create sub-micron devices. I describe both types of fabrication work in the following sections.

5.1 Lithography Techniques

This section will describe the processes used to create the device topography. Each fabrication step had an initial step of defining the pattern that was either removed by etching or had additional material deposited into it. In all cases, this was done by applying a photosensitive polymer (photoresist) and then exposing that polymer to either ultraviolet light or to electrons. A developer bath then removed photoresist where it has been exposed. The photoresists that I used in this work are sensitive to both deep ultraviolet light (DUV) and electrons (e-beam). This allowed me to use a combination of both DUV e-beam exposures to create my desired pattern. I will describe this process in detail below.

5.1.1 Photoresist Application

The lithographic process for each wafer was begun by de-greasing the sample with acetone and isopropyl alcohol. The sample was then heated to 120-150°C to remove water from the sample surface. The sample would then be placed onto a spinner to ensure even application of the photoresist polymer.

I used several different photoresists in this work, though they were all some variety of (poly)methyl methacrylate (MMA or PMMA). Different combinations of photoresists produced different profiles after exposure and development – these were useful for creating different types of device structures. The resulting photoresist profile was dictated by the molecular weight of the PMMA, the ratio of PMMA in the solvent, and the spin speed used to evenly distribute the polymer. Although several resists may have been tried for any given fabrication step, I will describe the type of photoresist profile desired for each step below.

5.1.2 Deep Ultraviolet Exposure

The PMMAs that I used have a low sensitivity to deep ultraviolet light [45], but can be patterned using DUV as a source regardless. After the photoresist was applied as described in Section 5.1.1, the sample could be placed into a mask aligner. For each step, I had a pattern etched in a chrome coating on a quartz plate. Where the chrome coating had been removed, DUV light would pass through and expose

the underlying photoresist. The light source that I used produced light around 250nm with an intensity in the range of $1.5\ mW/cm^2$ to $2.5\ mW/cm^2$. I found that I typically needed a dose of around $5000\ mJ/cm^2$ which resulted in exposure times exceeding 30 minutes. Furthermore, the DUV had a resolution limit of 2-3µm and a similar limitation in lateral feature placement. For devices much larger than these limits, all the work could be done using a DUV aligner, however I also wanted to create devices with submicron dimensions, and another exposure source had to be found.

5.1.3 Electron-Beam Exposure

As noted before, the photoresists used in this project are sensitive to both DUV light and to electron beams. The electron source I used was the Raith E-Line system in MRL. The sample was placed on a holder and loaded into vacuum chamber. The sample holder could be moved with great precision, allowing for nanometer accuracy in feature placement. The size of the features depended greatly on user parameters: the choice of photoresist, the accelerating voltage of the electrons, the aperture choice, development time, as well as my skill at focusing the beam on the surface of the photoresist. In any case, the devices fabricated in this work did not push the lower limits of what the e-beam system can resolve; I was typically making 700 nm – 2 µm devices. While this was beyond what I could achieve with DUV exposure alone, it was well above the 50 nm resolution limit of the Raith system.

There was a price to be paid for the increase in resolution by moving to e-beam exposure. The e-beam system works by accelerating electrons to 10-30kV, diverting them through an aperture, and using electromagnetic lenses to focus the electrons at the surface of the sample. The end result is a small beam of electrons, a natural result given the desire to create fine features. However, this is not truly ideal when trying to create large (millimeter-scale) features such as contact pads. While the beam can be made larger by choosing larger apertures and lower accelerating voltages, it can still take hours to fully expose these large features. Additionally, my substrate was sapphire, which is highly insulating; electrons incident on this substrate had difficulty finding a path to ground, leading to a build-up of charge on the surface. This surplus charge would damage the photoresist, defeating the benefits gained from using electrons as the exposure source. Although I was able to reduce the effect of charging by using the conductive polymer AquaSAVE, the polymer was not conductive enough for me to be able to expose my contact pads using the e-beam source. Instead, I used hybrid exposures, where the sample would have both a DUV exposure and an e-beam exposure. This technique is described below.

5.1.4 Hybrid Exposures

There were two ways in which hybrid exposures were used to create photoresist patterns. The first, and simpler, way was for the sample to be placed in the e-beam system, exposed, and then placed in the mask aligner and exposed again. The wafer would then be immersed in developer and both the e-beam pattern and the DUV pattern would be developed simultaneously. This method worked whether the fine features were in isolation from the DUV pattern or not. If the e-beam pattern needed to have an overlap with the DUV pattern, all that was required was that a sufficient overlap (several microns) was built into the exposure file in the Raith. However, this method would only work on a sample that had distinguishable features on it. Without any pattern on the sample, it wasn't possible to correctly relate the DUV pattern and e-beam pattern.

In the case of an unpatterned/blank film, I would first complete the DUV exposure. I would then develop the exposed photoresist, creating a pattern in the photoresist. Since the photoresist was applied on top of a metallic niobium surface, the contrast between the conductive metal and insulating PMMA was sufficient for me to be able to coordinate the two separate exposure methods. After developing the PMMA, the sample would be placed in the Raith, exposed, and then re-developed. There was some risk of feature broadening because of the dual developments, but the DUV pattern was created so that the edges would be several microns from where the final fine features would lie, minimizing the chance of feature broadening.

5.2 "Via"-Process Fabrication Techniques

This section will describe the four-step fabrication process initially used to create devices. Each step used the lithographic techniques described in 5.1 but I will list in additional detail the photoresist choices made for each step. I will also describe the methods of removing material and depositing either insulating layers or metallic connecting wires.

5.2.1 Device Isolation

The first stage of device fabrication required isolating each device from its neighbor. At this point, the photoresist pattern was used to protect the underlying film from being attacked by the various etchants. While the photoresist was mostly able to withstand the two reactive ion etches, a thick layer of PMMA was preferred to protect the surface. Photoresist was typically applied twice to create a thick enough layer. While thick photoresist can limit the resolution of the features, at this stage the smallest feature size was usually >5µm, which was too large to be affected by the thickness of the PMMA.

The main reason the photoresist needed to be so thick is because it had to remain intact while I removed the top niobium electrode, the insulating aluminum oxide and aluminum buffer layers, and finally the bottom niobium electrode. In particular, the reactive ion etch process used to remove the niobium could damage the photoresist. The reactive ion etcher (RIE) used both chemistry (hence "reactive") and physical bombardment (hence "ion") to effect an etch on the material in question. In the case of niobium, a gaseous mixture of sulfur hexafluoride (SF_6) and oxygen was pumped into the RIE chamber. An RF field was applied to the chamber, removing electrons from the gas molecules and creating a plasma. The plate holding the sample was held at a negative voltage, driving the positive ions toward the samples themselves. The fluorine atoms chemically attack the niobium film, but the bombardment of the sample also causes material to be sputtered away as well.

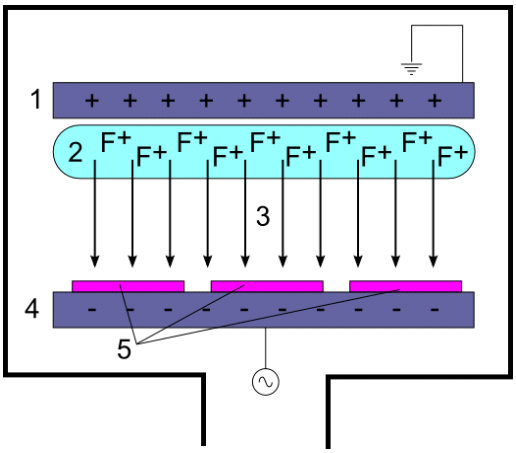

Figure 33 Schematic of a RIE chamber using fluorine chemistry.

This sputtering action was not limited to removing only the film, but could also remove the photoresist mask, thus coating the walls of the chamber with plastic residue. The residue had the unwanted effect of changing the effectiveness of my etch recipe. To compensate for this, I ran a high power oxygen cleaning process for ten minutes before loading any of my samples. I would then run my niobium etch recipe with the chamber empty; this was to ensure that all the gas lines had been fully purged before my etch recipe ran. I would then load my sample into the chamber for my etch process.

While there was some risk to my photoresist mask from the RIE, the reactive ion etch process produces a more anisotropic etch than a simple chemical bath, allowing for more control over feature size. In my case, I used a PlasmaTherm RIE to etch down 1000Å thick niobium electrodes in 2 minutes using a mixture of 20sccm SF_6 to 0.6sccm O_2, keeping the chamber at a pressure of 60mTorr and applying a RF power of 60Watts. This etch was used twice in the device isolation process, once to remove the top electrode, then again to remove the base electrode.

After the initial RIE etch, I would use a chemical bath to remove the aluminum oxide and metal layers. There are many acids and bases that attack aluminum; I primarily used a commercial blend Transene Aluminum Etchant Type A or Transene Aluminum Etchant Type D. These two blends differ only in that Type D was a nitric acid free mixture, but neither etchant attacked the photoresist mask. The aluminum

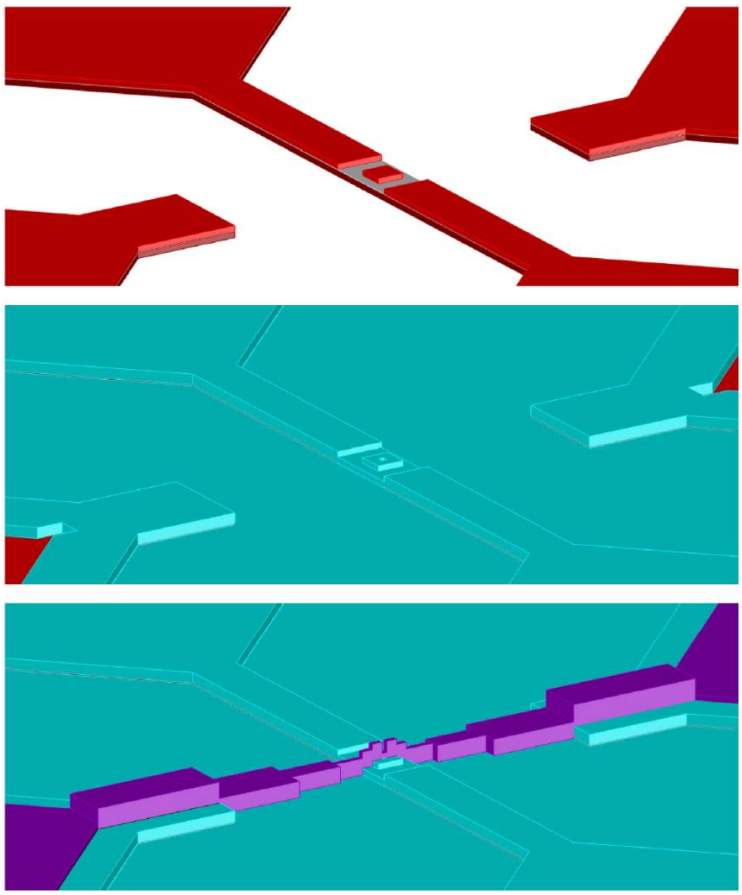

Figure 34 3D image of the processing steps used in the "via" process. In all the images, the vertical scale is exaggerated for clarity. (Top) The device state after the mesa etch step. The Josephson junction is defined by the small niobium (red) box on top of the alumina barrier (gray). (Middle) The device state following the SiO_2 deposition and via hole etch. The via hole is just visible atop the junction, and the holes atop the contact pads are also seen at the edges of the image. (Bottom) The niobium wire-up leads (purple) making contact to the top of the Josephson junction.

oxide surface was initially difficult to etch through, but once the oxide was removed the metal was removed within seconds. After thoroughly rinsing the wafer in deionized water, I would return to the RIE to remove the base electrode. The final step of the process would be to strip the photoresist mask by soaking the wafer in a bath of acetone.

5.2.2 Mesa Etch

After isolating the wires for each individual device, the next step was to define the actual Josephson junction. The Josephson junction was created by removing a ring of niobium from the device bar, leaving behind a cap of niobium which in profile resembles a mesa pyramid. For devices greater than $5\ \mu m \times 5\ \mu m$ the lithography was done entirely by DUV exposure, while for films where I attempted to create smaller devices the lithography was done entirely by Raith exposure. There were no hybrid exposures for this step. The typical photoresist for this step was the same as the photoresist used in the device isolation step.

After finishing the lithography, the sample would be loaded into the RIE, following the same procedures as those described in 5.2.1. While sulfur hexafluoride etched niobium in the RIE, the process did not attack aluminum. This allowed aluminum to act as the etch mask and the RIE etch would self-terminate after removing the top electrode. The mesa step would be finished by stripping the photoresist mask via an acetone bath.

5.2.3 Passivation Deposition and Via-Hole Etch

After the mesa etch step, the Josephson junction had been defined by the size of the superconducting cap atop the base electrode. Unfortunately, the junction needed contact to the top electrode to be probed and measured. If this were done by depositing a simple metallic lead, the lead would short to the base electrode, necessitating a passivation layer to protect the sides and top of the base layer. Thus, after finishing the mesa etch step and stripping the photoresist, I would then load the wafer into a vacuum chamber to deposit silicon dioxide or silicon nitride.

The source of silicon dioxide or silicon nitride was a Plasma Enhanced Chemical Vapor Deposition (PECVD) chamber in MRL's Microfabrication facility. To deposit silicon dioxide, the chamber walls are heated to 250°C and a mixture of silane and nitrous oxide is activated with a plasma. At the relatively high pressures (1000 mTorr) used in the PECVD, this plasma deposits a thick film of silicon dioxide on the surface of the sample. A few films were fabricated using silicon nitride as the passivation layer, but the work here is focused on samples which used silicon dioxide.

Now that 2000-3000Å of silicon dioxide had been deposited on the wafer, the entire junction structure was buried beneath an insulator. To make electrical contact, holes had to be opened in the silicon dioxide. Further, a very small hole had to be made which was aligned with the Josephson junction. While this step was initially done solely with DUV exposures, it was eventually migrated to a hybrid process where the small via-hole was exposed using the Raith and the large contact pads were exposed by DUV exposure.

The small via-hole was the most critical part of this step. Failure to fully expose the hole or to fully etch the insulating layer would result in an open circuit and an immeasurable device. If the hole was misaligned or over-large, the wire-up leads would short to the base electrode. As device sizes got smaller, these issues became ever more problematic, especially as my etch choices were severely limited.

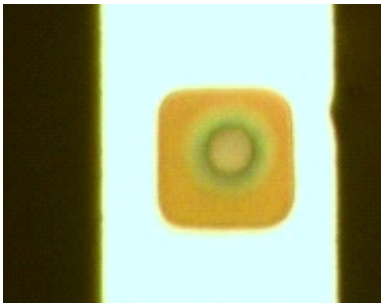

Figure 35 An example of a broadened via hole atop a niobium cap. The bright vertical bar is the device isolation bar, etched down to aluminum. The niobium metal is clearly visible inside the via-hole. With the exception of the hole, all features are capped with silicon dioxide. The ring around the via-hole is due to over-etching by the HF acid etch. Further misalignment or over-etching would destroy this device.

Silicon dioxide is commonly etched by fluorine chemistry in a RIE. Unfortunately, as discussed above, niobium is also etched by this chemistry. The etch rates for both silicon dioxide and niobium by RIE were too high for this to be a viable option for removing the silicon dioxide. Instead, a bath of Buffered Oxide Etchant (BOE) was used to remove the silicon dioxide. The BOE mixture I used was a combination of hydrofluoric acid (HF) and ammonium fluoride (NH_4F) in a ratio of 1:7. While the BOE mixture quickly removed the silicon dioxide, it did not noticeably affect the underlying niobium structure. However, as is typical with wet chemistry, the BOE etch was fairly isotropic, broadening the via-hole atop the junction. An example of such broadening is shown in Figure 35. It was not clear whether the broadening was due to a delamination of the photoresist from the surface of the silicon dioxide, or just a natural result of an isotropic wet chemical etch. In either case, the size of the broadening – approximately a micron – eventually led me to pursue another type of fabrication to create submicron Josephson junctions. This second type of fabrication will be discussed in Section 5.3, but I will first conclude the fabrication steps for the four-step via process with the deposition of the wire-up leads.

5.2.4 Wire-Up Deposition

The final step was to deposit electrical leads to make contact with the niobium cap through the via-hole. Unlike all previous steps, where the lithography was done before an etch step, for this step the photoresist was prepared for a lift-off process. After exposure and development, the wafer was placed in a vacuum deposition chamber and coated with niobium metal. When the photoresist was stripped, it would lift-off the metal atop it, leaving behind the pattern created by the exposure.

A photoresist stack of a MMA and PMMA allowed for a photoresist profile where the opening in the MMA was larger than the opening in the PMMA. This arrangement created an undercut in the photoresist which broke contact between the metal deposited in the opening and the metal deposited atop the photoresist. The thickness of the underlying MMA layer also

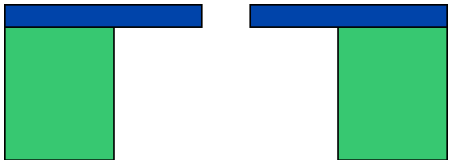

Figure 36 Example of a photoresist profile. The blue layer is the PMMA and the green layer is the MMA with an undercut to aid liftoff. Not to scale.

aided in the lift-off process, so spin speeds for this layer were usually slower than the spin speeds for the PMMA layer, resulting in a thicker base.

As with the via-hole step, the lithography for the wire-up process was done using a hybrid exposure process. The hybrid process was necessary because this step included opening large features for the contact pads as well as small features contacting the superconducting Josephson junction. A small electrical lead was desired at the cap to minimize any stray capacitance through the passivation layer to the base electrode. Further, the wire-up layer had to be precisely aligned to the via-hole previously made, which was considerably easier to achieve using the Raith e-Line system.

After finishing the lithography for the wire-up step, the wafer was loaded into a sputtering chamber in Dale Van Harlingen's laboratory. This sputtering chamber had three separate sputtering sources that could be powered with either a DC or RF power source. The chamber also featured an ion mill to allow for *in-situ* milling of samples before deposition. The chamber could be filled with argon, nitrogen, oxygen, or a mixture of gases, depending on need. The base pressure in the chamber was 1×10^{-7} Torr, which was sufficient to provide clean-enough metals for the wire-up process.

Before depositing my niobium wire-up leads, I always ion milled my sample surface. Argon atoms were introduced into the chamber, and were ionized by a Kaufman-type ion beam source. The ions were accelerated to 400 eV with a filament current around 15 mA and would slowly remove material across the wafer surface. By leaving the beam incident on the wafer for 10-30 seconds, I was able to remove any debris from the surface and mill away the topmost monolayers of material, leaving a fresh surface for the wire-up metal to bond with.

Immediately following the ion mill step, the sputter gun would be activated. Argon was also the gas I used for the sputtering process, and the chamber pressure was raised to 10mTorr. The argon atoms would be ionized into a plasma, removing material from the sputtering target and re-depositing that material on my film. I would allow the sputter gun to run for 3-5 minutes before depositing any material on my sample to ensure cleanliness of the finished product. I would typically deposit 1000Å in approximately 2 minutes. The wafer was then removed from the deposition chamber and immersed in an acetone bath. At this point, the niobium coating the photoresist was lifted-off as the photoresist was simultaneously stripped from the sample.

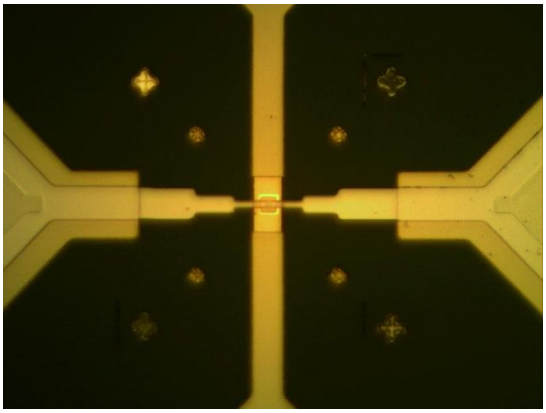

Figure 37 Typical finished Josephson junction device layout. This particular device is e423 Device 11. The device isolation bar runs vertically and the wire-up bar runs horizontally. The eight crosses are alignment marks.

Niobium was the metal that I chose for the wire-up leads because of its superconducting properties. The niobium produced in the Van Harlingen sputtering chamber usually had a critical temperature of around 8.0 – 8.5 Kelvin, which was substantially lower than the critical temperature of 9.2 Kelvin which was typical of MBE-grown films. However, even this lower transition temperature was sufficiently high for the leads to be completely superconducting at the measuring temperature of 4.2K. By ensuring that the leads were in the zero-voltage superconducting state, my electrical measurements effectively had four-point contacts even though the top and bottom electrodes were continuous metal wires (see Figure 37).

After rinsing the wafers in isopropyl alcohol, they were ready to be inserted into a cryostat and measured. However, as noted above, there were several issues with the fabrication steps described above. The first issue was the complexity of the process; there were four quite lengthy steps to complete for each wafer with many points of failure. This was not an insurmountable hurdle, as it merely necessitated careful handling on my part. The larger issue was the broadening of the via-hole described in Section 5.2.3. Moreover I ultimately wanted to be fabricating devices in the 100-200nm range, and it was far from obvious that I would even be able to make contact through a via-hole to the underlying junction, even if I could successfully make a hole small enough for the job. An entirely new fabrication process was developed that featured only two steps and no similar limitation on the junction size. This process is described below.

5.3 Self-Aligned Fabrication

This section will describe the two-step fabrication process used to create submicron Josephson junctions. In the following sub-sections I will detail the methods used to isolate each device, protect the base electrode from shorting to the wire-up leads, and finally the way in which the Josephson junction was defined.

5.3.1 Device Isolation

The primary problem with the four-step via-hole fabrication process was the wet etch through the silicon dioxide passivation layer. The two-step self-aligned process was envisioned when it was realized

that we could use a different material, and more importantly a different technique, to achieve the same end: protecting the base electrode from making contact to the wire-up niobium. Any sufficiently insulating material would do; I opted to work with aluminum oxide deposited by atomic layer deposition (ALD).

The self-aligned process began with the same base structure: an unpatterned wafer composed of a base niobium electrode, an aluminum buffer metal, a thin aluminum oxide barrier, and a top niobium electrode. I then coated the wafer with a thick layer of photoresist and placed it in the DUV aligner. The lithography for the device isolation step was a hybrid process as explained in 5.1.4 for unpatterned films. After the DUV exposure, I would be left with a 10µm-wide bar from which I would carve out the junction. I would then use e-beam lithography to expose the photoresist, leaving only a thin strip that was the approximate width of my final junction.

The next step in the device isolation process was to etch down to the substrate in all the places left exposed by the photoresist. For the self-aligned wafers, I used a March RIE to remove the niobium layers. In the March RIE, the recipe to etch niobium was 20 SCCM of SF6, at a pressure of 200mTorr and a power of 30Watts; this recipe would remove 1000Å in 80 seconds. As in the via-hole process, the aluminum oxide layer and metallic aluminum buffer layer were removed by a wet chemical etch. The wafer was then returned to the March RIE for the removal of the base electrode.

With the MBE structure removed, the wafer was then placed into the ALD chamber with its photoresist mask intact. ALD works by placing the wafer into a low vacuum chamber, which can be filled with various gases. The chamber floor and walls can be heated, creating a reactive environment. A precursor gas of water is released into the chamber, coating the sample and creating dangling hydrogen bonds for the next gas to bond with. The second, “deposition”, gas is then released into the chamber. The methyl groups of the deposition gas bond with the hydrogen atoms, and the new molecules are then pumped away, leaving behind a coating of aluminum oxide over the entire structure.

Because ALD used a diffusive gas as the deposition source, there was no line-of-sight limitation on where new material was deposited. As a result, there was an equal coating of material on the sidewalls of features as on the top. For my devices, I simply needed a passivation layer protecting the base electrode from contact to the wire-up metal. Furthermore, the layer I deposited only needed to be thick enough to prevent transmission through the passivation layer, so coatings were only about 50Å thick.

Of course, using ALD to coat the sidewalls of my MBE structure also resulted in an alumina coating on the top of the photoresist capping the structure. When the photoresist was removed, the ALD alumina would be lifted-off, leaving the top of the MBE structure exposed. However, just as the ALD alumina coated the sidewalls of the MBE structure, so too did it coat the sides and top of the photoresist.

I developed a method of chemical-mechanical polishing to remove the photoresist and lift-off the encapsulating ALD alumina. After the ALD process, the sample was immersed in a bath of 70C MicroChem Remover PG for 10 minutes. Then the sample was abraded with cotton swabs and returned to the Remover PG bath for another 10-minute soak. The abrasion was then repeated, and then the wafer was visually inspected. The photoresist was thick enough for there to be a clear color difference relative to the niobium top electrode – if necessary, the sample would be returned to the PG remover bath for a third soak and abrasion.

The device isolation step relied on the successful lift-off of the ALD alumina. To this end, the choice of photoresist was critical. The photoresist had to stand up to both of the RIE etches and still be thick enough for the lift-off process. The PMMA produced from MicroChem comes in two molecular weights (495,000 and 950,000) and several solid concentrations in anisole. After much trial and error, I found that the best photoresist mask to use was a stack of three layers. The bottom layer was 495K PMMA A2 (2% solids in anisole), the middle resist was another layer of 495K PMMA A2, and the top layer was 950K PMMA A4 (4% solids in anisole). The thickness of each layer of PMMA A2 was 700Å and the PMMA A4 had a thickness of 3000Å [45]. This photoresist stack was sufficiently thick to withstand the RIE and lift-off the alumina, but was not so thick to limit the size of the device.

5.3.2 Wire-Up Deposition and Etch Back

Following the lift-off of the ALD alumina, the device isolation step was complete. The wafer would then move on to the wire-up deposition and etch back step. This section will describe the lithography and fabrication steps involved in the completion of the Josephson junction circuit.

I began the wire-up process by spinning a stack of MMA and PMMA. As with the wire-up lithography described in Section 5.2.4, the purpose of this bi-layer was to aid in the lift-off of the metallic contacts. The sample was then exposed by a hybrid exposure of electron beam and deep UV lithography. The electron beam lithography was used to define the contacts at the device isolation bar, while the deep UV lithography opened up the areas at the contact pads.

In the self-aligned process, the width of the Josephson junction was defined by the device isolation step; the length of the junction was defined at the wire-up step. This was why it was crucial to use electron beam lithography to expose the photoresist at this step – the width of the metal contact became the length of the junction. The narrower the metal contact, the smaller the junction ultimately was.

After the lithography, the sample was loaded into the Van Harlingen sputtering system described in Section 5.2.4. As in the four-step process, the wafer was ion-milled for 3-10 seconds. It then had 1000Å of niobium deposited by sputtering onto it. I would then deposit 200Å of aluminum, again by sputtering, on top of the niobium. The aluminum was deposited to be a mask for the niobium for the final etch back step in the process. The photoresist and overlying metal would then be lifted-off in an acetone bath.

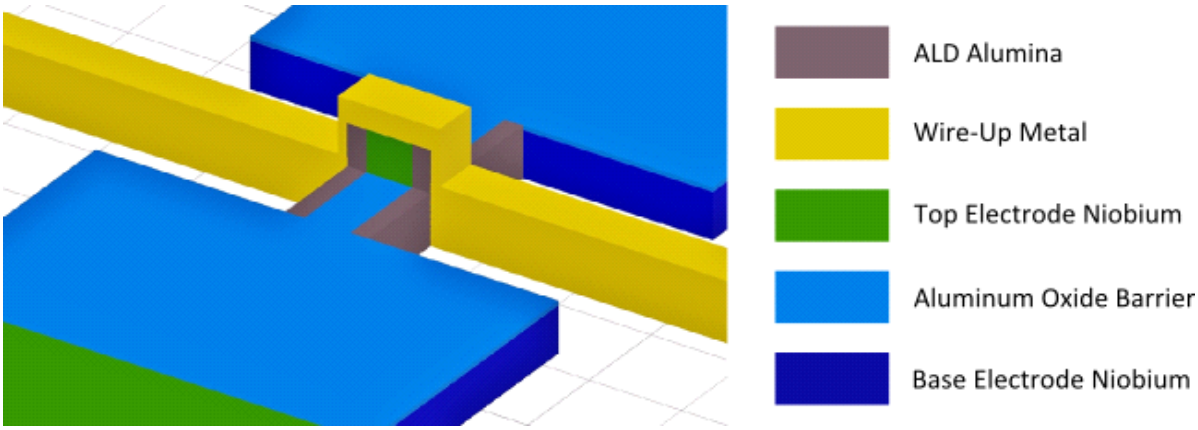

Figure 38 3D view of a Josephson junction made using the self-aligned process. The junction is formed by the top niobium electrode (green) at the center of the image. The conformal ALD alumina (gray) is only shown near the barrier – in reality it coats all of the edges of the device, but this is not shown here so that the MBE film structure is visible.

Following the metal deposition and lift-off, the sample would then be placed into the March RIE. Anywhere that hadn't been covered with metal deposition would either be sapphire substrate, or niobium top electrode. The top niobium electrode was thus removed in all places not covered by the wire-up contact. It was this "etch back" step that finally defined the size of the Josephson junction and the junction circuit was complete.

Compared with the four-step fabrication process, the two-step self-aligned process had several benefits. The first was the relative simplicity of the process – there were only two lithographic steps as opposed to four steps with the via-hole method. A second benefit was that at no point was there a critical alignment of patterns. In the four step process, there were two critical alignments: the via-hole had to be precisely aligned atop the mesa and then the wire-up had to be precisely aligned to the via-hole. Because the self-aligned process is just a cross pattern, there isn't a comparable necessary alignment. Finally, the third benefit of the self-aligned process was that the size of the junction was determined by

two lithographic steps. This meant that I could focus on shrinking the size of each lateral dimension separately, making it easier to create smaller junctions.

5.4 Resonators and SQUIDs

In addition to the single Josephson junctions, I also made coplanar waveguide resonators and SQUIDs to evaluate noise in MBE-grown materials. The resonators were made solely from 3000Å thick niobium films and the fabrication process used the device isolation RIE etch described in Section 5.3.1. Measurements of the quality factor from these resonators revealed information about the niobium used in making Josephson junctions and qubits. The results of the quality factor measurements are discussed in Section 6.1.

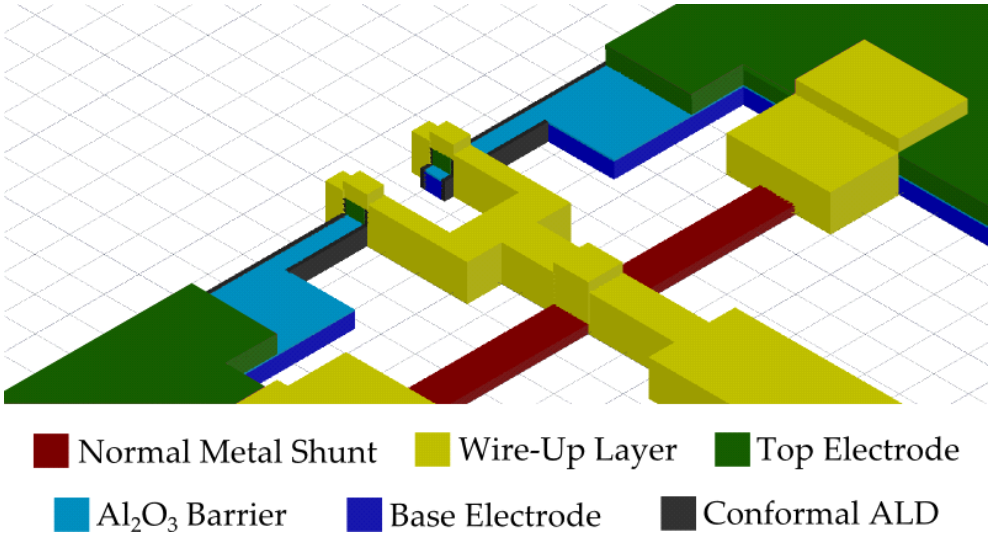

Figure 39 3D detailed view of the Josephson junctions in a dc-SQUID made using the self-aligned process. The color scheme is the same as in Figure 38 with the resistive shunt added in red. The SQUID loop extends to the upper left (not shown).

Two-junction SQUIDs (also referred to as dc-SQUIDs) were used to measure noise due to magnetic impurities, an important concern in flux qubits. The fabrication process was nearly identical to that of a simple Josephson junction, with the addition of a resistive shunt. The shunt was a multilayer stack of titanium and gold such that it would remain normal down to temperatures as low as 10 milli-Kelvin and still have several Ohms of resistance. The shunts were used to remove hysteresis from the transport characteristic for the device and allow noise measurements to be made. SQUIDs were made with both the via-process and with the self-aligned process; the results of the noise measurements are discussed in Section 6.3.

Chapter 6 Transport Characteristics and Noise Measurements

In this chapter I will discuss measurements made from resonators, Josephson junctions, and SQUIDs. While microscopy measurements such as those in Chapter 4 can identify individual defects that might give rise to noise, transport measurements measure the effect of these defects on the characteristics of the superconducting device. These can be correlated with growth conditions and film microscopy to identify the physical mechanism giving rise to the noise. The transport measurements are then used to determine the ideal growth parameters (such as temperature, oxygen pressure, and aluminum flux rate) for flux qubits and transmons.

I shall begin by describing internal-Q measurements from resonator structures made from niobium films, which are a simple structure that can be used to evaluate just the niobium quality. I then present Josephson junction measurements, paying particular attention to the temperature dependence of both the superconducting order parameter and the critical current. I also report critical current noise measurements from some of these junctions. Finally, I return to the question of flux noise and examine dc-SQUIDs made using two separate processes to understand the effect of epitaxy on flux noise.

6.1 Resonators

Qubits are operated at microwave frequencies, and the materials used to construct a qubit may behave differently in these rapidly changing electromagnetic fields compared to DC measurements. Building a resonator circuit with a resonant frequency in the 1-10 GHz range allowed me to evaluate both my qubit materials as well as my fabrication techniques. These devices were only built from niobium films, this allowed me to focus on noise sources only from the niobium.

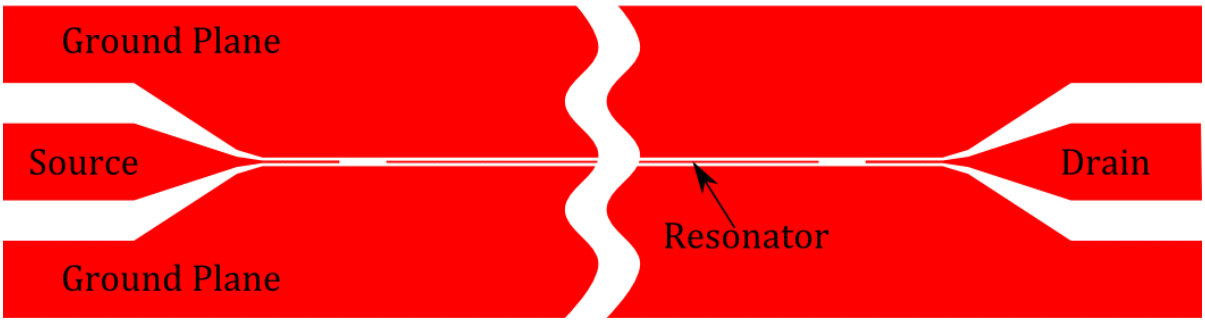

Figure 40 A cartoon of a co-planar waveguide resonator. The areas in red are metallic, the white areas are etched down to the sapphire substrate. The length and width of the resonator determine the resonant frequency -- different resonators have different geometries to produce the desired resonance point. The gaps between the source, drain, and resonator are exaggerated for clarity.

These samples used a co-planar waveguide structure as the resonator. Co-planar waveguides are fabricated with the resonator separated from a source and drain by a gap which determines the coupling strength for the device. Metal structures on either side of the resonator are grounded – these structures may or may not have holes in them to fix vortices in the superconductor.

These resonators are evaluated by measuring the quality factor. In general, the Q-factor has the property $Q(\omega) = \omega \cdot \frac{\text{Maximum Energy Stored}}{\text{Power Loss}}$, so the sharpness of the resonance peak relates to losses in the material. Defects which can lead to power loss should then result in lower Q-factors.

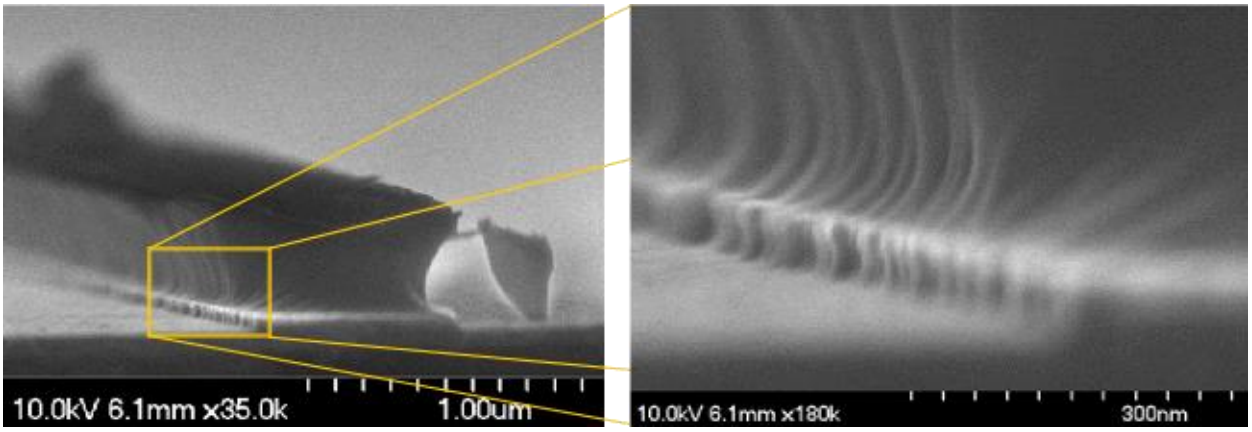

Figure 41 SEM micrographs of photoresist patterns being transferred to the underlying film. (Left) Gross view of the photoresist pattern on silicon. (Right) Zoom view of the boxed region showing wave pattern in the photoresist and metallic film.

Work done by Wenner et al [46] suggests that noise in resonator circuits is primarily contributed by interactions at the surfaces of the samples with the most significant contribution coming from the substrate-air interface. Based on this work, it is likely that the limiting factors on Q-values for my resonators come from two sources: the substrate-air interface and the smoothness of the resonator lines themselves. With regards to the second concern, I studied the structure of the PMMA/MMA resist by growing sputtered niobium-aluminum-niobium trilayers on silicon wafers. The samples went through the same fabrication process as the MBE-grown sapphire chips, but were then broken into small pieces so that cross-sectional SEM images could be taken. The PMMA resist shows distinct wave patterns, which translates down into the etched metal film. This roughness may distort the electric fields in the resonator and lead to loss. Unfortunately, while I was able to determine that I have a rough resist lithography, I was unable to find a smoother PMMA recipe, so I was not be able to determine the effect of this roughness on my resonators. I note it here to make the reader aware of this potential issue.

In the data presented below, there is a distinction between external-Q and internal-Q. The external-Q is based on the coupling between the resonator and the source/drain; high external-Q corresponds with

weak coupling. The external-Q for the resonators that I fabricated is approximately 200,000. The internal-Q is intrinsic to the losses in the material under study, and is the value of interest in these measurements. Measurements of the Q-factors were made by Steven Weber and Kater Murch, in Irfan Sidiqqi's group at the University of California, Berkeley.

SAMPLE	GROWTH TEMP	Q_{int} $(\times 10^5)$	RMS ROUGHNESS	CLEANED	FABRICATION ISSUE
E410	652	3.2		No	
E439	610	1.5	2.389	No	
E440	675	2.4	0.214	No	
E444	676	0.7	0.333	No	Sapphire dust
E445	681	2.4		No	
E545	667	1.3	0.234	Yes	Ti Backside Coating
E547	667	2.5		Yes	

Table 2 A list of the internal-Q for the co-planar waveguide resonators. Also listed are the RMS roughness from the AFM microscopy and notes about the device fabrication.

In Table 2, I report the internal-Qs for seven different resonator films. With the exception of wafer e547, the niobium thickness was 3000Å. This relatively thick film was deposited in two stages – 1500Å of material was deposited, then the system was allowed to cool before the second half of the film was evaporated. The temperature measurements reported here are the growth temperature recorded at the midpoint of the growth. The final measurement of temperature on the film was within 10°C of the

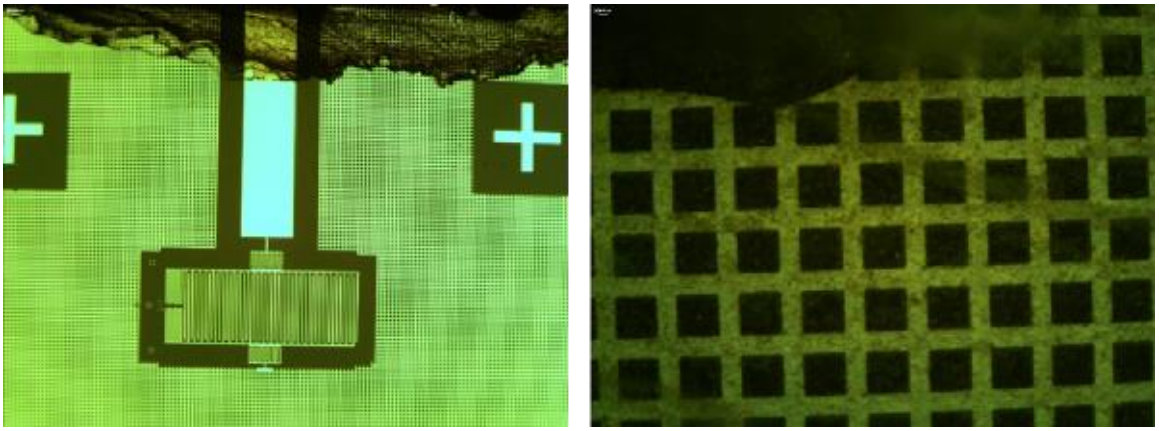

Figure 42 Sapphire dust residue remaining after dicing the wafer into smaller chips. (Left) Low magnification zoom of a flux qubit resonator structure (qubit structure not otherwise discuss in this work). The residue covers the top 15% of the image. (Right) 80x zoom of the ground plane at the top of the chip.

midpoint measurement. The sole exception, film e547, was not originally intended to be a resonator film and was only 350Å thick.

Inspection of the data shows that neither the growth temperature nor the surface roughness (RMS roughness) is correlated with internal-Q. The sample with the worst internal-Q, e444, is relatively smooth but grown at the same temperature as one of the samples with the highest Q, wafer e440. However, both samples e444 and e545 had issues which may have contributed to the low Q. To measure each resonator, the 15 x 15 mm^2 wafers were diced into smaller chips. This dicing process produced a tremendous amount of sapphire dust, which tended to settle on the chips as in Figure 42 Sapphire dust residue remaining after dicing the wafer into smaller chips. (Left) Low magnification zoom of a flux qubit resonator structure (qubit structure not otherwise discuss in this work). The residue covers the top 15% of the image. (Right) 80x zoom of the ground plane at the top of the chip.. This dust led to loss in the resonators and resulted in a lower Q. The debris could be limited by using a thick AZ 5214 resist to protect the sample during the dicing process, but was not always successful.

Another problem which could lead to loss was incomplete removal of the titanium backside coating. This coating (see Section 3.3) was required for the high-temperature growth of the niobium. However, if the coating were to remain in place, it would act as an additional (unwanted) ground plane for the resonator. The titanium could be removed by a 3-5 minute etch in BOE and was typically removed while the sample was still protected by the AZ photoresist used during the dicing process. Incomplete removal of the titanium was reported to have been a problem with sample e545, while sample e444 had issues with sapphire dust. Moreover, as most of these wafers were grown at approximately the same temperature, the data suggest that the device fabrication plays a critical role in improving the internal-Q of the resonators.

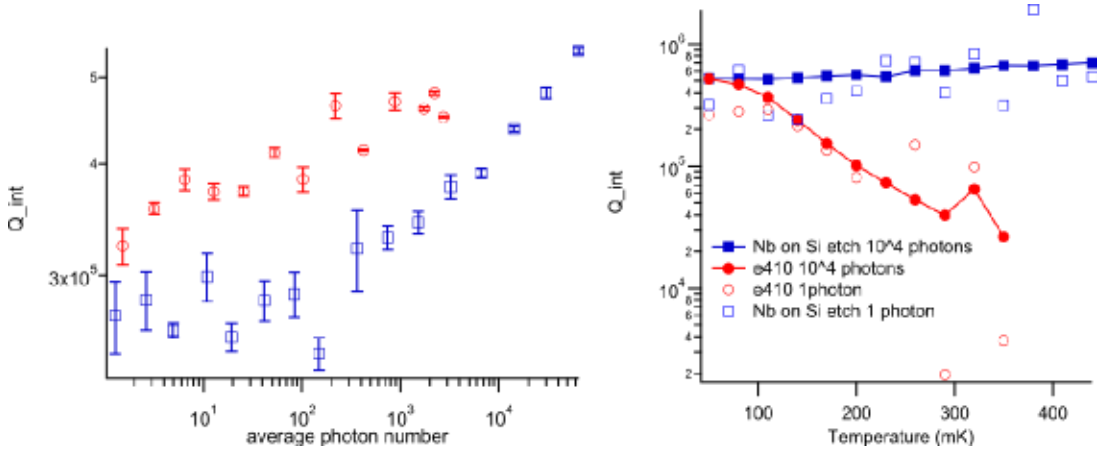

Figure 43 Internal-Q measurements comparing epitaxial niobium on sapphire to polycrystalline niobium on silicon. (Left) Comparison of the epitaxial sample (red) to the polycrystalline sample (blue) at various photon numbers measured at a temperature of 20 milli-Kelvin. (Right) Temperature dependence of the internal-Q's for the epitaxial sample and polycrystalline sample. Data was taken at photon numbers of 1 and 10,000 for each sample. Figures courtesy Irfan Siddiqi, UCB.

The sample with the highest internal-Q, wafer e410, had internal-Q measurements made at multiple photon powers and temperatures. These measurements are contrasted with a wafer of niobium-on-silicon grown and fabricated at the University of California, Berkeley. The internal-Q of e410 is only moderately better than the niobium-on-silicon sample, although e410 was grown as a single crystal film on sapphire, and the niobium-on-silicon wafer is polycrystalline. This is similar to work done by Megrant *et al* studying aluminum resonators grown by sputtering or by MBE [47]. Megrant is primarily concerned with finding what substrate preparation is required to achieve an internal-Q of over one-million but he also compares sputtered, polycrystalline aluminum films with single-crystal MBE-grown aluminum films with no cleaning. For these films he sees nearly a factor of 2 improvement in the internal-Q at single photon powers (from 4×10^5 to 7.6×10^5) for the resonators with comparable widths and resonant frequencies to my devices. My crystalline niobium film only has an internal-Q which is 1.23 times better than the polycrystalline film at single photon powers. It is not obvious why there is so little improvement in the internal-Q for the niobium samples. One possibility is that the niobium growth temperature wasn't optimized for good characteristics. However, TEM and RHEED show good crystallinity for samples grown at ~670°C (see Sections 4.3 and 4.4.1), so this seems unlikely. Another possibility is that the substrate preparation isn't optimal; this in particular was studied by Megrant *et al* and allowed them to achieve internal-Qs of over 1 million in their aluminum resonator [47]. Further work studying substrate preparation might make similar improvements for niobium.

6.2 Josephson Junctions

In qubits, the Josephson junction is merely a component, albeit it a crucial one. However, Josephson junctions by themselves can reveal useful information about the materials composing them by measuring the current-voltage characteristics and studying the critical current dependence on temperature. Measurements that were made at 4.2 Kelvin were measured in the Eckstein lab in a simple helium Dewar with no magnetic shielding. Battery-operated preamplifers were used with either a PC generated voltage signal or a function generator, depending on the vintage of the device. These measurements are used to analyze the subgap behavior of niobium based Josephson junctions. Measurements at temperatures from 9 Kelvin down to 33 milli-Kelvin were made while working with students in Dale van Harlingen's group. For those measurements, three devices from two wafers were measured in an Oxford Instruments Helium-3 system which was primarily used to measure critical-current fluctuations in Josephson junctions [48]. The system used battery-operated preamplifiers and potentiometers to measure the transport characteristics of my devices. Both the critical current and the

superconducting gap can be found from these measurements; I will show that I see different temperature dependences from these two measurements.

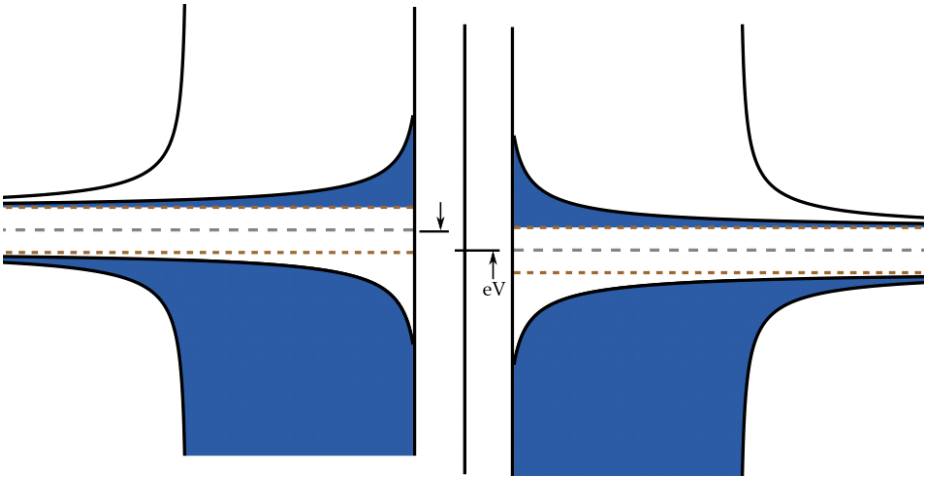

Figure 44 Semiconductor diagram showing transport between two insulators. Energy is plotted vertically, while the density of states of the superconductors are horizontal. Blue indicates filled states; filled states above the superconducting gap are excitations due to T > 0.

6.2.1 Superconducting Order Parameter and Subgap Conductance

Transport in Josephson junctions can be modeled using a "semiconductor" diagram as shown in Figure 44. In a normal metal, electrons fill the states up to the Fermi energy, and tunneling only occurs when an electron (or hole, if tunneling from right-to-left) can fill a vacant state. The same process occurs for superconductors, but with different transport characteristics due to the density of states of the superconductor.

Here, transport is due to "quasiparticle" tunneling between the superconductors. Quasiparticles are generated either from the breaking of a Cooper pair, or from excited electrons [49]. For identical superconductors, as the voltage bias increases, the available vacancies decreases and a negative resistance can be observed. When the voltage bias is twice the superconducting gap, the filled states below the gap can tunnel to the other electrode and there is a sudden increase in the tunneling current. The tunneling current can be calculated from the following equation [49, 50]:

$$I_{ss} = \frac{G_n}{e}\int_{-\infty}^{\infty} \frac{|E'|}{\sqrt{E'^2 - \Delta^2}} \frac{|E' + eV|}{\sqrt{(E' + eV)^2 - \Delta^2}} [f(E') - f(E' + eV)]dE'$$

Here, G_n is the normal-state conductance of the device and Δ is the superconducting order parameter. The first two fractions in the integrand are the density of states for the two superconductors, and the

final term is the Fermi distribution because the carriers are electrons. There is an inherent temperature dependence in the superconducting order parameter and in the Fermi distribution function.

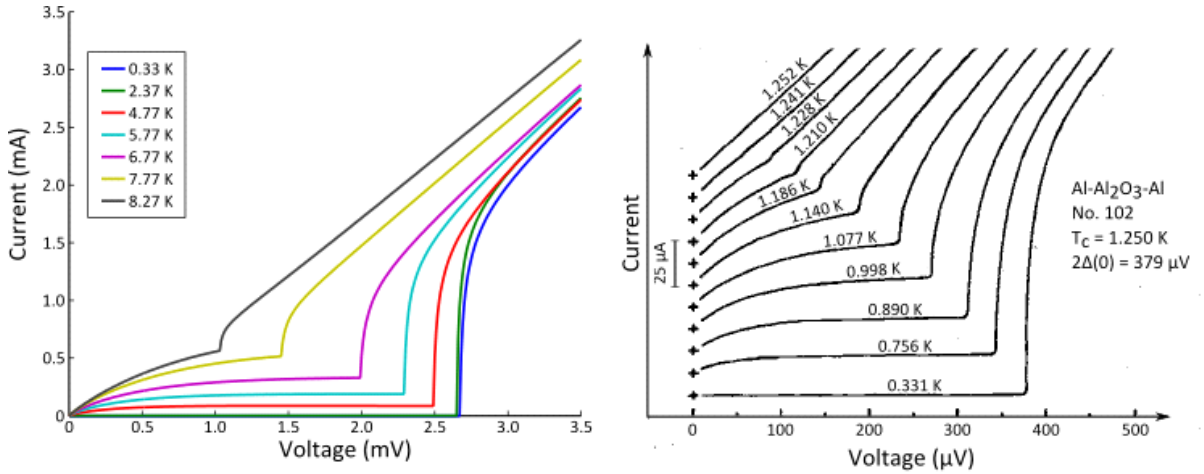

Figure 45 (Left) Tunneling currents calculated by numerical integration in Matlab for niobium electrodes. Temperatures used in the model were from the measurements of wafer e514 (discussed below). A representative subset of temperatures is shown. (Right) Transport of aluminum – aluminum oxide – aluminum junction taken from [51]. Each curve has a vertical offset for clarity.

I ran numerical integration in Matlab to solve for the tunneling currents to compare with the voltage-current measurements; these are seen for positive voltages in Figure 45. The tunneling current is essentially flat until the sharp turn-on at 2Δ. In Figure 45, I also show data from all aluminum junctions, as measured by Blackford and March [51]. A visual comparison of the experimental data shows close agreement with the theoretical curves; my experimental voltage-current curves for niobium-based junctions do not follow this theory as precisely. I refer to the flat region as the "subgap" and the vertical section as the "quasiparticle current" which is simply a shorthand for the quasiparticle current turn-on region, since all of the carriers here are quasiparticles.

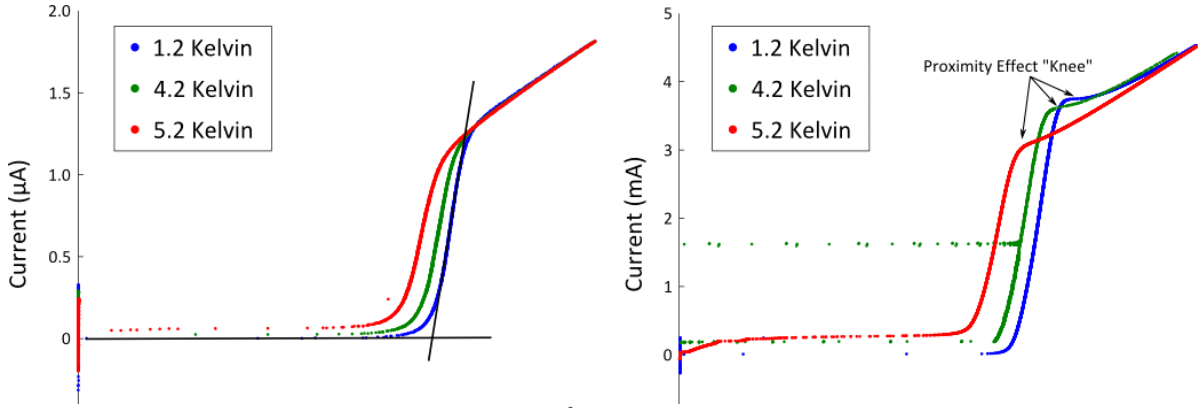

Figure 46 (Left) Data measured from film e514, with a 50Å aluminum buffer layer. Black lines are added by hand to determine the superconducting order parameter. The intercept point of the black lines is twice the superconducting order parameter (see text). (Right) Data measured from film e302, with a 200Å aluminum buffer layer. The "knee" at the top of the quasiparticle turn-on is due to the proximity effect.

My niobium-based junctions show a flat subgap, as in the theory but the quasiparticle current turn-on at approximately 2.5 mV is concave, rather than having a discontinuity like in the all-aluminum junctions. Furthermore, the quasiparticle current has a noticeable slope, in contrast with the predicted tunneling current which is essentially vertical. In the following, I shall discuss several possible explanations for this behavior, including thermal fluctuations, the superconducting proximity effect, the strong-coupling nature of niobium, and impurities/material defects in the superconductor.

Thermal fluctuations in the superconductor are certainly capable of "smearing" the density of states and providing additional current paths close to the gap. However, if the concavity was due to thermal fluctuations, it should be a small (if not negligible) effect at 0.33 Kelvin. The concavity is equally strong at both 5 Kelvin as at 0.33 Kelvin, so thermal fluctuations do not explain this behavior.

It is also possible that the superconducting proximity effect could be modifying the behavior of my junctions. If present, the proximity effect would arise from the aluminum buffer layer between the base niobium and the insulating aluminum oxide. Indeed, samples with thick aluminum buffer layers show a "knee" at the top of the quasiparticle turn-on due to the proximity effect of the aluminum film [52, 53, 54], something which I see in films with 200Å thick aluminum buffer layers but which is not present for films with 50Å thick buffer layers (see Figure 46). However, I observe similar concavity and relative slopes for the quasiparticle current regardless of the buffer layer thickness. If these features were due to the proximity effect, they should vary with the metallic aluminum thickness.

There is some evidence for an anisotropic gap in niobium films [55] which could theoretically explain this behavior. However, observation of this gap in a tunnel junction would require a crystalline interface with the niobium film. Both the amorphous nature of the aluminum oxide and the buffer layer of aluminum should preclude observation of this effect. Moreover, the effect is more pronounced with magnetic fields – the magnitude of the concavity at zero field is not consistent with gap anisotropy.

Another possible explanation for the concavity of the quasiparticle current is that I haven't properly accounted for the strong-coupling nature of niobium. The tunneling current theory assumes a weak-coupling superconductor, which I have applied here unmodified as a guide. However, modifying the tunneling theory for a strong-coupling superconductor is a small correction [56]. Moreover, other experimenters have also observed some concavity in niobium-based junctions [57, 52] while others [53, 54] see tunneling behavior consistent with the theory discussed above. This suggests that sample

preparation is the most important factor in controlling the voltage-current characteristics of niobium-based Josephson junctions.

Much of the literature that examines the voltage-current characteristics is concerned with heating the substrate when applying the aluminum [57, 58, 53]. The concern is that a hot niobium substrate is difficult for the aluminum to wet which will lead to leaky tunnel barriers. In the cited works, the authors are concerned with any subgap leakage, and do not address any discrepancies between the quasiparticle current turn-on and theory. However, my fabrication procedure is similar to the methods in [57, 58, 53], waiting 180 minutes after finishing the base electrode deposition before beginning to deposit the aluminum. My ability to obtain RHEED oscillations is also an indicator that I have flat aluminum films which wet the niobium well. I therefore find that residual substrate heat is not an adequate explanation of the observed behavior.

The main deviation between my samples and the literature is that my niobium was grown epitaxially, whereas the cited works had polycrystalline or amorphous niobium electrodes. As discussed in Section 3.3, to grow epitaxial niobium, I heated my sapphire substrates to temperatures greater than 600°C. Using STEM and EELS, I have shown in Section 4.4.3 that at these elevated temperatures oxygen sublimates from the sapphire into the niobium film. Some species of niobium oxide are superconducting [59], but at lower critical temperatures (and hence with a different density of states) than metallic niobium. I suggest that it is because of this adsorbed oxygen in the niobium film that the quasiparticle current is more concave than expected from theory.

This concavity is important because it impinges on my ability to measure the superconducting order parameter for each device. When the quasiparticle turn-on is sharp, the discontinuity occurs at twice the superconducting order parameter. Because of the roundness of this transition, I have had to find a consistent measure of the quasiparticle turn-on to determine the temperature dependence of the superconducting order parameter. I manually added lines to each current-voltage characteristic to measure the quasiparticle turn-on point. One line follows the flat subgap region and is merely an extension of this flat zone. The second line follows the tangent of the quasiparticle current; as it is not clear whether the bottom or the top of the quasiparticle current is the most appropriate choice for the superconducting order parameter, by choosing the tangent I average bottom and top. Furthermore, choosing the tangent allows me to be very consistent in my measurement. Twice the superconducting order parameter is the intercept point between the flat subgap line and the tangent of the quasiparticle current. I estimate I find this point with less than a 5% margin of error.

I then fit my experimental values to the superconducting order parameter which can be found from the implicit relation from BCS theory [60]

$$\frac{1}{N(0)V} = \int_0^{\hbar\omega_D} \frac{\tanh\left[\left(\epsilon^2 + \Delta^2(T)\right)^{1/2}/2k_BT\right]}{\left(\epsilon^2 + \Delta^2(T)\right)^{1/2}} d\epsilon$$

In this integral, $N(0)$ is the density of states at the Fermi level, $\Delta(T)$ is the superconducting order parameter, and ω_D is the DeBye frequency for niobium. I used numerical integration in Mathematica to solve for the superconducting order parameter for niobium with a critical temperature of 9.3 Kelvin and for niobium with a critical temperature of 8.77 Kelvin. The base electrodes typically have critical temperatures greater than 9.0 Kelvin (see Section 4.1), but the top electrodes and wiring circuits have lower critical temperatures; the critical temperature of the Josephson junction depends on all the components in the circuit.

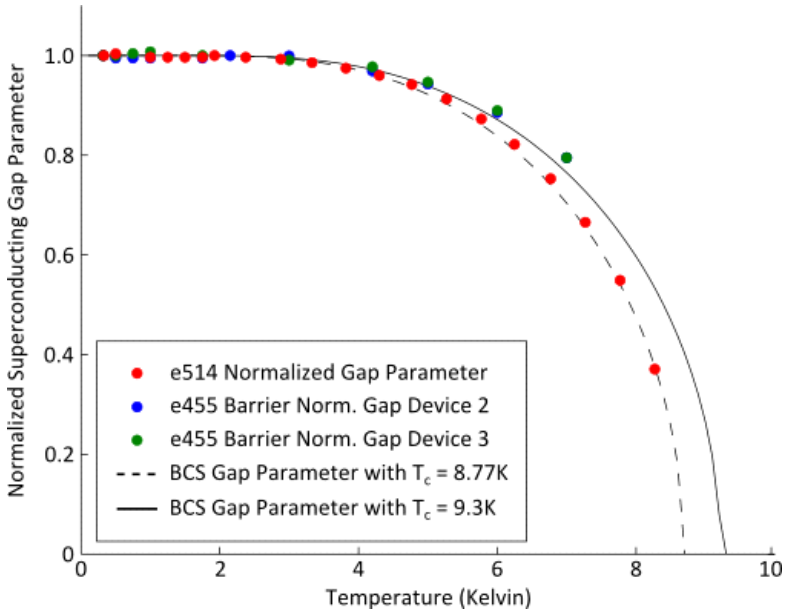

Figure 47 Experimental superconducting order parameter temperature dependence. The solid black line shows the BCS superconducting gap for niobium with a critical temperature of 9.2 Kelvin. The dashed black lines shows the same BCS gap but with the critical temperature set to 8.77 Kelvin, which is the critical temperature of the e514 device (red data). The red dots show the superconducting gap measured from film e514 (see text for details on growth). The blue and green dots show the superconducting gap measured from two devices on film e455.

Figure 47 shows the superconducting order parameter temperature dependence for three devices measured down to 320 milli-Kelvin. The BCS order parameter theory is shown as the black lines – the solid line is for niobium with a critical temperature of 9.3 Kelvin, and the dashed line is for niobium with a critical temperature of 8.77 Kelvin. The two wafers that were measured both had 1000Å thick niobium base electrodes; wafer e455 (blue) was grown at 637°C while wafer e514 (red) was grown at 706°C.

Wafer e455 had a 200Å thick aluminum buffer layer whereas wafer e514 only had a 20Å thick aluminum buffer layer (see section 3.4 for additional information about why this buffer layer was used). It is possible that this aluminum layer might alter the effective superconducting gap via the proximity effect, but I see good agreement with BCS and will neglect the proximity effect for this work. The maximum percent difference between the theory and experimental values is 3.73% and the average percent difference between theory and experiment is less than 1%.

In spite of the concavity in the quasiparticle current that suggests some type of defect in the junction (possibly due to depleted oxygen in the niobium electrode), there is good agreement between the temperature dependence of the superconducting order parameter and the BCS theory. This is an indication of high quality films with long range order that should have little noise if placed in a qubit circuit.

Examining the transport of the subgap region provides another way to fit for the superconducting order parameter, as well as qualitatively measuring how well the experimental data follows the theoretical subgap current. Recall that, as per the semiconductor picture of tunneling between the two electrodes, transport below the superconducting gap should only occur due to thermal excitations of the density of states. Excess current indicates additional channels for tunneling in these devices, likely due to oxygen vacancies in the barrier.

For the voltage region below $2\Delta/e$, the current is approximated by the function

$$I = \frac{2G_n}{e}\, e^{-\Delta/k_BT} \sqrt{\frac{2\Delta}{eV+2\Delta}}\,(eV+\Delta)\sinh\left(\frac{eV}{2k_BT}\right) K_0\left(\frac{eV}{2k_BT}\right)$$

where K_0 is the zeroth-order modified Bessel function, Δ is the superconducting order parameter, and G_n is the normal-state conductance of the junction [49]. The superconducting gap has a temperature

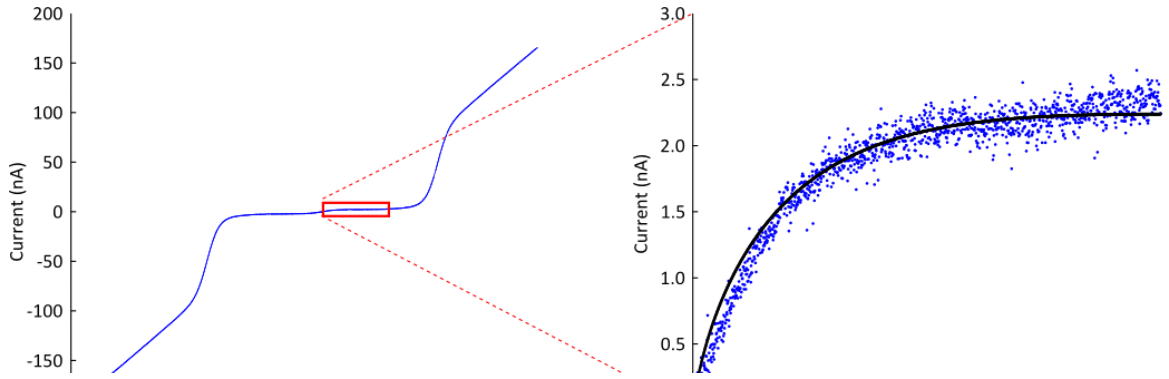

Figure 48 (Left) Example of transport measurement for sample e512 device B2. (Right) Zoom-in of boxed area in Left image, showing the subgap conductance in detail. The black line shows the fit function.

dependence, but at T = 4.2 Kelvin, it is reasonable to estimate the gap for niobium using the equation which holds at zero temperature; this has a 3% error. For films with a critical temperature of 9.2 Kelvin, this should result in a superconducting order parameter of 1.4 mV.

In Figure 48, I show an example of transport measured at 4.2 Kelvin. For this device, the supercurrent has been suppressed by applying a transverse magnetic field, which allowed for careful measurement of the subgap characteristics. The data was then averaged to reduce noise before being fit to the equation above. The normal conductance was obtained by measuring the current above 4mV and was 3.19 Siemens, and the only fit variable was the superconducting order parameter. In the example in Figure 48, the superconducting order parameter is 1.15 mV. Using the intercept method, I find the superconducting order parameter for this device to be 1.22 mV. Fitting the subgap region for eighteen devices from seven separate wafers results in an average superconducting order parameter is 1.2 mV ± 0.4 mV, which is in good agreement with other reported measurements [35].

The junction shown in Figure 48 had a 2.5 monolayer aluminum oxide barrier (approximately 6Å) which was grown at a rate of 41.0 seconds/monolayer barrier at a pressure of 5×10^{-6} Torr. Devices which were grown at faster rates and at lower pressures typically had excess subgap currents, which I attribute to oxygen vacancies in the barrier. The most common type of deviation that I observed was a slow increase in the subgap conductance that I interpret to be a high resistance short across the barrier. Although devices with this type of short have a Josephson critical current, the non-ideal subgap conductance is a key indicator that these devices would have a small coherence time if placed in a qubit configuration.

Analysis of voltage-current characteristics of Josephson junctions allows measurement of the BCS superconducting order parameter for the niobium superconducting electrodes. The temperature dependence of the order parameter is in good agreement with the BCS theory, which suggests that these electrodes should have low noise if placed into a qubit circuit. Subgap measurements for high-oxygen dose samples are also well-characterized by tunneling theory indicating no additional transport channels between the superconductors for these films. Samples with lower oxygen doses have high-resistance shorts that signal superfluous conduction paths in the barrier which may give rise to critical current noise. In the next section, I continue this analysis to show exactly how much oxygen is required for clean, low-noise Josephson junction barriers.

6.2.2 Josephson Junction Critical Current Measurements

In this section I will discuss what measurements of the Josephson critical current tell us about the nature of the materials used in my devices. I will discuss both the general effect of different oxidation doses on the critical current and the temperature dependence of the critical current for three devices grown under different oxidation regimes. I will finish by reporting critical current noise measurements comparing niobium epitaxial junctions with aluminum shadow-evaporated junctions.

The temperature dependence of the critical current in a Josephson junction was first modeled by Ambegaokar and Baratoff [61, 62] and has the functional form

$$I_c = \frac{\pi G_n}{2e}\,\Delta(T)\tanh\left(\frac{\Delta(T)}{2k_B T}\right)$$

At low temperatures, the hyperbolic tangent approaches 1, and the relation may be rewritten

$$J_c R_n A \approx \frac{\pi\Delta(0)}{2e}$$

Here J_c is the critical current density, R_n is the normal state resistance, and A is the area of the Josephson junction. As noted previously, the error in this approximation is around 3%. Inspection of this equation reveals that the right hand side is a constant for a given material, and the $R_n A$ product may be used as a proxy for the supercurrent. Since the critical current is sensitive to small magnetic fields but the normal state resistance is not, I use this proxy to compare the effect of oxidation doses from more than 50 devices.

In the course of this project, I grew wafers at multiple oxygen pressures and with different aluminum deposition rates. I have found that controlling both these parameters in addition to the barrier thickness is necessary in order to produce a desired critical current density. The slope of each plot series in Figure 49 is key to understanding the effect of different oxygen doses on the Josephson junction. There are three plot series with approximately equal slopes: the 30 second/monolayer 1.6x7 series, the 30 second/monolayer graded dose series, and the 40 second/monolayer 5x6 series. These samples all had a slow aluminum growth rate (.06 Å/s – .08 Å/s) with varying oxygen pressures. The growth rate determines the slope of the supercurrent density versus thickness plot, and the pressure determines the intercept. Lower pressures, naturally, require thicker barriers to produce the same critical current density. It is interesting to note that gradually increasing the oxygen (the graded dose series) is essentially equivalent to growing at 1.6×10^{-7} Torr pressure.

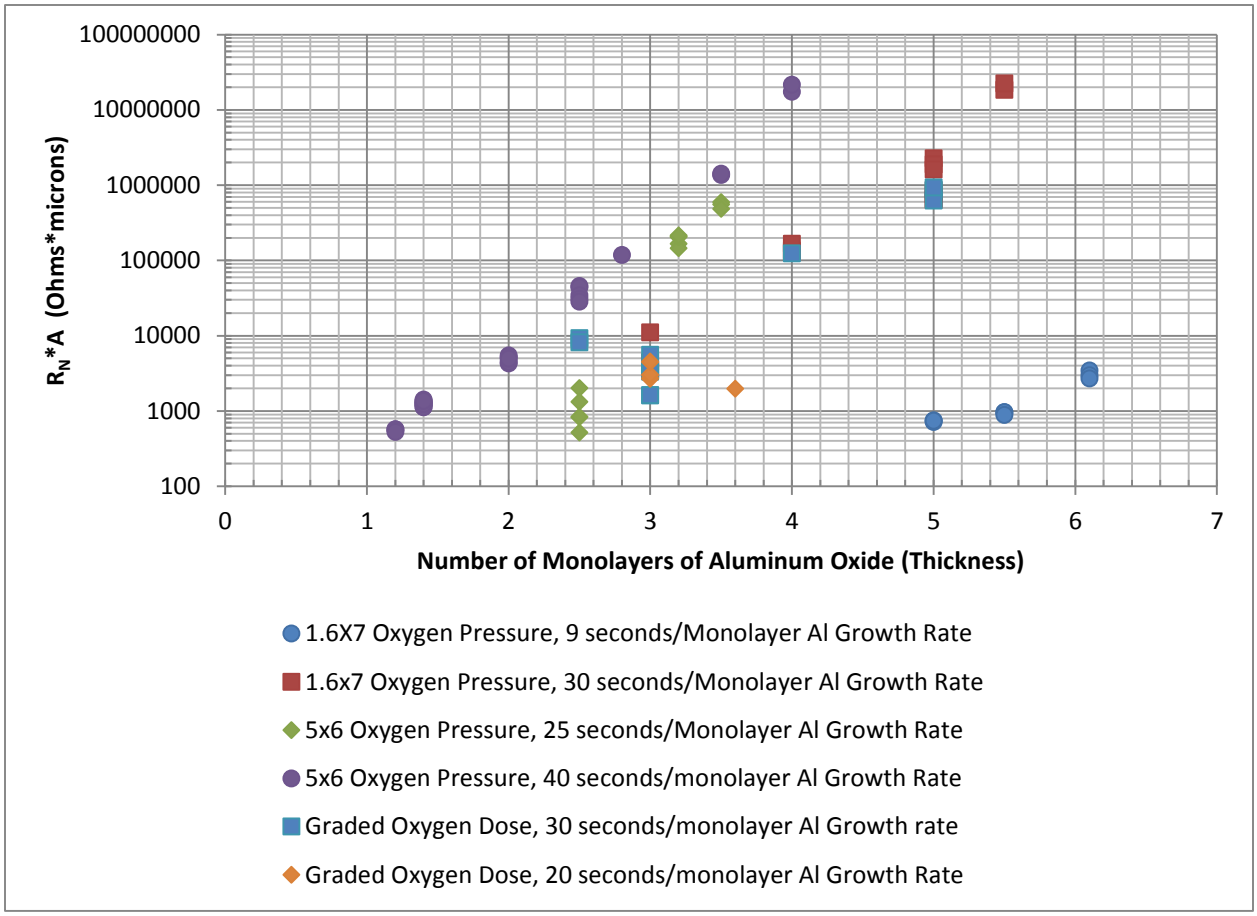

Figure 49 R_nA products for more than 50 devices with different oxygen doses. The horizontal axis plots the thickness of the barrier, while each plot series shows the different critical current scaling factors given different oxidation regimes. The notation "a x b" is a shorthand for $a \times 10^{-b}$ Torr. The "graded" oxygen dose refers to growth that started at an oxygen pressure of 5×10^{-7} Torr and ended at a pressure of 5×10^{-6} Torr.

The other three series were grow with a faster aluminum rate (.09 Å/s – .25 Å/s). The 20 second/monolayer graded dose series has nearly no dependence on thickness; this may simply be due to a lack of widely spaced data points. The 25 second/monolayer 5x6 series has considerable spread in the data at 2.5 monolayer thickness – if there was less spread it would be possible to determine if the slope for this series is truly steeper than for the slower growth rates. The final data series, 9 seconds/monolayer at 1.6x7 pressure, has a more gradual slope than any of the other data series. I believe this is because the aluminum oxide here is being not being fully oxidized before being buried under new material, which will lead to defects in qubits made from this material. Since the slopes do not vary with pressure for the samples grown at a slower rate, this indicates that the aluminum oxide can be fully oxidized before being buried beneath new aluminum. The difference in the y-intercept with pressure is likely due to oxygen diffusing into the buffer aluminum at higher pressures and increasing the effective barrier thickness.

Using the numerical fits for the superconducting order parameter, I have also examined the temperature dependence of the Josephson junctions' critical current. Figure 50 shows the critical current temperature dependence for the same three devices measured in Figure 47. Each device was normalized to its respective critical current measured at 0.33 Kelvin.

The superconducting order parameter for the e514 device is well characterized for a critical temperature of 8.77 Kelvin which is also the measured critical temperature of this device. However, the measured critical currents are not well fit by the Ambegaokar-Baratoff theory when the critical

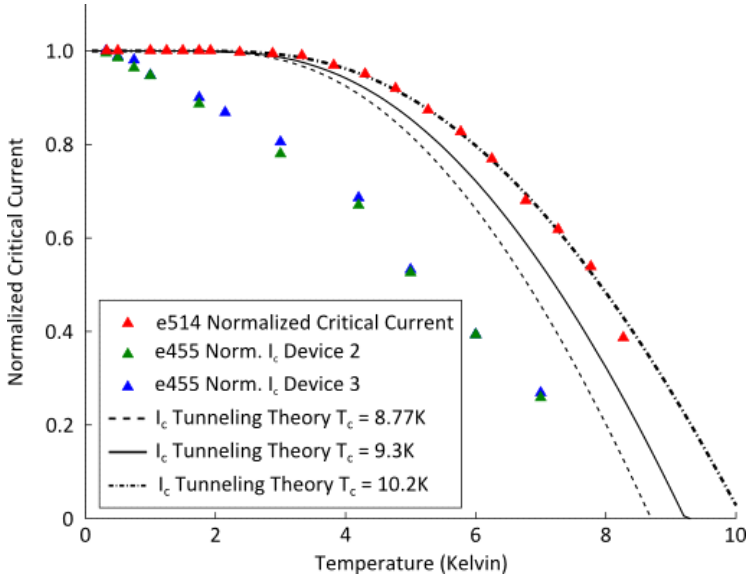

Figure 50 Experimental Josephson junction critical current temperature dependence. The dot-dash black lines shows the Ambegaokar-Baratoff tunneling theory assuming a critical current of 10.2 Kelvin. The solid black line shows the same theory for a film with a critical temperature of 9.3 Kelvin, and the dash black line shows the tunneling theory for a critical temperature of 8.77 Kelvin. The red triangles show the critical currents from film e514. The blue and green triangles show the critical currents from two devices on wafer e455.

temperature is 8.77 Kelvin. I found the best fit was with a critical temperature of 10.2 Kelvin with a mean percent difference of only 0.47%. In contrast, neither device from film e455 is well fit by the Ambegaokar-Baratoff theory; these devices appear to have a nearly linear dependence with temperature. To understand why the critical currents behave so differently, I will describe the differences in the fabrication of the two barriers.

It should be noted that wafer e455 was fabricated into Josephson junctions using the "via" process described in Section 5.2 whereas wafer e514 was fabricated into junctions using the self-aligned process described in Section 5.3. Although the junction fabrication process had an impact on device yield, I assert that it is irrelevant with respect to the critical currents and superconducting order parameters of

interest in this chapter. The most important factor affecting the critical currents and the superconducting order parameter is the fabrication of the film and barrier.

The primary difference between the two wafers is the thickness and oxidation dose of the barrier. Wafer e455 has a 6 monolayer thick aluminum oxide barrier (approximately 16Å) which was deposited at a rate of 20.6 seconds per monolayer of aluminum atoms in a pressure of 1.6×10^{-6} Torr of oxygen. In contrast, wafer e514 has a 2 monolayer thick aluminum oxide barrier (approximately 5Å) which was deposited at a rate of 40.8 seconds per monolayer of aluminum atoms in a pressure of 5×10^{-6} Torr of oxygen. Each sample was transferred to an oxidation chamber for an oxygen diffusion of an equivalent 10 Torr for 1 hour (see Section 3.6). Concisely, wafer e455 has a thick barrier that was deposited quickly, while wafer e514 has a thin barrier which was deposited slowly and at a higher pressure.

I propose that the explanation for the linear temperature dependence in e455 is due to spins in the barrier caused by incomplete oxidation of the aluminum. When there is sufficient oxygen to form insulating AlO_x but not stoichiometric Al_2O_3, oxygen vacancies leave pinned electrons in the barrier. The dipole moment associated with these electrons essentially becomes a magnetic impurity in the barrier. While quasiparticle transport is predominantly insensitive to magnetic fields, the Josephson supercurrent is sensitive to magnetic impurities [63].

The Josephson supercurrent is the tunneling of Cooper pairs across the insulating barrier [64]. Since Cooper pairs are a spin-up/spin-down pair, a spin flip of one of the electrons will break the Cooper pair, reducing the supercurrent across the junction. Paramagnetic impurities in the either the superconductor or the insulating barrier will cause some of these electrons to scatter and spin-flip, breaking the Cooper pairs and reducing the critical current [56].

For wafer e455, I observed a reduced critical current in comparison with the Ambegaokar-Baratoff theory. Reduced critical currents with similar linear temperature dependencies have been modeled by others [63, 65], suggesting that dipole moments in the barrier is a reasonable physical model for the reduced supercurrent for these two devices.

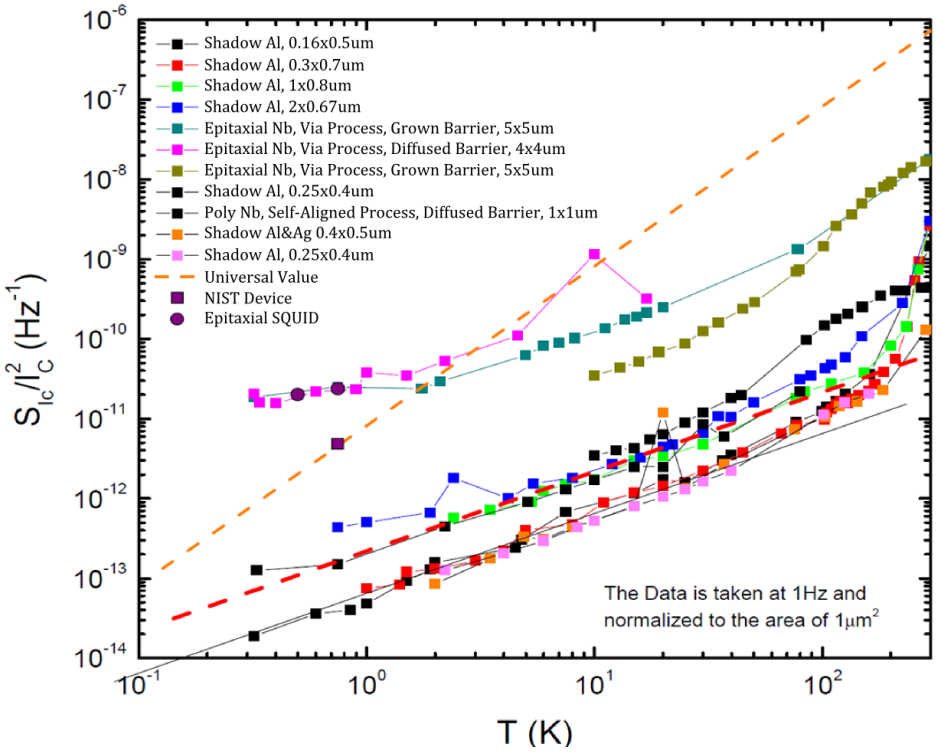

Figure 51 Critical current noise measurements of trilayer niobium junctions and shadow-evaporated aluminum junctions. Figure courtesy D. Van Harlingen.

The effect of oxygen vacancies can also be observed in measurements of the critical current noise. Working with students in Dale Van Harlingen's group, fluctuations in the Josephson critical current were measured for four trilayer samples fabricated by me and compared with aluminum-aluminum oxide-aluminum shadow-evaporated junctions. The process of shadow-evaporation is commonly used to fabricate aluminum superconducting qubits and is described in Appendix A. Let it suffice here that this fabrication technique permits much smaller junctions than what I was able to fabricate from trilayer structures.

The data in Figure 51 shows that the four niobium devices have more critical current noise than any of the aluminum junctions. Among the four niobium devices, the epitaxial niobium film with the diffused barrier (no oxygen was present in System E during the growth) has the most noise, which might suggest that better oxidation conditions produce less noise. However, this seems to be immediately contradicted by the polycrystalline niobium film, which has the lowest noise amongst the niobium films

and also a diffused barrier. However, the fabrication process between the epitaxial and polycrystalline devices also varied – the polycrystalline niobium film was made using the self-aligned process and was subsequently much smaller than the epitaxial films. In fact, although the noise measurements were normalized with respect to size, the smallest junctions consistently have lower noise than the larger ones. If the noise were simply due to some uniform density of defects, there should not be such a size dependence, leading some to speculate that strain may be causing critical current noise in these devices [18].

The critical current temperature dependence shows how oxidation dose affects the barrier characteristics. The suppression of the critical current for the sample grown with a lower oxygen pressure suggests that oxygen vacancies are leading to magnetic impurities within the barrier. Differences in the slope of the R_nA product also implies that low oxygen dose leaves oxygen vacancies in the Josephson junction barrier. These measurements indicate that oxygen vacancies may give rise to critical current noise in the qubits. To reduce critical current noise, the growth pressure should be in the range of 5×10^{-6} Torr and the aluminum flux should be no less than 40 seconds/monolayer.

Measurement of critical current noise itself also suggests that improvements in the barrier composition reduces the critical current noise from these junctions. A comparison of *only* the epitaxial niobium films shows that grown barriers have less noise than the diffused barrier. However, size of the Josephson junction plays a large role in the critical current noise measurements, in spite of area normalization. A true comparison of critical current noise measurements between niobium and aluminum junctions will require further improvements in fabrication techniques to make 100 nm x 100 nm niobium junctions.

6.3 Flux Noise Measurements of Epitaxial SQUIDs

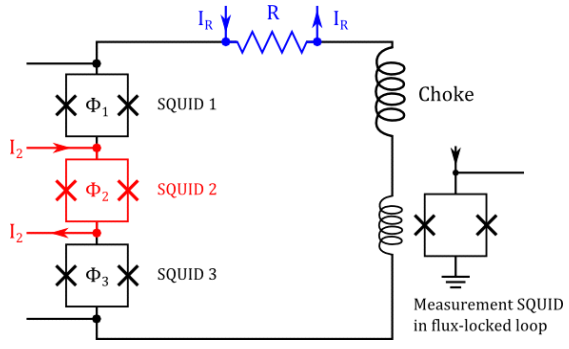

Figure 52 Flux noise measurement setup. Multiple SQUIDs were fabricated simultaneously on each chip. Only three SQUIDs are shown; in reality 6-8 SQUIDs were on-chip. Figure courtesy J. Clarke, UCB.

In this section, I will describe the measurements of flux noise measured in dc-SQUIDs. In Chapter 4, I demonstrated that my niobium films are grown epitaxially on the sapphire substrate, which should minimize the effects of MIGS as a noise source. I compare all-epitaxial SQUIDs to polycrystalline SQUIDs fabricated at NIST. Additionally, SQUIDs with epitaxial loops and shadow-evaporated junctions were fabricated and measured. The flux noise

measured from these SQUIDs was a new record low; this can be attributed to the use of epitaxial niobium in the SQUID loop construction.

To evaluate the effect of using single-crystal niobium, I fabricated dc-SQUIDs which are extremely sensitive to changes in magnetic fields. This sensitivity can be seen following the derivation in Tinkham [50]. For SQUIDs where the critical currents of the two junctions are equal, the dc average voltage across a SQUID is approximately

$$V = \frac{R}{2}\sqrt{I^2 - \left(2I_c \cos\left(\frac{\pi\Phi_{ex}}{\Phi_0}\right)\right)^2}$$

Here, R is the resistance of the shunt, and Φ_{ex} is the external flux. Biasing the current at $I = 2I_c$ produces the largest voltage swings for different fluxes through the loop, making it possible to measure small changes in the flux through the loop.

My SQUIDs were measured by John Clarke's group at the University of California Berkeley (UCB). The measurement set-up is shown in Figure 52. The current flowing through SQUID 2 biases it at a non-zero voltage and the current running through the resistor nulls any current flowing through the measurement circuit. Flux noise in SQUID 2 produces a current noise in the measurement SQUID which is biased in a flux-locked loop (see Reference [20] for an explanation of flux-locking a SQUID).

The figure of merit for evaluating flux noise is the spectral density $S_\Phi(f)$ of the magnetic flux noise. The spectral density goes as $1/f^\alpha$ with respect to frequency, hence the name "1/f" noise. The 1/f noise behavior has been observed for many materials and a multitude of device geometries with only a spread of a factor of two [19]. This provides a reference point for the SQUIDs that I fabricated from epitaxial

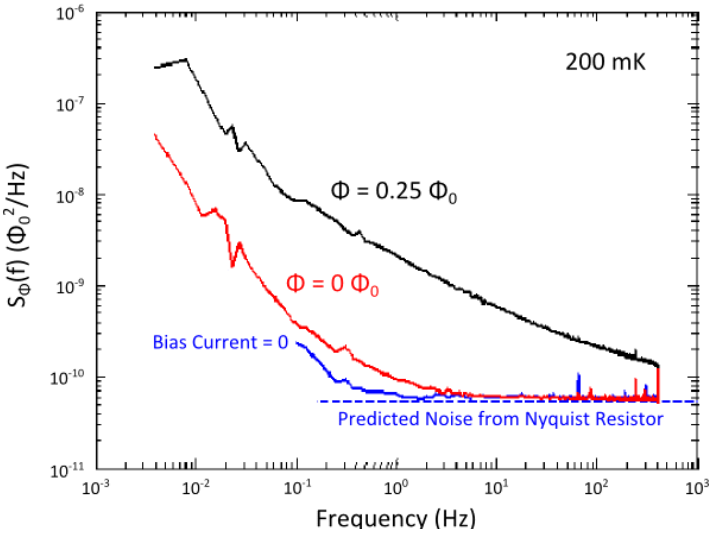

Figure 53 Spectral noise density for a typical epitaxial SQUID measured by John Clarke's group at the University of California Berkeley, measured at 200 milli-Kelvin and in external fluxes of 0 flux quantum and ¼ flux quantum. The lines in blue show noise from a Nyquist resistor (the noise floor). Figure courtesy J. Clarke, UCB.

niobium. A typical spectra for one of my MBE-grown SQUIDs is shown in Figure 53. In Reference [19], the lowest value of flux noise at 1 Hz ($S_\Phi(1\ Hz)$) is approximately $16\ \Phi_0^2/Hz$; the lowest value of flux noise in my SQUIDs is approximately $1.6\ \Phi_0^2/Hz$, an order of magnitude lower.

John Clarke's group at UCB also compared the SQUIDs that I made to polycrystalline niobium SQUIDs made at NIST. My SQUIDs all had a uniform geometry, whereas the NIST SQUIDs varied with linewidth and size. This difference in size likely accounts for some of the spread in the spectral density measurements for the NIST SQUIDs. The list of the different sizes and the geometry is in Table 3 and range from "fat" loops to "skinny" ones. My devices are at the middle of the range.

The flux noise spectral density was measured at temperatures from 30 milli-Kelvin up to 4.2 Kelvin for the six NIST polycrystalline SQUIDs as well as my epitaxial SQUIDs. For temperatures below 1 Kelvin, the epitaxial SQUIDs were quieter than any of the NIST SQUIDs, and for temperatures up to 1.7 Kelvin, the epitaxial SQUIDs were in the middle of the range again.

SQUID #	D (μM)	W (μM)	D/W	
1	370	160	2.31	"Fat"
2	210	80	2.63	
3	130	40	3.25	
UIUC	110	30	3.66	
4	90	20	4.50	
5	65	10	6.50	
6	50	5	10.00	"Skinny"

Table 3 List of the sizes (D) and linewidths (W) for the NIST SQUIDs (1-6) and my devices (labeled UIUC).

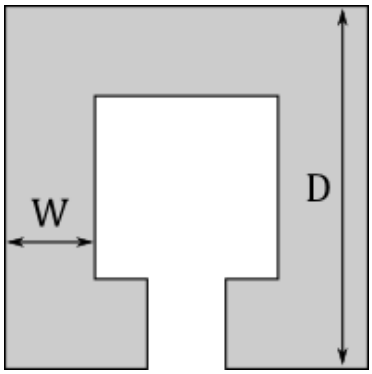

Figure 54 Cartoon of the SQUID loop identifying D and W (junctions not shown).

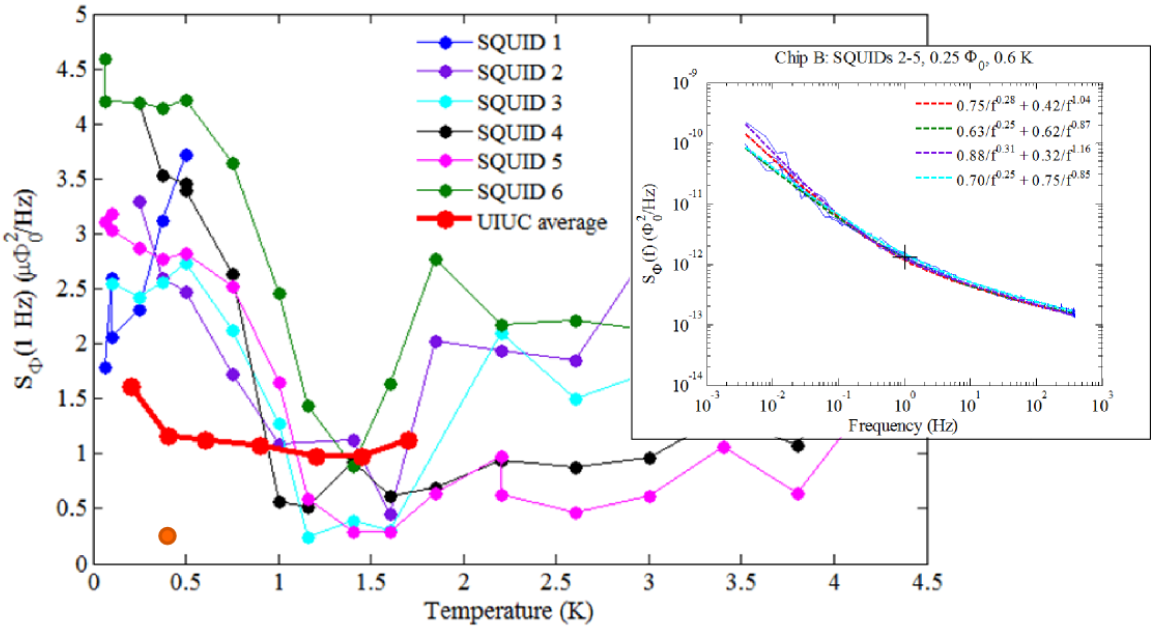

Figure 55 Flux noise spectral density measured at 1 Hz for the 7 devices listed above. Orange dot shows the spectral flux noise for the SQUIDs fabricated with shadow-evaporated junctions, discussed below. The inset shows the spectral noise density frequency dependence at 600mK for four of the SQUIDs averaged in the temperature dependent measurement. Note the consistency between all four devices; the 1Hz point on for this temperature is marked by the cross. The figures are courtesy J. Clarke, UCB.

This work is not primarily concerned with understanding the noise in the NIST devices or why there is so much spread in the spectral noise density. The key result here is that, regardless of geometry, the epitaxial SQUIDs have the lowest noise at the temperatures where qubits are operated. This indicates

that the epitaxial nature of the materials is reducing the flux noise in these SQUIDs, which supports the notion that MIGS are causing noise at the niobium-sapphire substrate. Furthermore, critical current noise was also measured in these SQUIDs. The measured value of the critical current noise was $300\ (pA)^2/Hz$. This value is less than the expected value of $2500\ (pA)^2/Hz$ for Josephson junctions of this size; the expected value is calculated from the formula for the universal noise found in Reference [11]:

$$S_{I_c}(1\ Hz, 4.2\ K) \approx 144\ \frac{(I_c/\mu A)^2}{A/\mu m^2}\frac{pA^2}{Hz}$$

Both the low critical current noise and low flux noise values suggest that the epitaxial structures used to make these SQUIDs are reducing noise in the junctions, most likely due to reducing MIGS caused by lattice mismatch at the interfacial boundary between the sapphire and niobium.

In addition to the all-epitaxial MBE SQUIDs, I also worked with Chris Nugroho at UIUC to produce epitaxial SQUID loops with shadow-evaporated junctions. Since the trilayer niobium junctions were larger than any of the shadow-evaporated aluminum junctions, it was possible that any noise due to lattice defects was masked by noise due to the large junctions. Incorporating shadow-evaporated aluminum junctions into the design allowed for much smaller junctions and which would make flux noise due to the SQUID loop to be the dominant noise source in the device. At the same time, *in-situ* gold was deposited on the film to prevent the top surface of the niobium from oxidizing when exposed to air. This allowed me to evaluate both the epitaxial niobium SQUIDs and to test whether gold cladding layers would reduce noise in these devices. I grew a niobium film capped with 100Å of gold and etched the film into SQUID loops. Three of the SQUID loops were covered with photoresist and the gold was then removed from the rest of the film. At this point the aluminum junctions were deposited. A micrograph of these SQUIDs is shown in Figure 56.

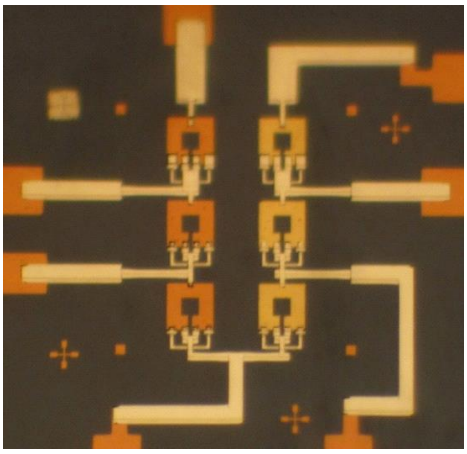

Figure 56 Micrograph of niobium SQUID loops with shadow evaporated aluminum junctions. The SQUIDs on the right side of the image also have a thin gold layer (deposited in-situ) to try to determine if a cladding layer would reduce noise.

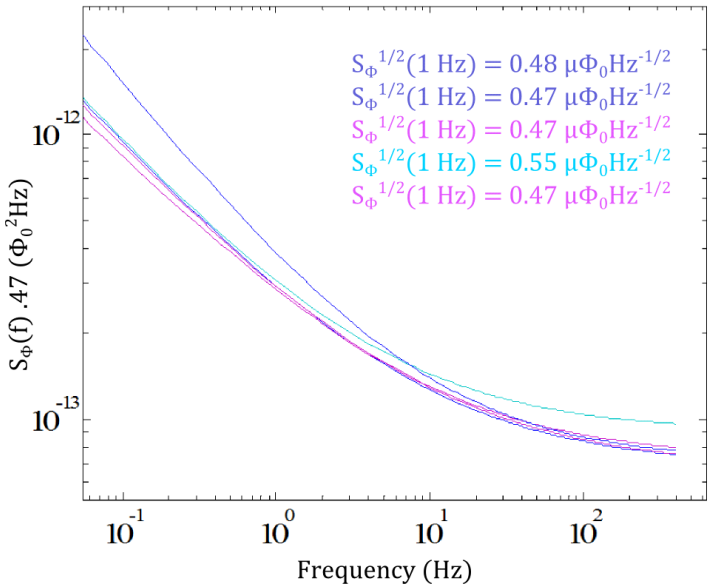

Figure 57 Flux noise spectral density for SQUIDs fabricated with shadow-evaporated junctions. The RMS excess noise for each SQUID at 1 Hz is listed in the inset for better comparison with Reference [19]. Figure courtesy J. Clarke, UCB.

The magnetic flux noise spectral density is shown in Figure 57. The noise measurements here represent a new record low flux noise measurement in dc-SQUIDs. This record low value for the spectral noise density is due to the use of epitaxial materials in the SQUID loop, validating the view that MIGS cause flux noise in polycrystalline films. However, no difference is observed between SQUIDs coated with gold and SQUIDs without the gold cladding. The gold should prevent spins due to oxidation on the surface of half the measured SQUIDs. Since no effect of the gold is observed, spins on the surface of the niobium loop due to oxidization (as seen in Reference [22]) are not the primary source of noise in these devices.

This section has used the measurements of flux noise in dc-SQUIDs to identify the best parameters for growth of low-noise flux qubits, which are sensitive to changes in the magnetic field. The all-epitaxial niobium SQUIDs have less noise than comparable polycrystalline SQUIDs at the very lowest temperatures, and epitaxial SQUID loops with shadow-evaporated junctions have the lowest flux noise ever measured. The decrease in flux noise is likely due to reductions of crystal defects which give rise to MIGS. In contrast with these improvements, surface cladding to prevent oxidation at the top electrode does not show any appreciable decrease in flux noise, suggesting that surface spins are not (yet) a dominant source of flux noise.

Chapter 7 Sample Cleanliness

In Section 6.1 I described the physics of resonators and how these simple devices can be used to identify losses in materials. In particular, work by Wenner [46] suggests that losses in co-planar waveguides are primarily due to defects at the substrate-air interface with the metal-substrate and metal-air interfaces of decreasing importance. One of the ways in which defects may affect these surfaces is by fabrication residue being left over after sample fabrication. In Figure 58 this is hypothesized to be photoresist residue remaining on the metal and substrate surfaces. The dielectric coefficient of the residue is different than the air/vacuum surrounding it and will have a different response to the alternating electromagnetic fields as a result. The residue may be a source of two-level fluctuators, or may simply be a dissipative mechanism in the cavity in which the qubit is operated. In either case, removal of any fabrication residue should result in longer coherence times.

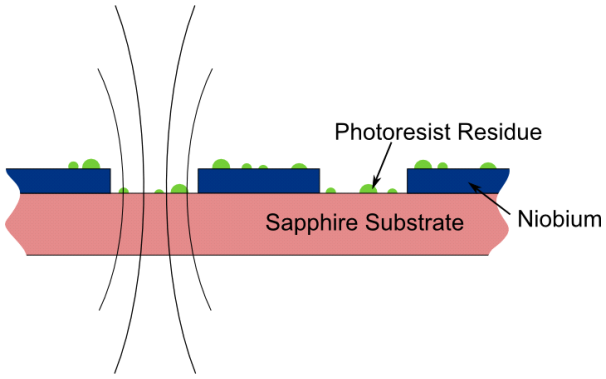

Figure 58 Side view of a coplanar waveguide resonator, showing the electromagnetic field lines (black lines) and nonconductive photoresist residue on the surfaces of the wafer.

In this chapter, I will provide AFM images of processed and unprocessed films, showing both the sapphire and niobium surfaces. The processed films are simply the result of fabricating coplanar waveguides; AFM images of these surfaces reveal photoresist residue decorating both the sapphire and niobium films. To remove this residue, I devised a series of etching tests. In the following sections I will describe the various etching methods used to remove the photoresist and I will show AFM images of both the sapphire and niobium surfaces following each etching attempt. I will conclude with the etch recipe that leaves both surfaces in the most pristine condition.

7.1 AFM of Unprocessed and Processed Films

I can only make an AFM measurement of the clean sapphire surface by scanning the wafer before growth of the niobium film. Scans of the sapphire surface, as in Figure 59, show very flat films with atomic step-edges. I will show that I can return to such a surface with an appropriate etch recipe. To measure the clean niobium surface, I did an AFM scan of the surface immediately following the growth of the film. The surface RMS roughness of the niobium surface shown in Figure 59 is 300 pm.

After scanning the initial, clean niobium surface, I fabricated the niobium films into transmission line

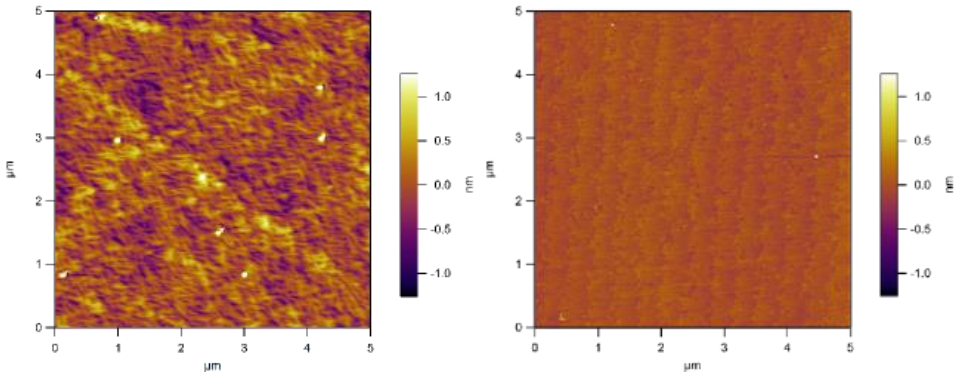

Figure 59 (Left) AFM image of the niobium surface of wafer e542 immediately after growth, before any processing occurs. (Right) AFM image of a typical sapphire surface. Atomic step-edges are clearly visible.

resonators. This was done by spinning a layer of PMMA and then exposing the wafer in the deep UV mask aligner. After development, the sample was etched in the March RIE system using the standard SF_6 niobium etch recipe. The wafer was diced into 11 separate pieces, each featuring at least one resonator device. Because sapphire is such a hard material, the dicing had to be done with a diamond-edged saw blade. Furthermore, the dicing saw produced a large amount of dust and debris that coated the wafer. To preserve the devices, a protective layer of MMA and AZ 5214 was spun onto the sample before it was diced. The photoresist was stripped by immersing the devices into acetone, but also by applying some sonication to the device. The sonication helped ensure that the debris did not settle back down onto the wafer, leaving it free of visible dust after the dicing.

While visible debris can be removed by sonication in an acetone bath, a residue still remains on the film surface. AFM images of both the niobium and sapphire surfaces reveal the extent of this residue. Whereas the RMS roughness of the niobium surface was 300 pm before processing, the RMS roughness of the surface is 645 pm after the processing steps and acetone strip. The RMS roughness of the

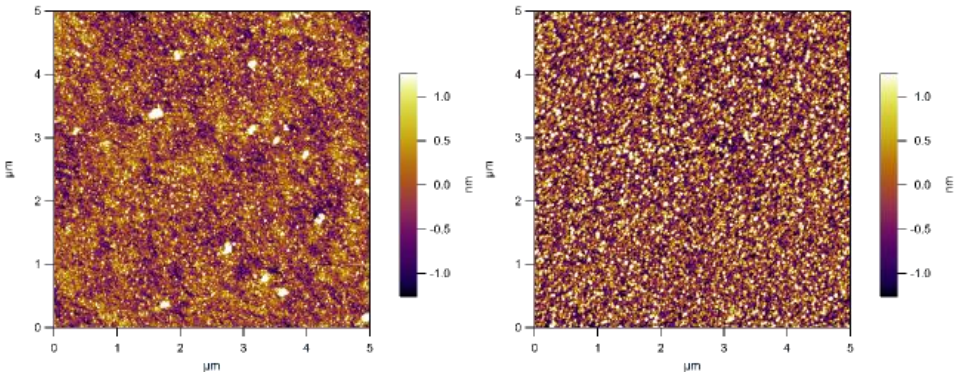

Figure 60 (Left) AFM of the niobium surface after device fabrication. (Right) AFM of the sapphire surface after device fabrication.

sapphire surface is even larger: 740 pm over the 5 micron scan area. In the next section I will describe the different attempts made to remove this residue and the resulting images from the various attempts.

7.2 Cleaning Etches and Resulting Microscopy

To determine a method of cleaning the niobium film, wafer e542 was diced into 11 separate pieces. I devised 9 different chemical and reactive ion etches to try to remove the photoresist residue from the niobium and sapphire surface. After each etch, I performed microscopy on both the niobium and sapphire surfaces. In this section I will first list the various etches and the resulting RMS roughness values for both relevant surfaces. I will then provide the details of each individual etch and the resulting AFM images.

FILM STATE OR ETCHANT USED	NIOBIUM SURFACE RMS ROUGHNESS (PM)	SAPPHIRE SURFACE RMS ROUGHNESS (PM)
UNPROCESSED FILM (INITIAL STATE)	300	90
PROCESSED FILM (ACETONE ONLY)	645	745
ACETIC ACID	695	530
BUFFERED OXIDE ETCHANT + PEROXIDE	560	1100
OXYGEN ASH + 10 SECOND BOE	340	100
OXYGEN ASH + 60 SECOND BOE	225 – 4000	75 – 120
OXYGEN ASH + SF_6 RIE ETCH	490	800
"PIRANHA" ETCHANT	1000	1250
"RCA" CLEAN	525	735
TRANSENE ALUMINUM TYPE D + PEROXIDE	850	700
TRICHLOROETHLYENE SOLVENT BATH	540	380

Table 4 List of the RMS roughness values for both the niobium and sapphire surfaces for various etches. The green color is used to emphasize the measurements of the initial film, while the orange color is used to denote the roughness measurements of the film after processing using only acetone to strip the photoresist.

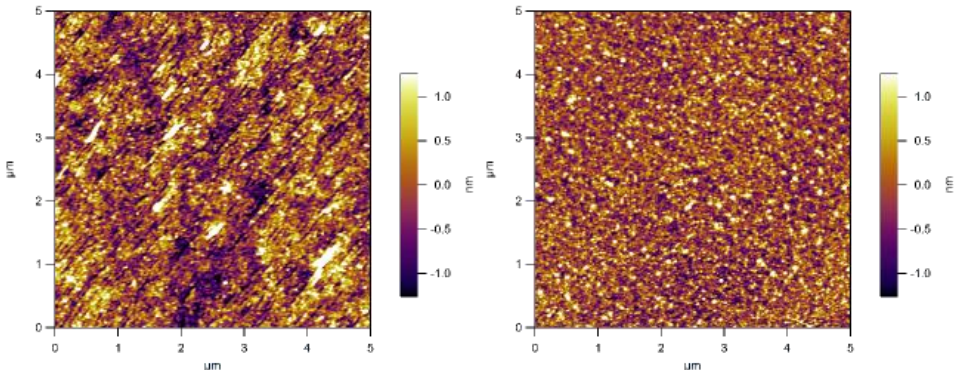

Figure 61 (Left) AFM of the niobium surface after the acetic acid etch. (Right) AFM of the sapphire surface after the acetic acid etch.

A literature review of photoresist strippers revealed that acetic acid can be used to remove photoresist in some semiconductor applications [66]. For this test the sample was placed in a room temperature solution of acetic acid for 60 seconds. The acetic acid did appear to remove some of the photoresist residue, but also appeared to attack the niobium surface (see Figure 61). Given that the acetic acid was only marginally removing the residue and was in fact damaging the niobium surface, I moved on to other etches.

After the acetic acid etch, I tried a mixture of 1:10 hydrogen peroxide to Buffered Oxide Etchant (BOE). The BOE was itself a mixture of 1:7 hydrofluoric acid (HF) and ammonium fluoride (NH_4F). The sample was held in the solution for 60 seconds before being thoroughly rinsed with water. Upon imaging the surfaces, I observed that there were deep pits etched into the niobium surface (see Figure 62).

These pits did not have the typical characteristic of pits caused by physically scratching the surface, and were found in multiple scans across the niobium surface. Moreover, the sapphire surface was left significantly rougher than its prior state with only an acetone sonication. Because of both of these downsides, a peroxide-BOE mixture is not a suitable cleaning procedure for niobium films.

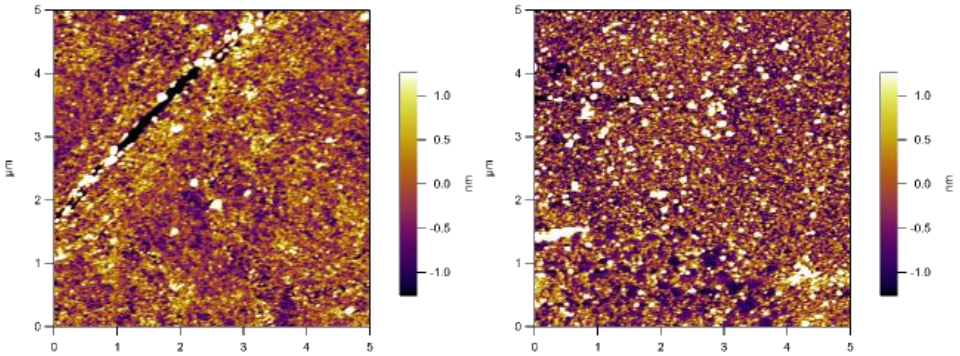

Figure 62 (Left) AFM of the niobium surface after the BOE + peroxide etch. (Right) AFM of the sapphire surface after the BOE + peroxide etch.

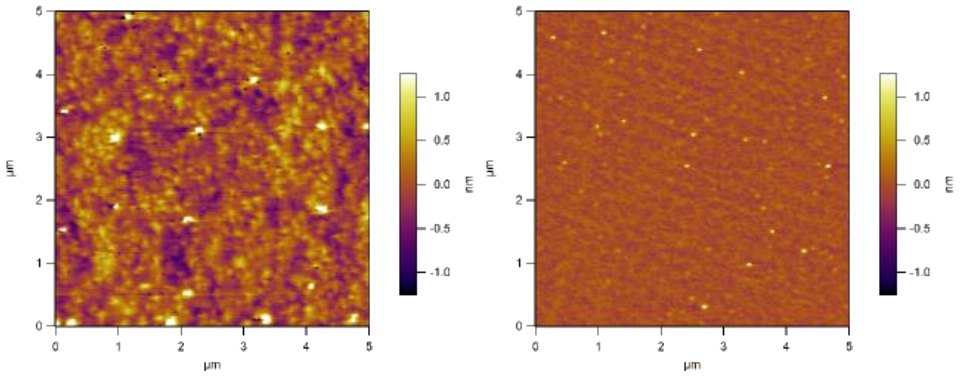

Figure 63 (Left) AFM of the niobium surface after the oxygen ash and 10 second BOE etch. Comparing this image with Figure 59 shows the original surface structure with some additional roughening. (Right) AFM of the niobium surface after the same etch. Atomic step-edges are again visible.

I also used a two-step cleaning process on some of the pieces of e542. This involved using an oxygen plasma to attack the photoresist residue and then a subsequent etch to remove any native oxide the oxygen ash created. The sample was placed in the March RIE and subjected to a 200mTorr oxygen atmosphere. The RF power was 175 Watts and the sample was etched for 20 seconds. Following the oxygen ash, I then immersed the sample into a BOE bath for 10 seconds. Microscopy of the niobium surface is shown in Figure 63. When compared to the unprocessed niobium surface shown in Figure 59, the two surfaces are seen to be very similar, with perhaps more point defects or debris remaining in the cleaned surface. Similarly, there is some residue still remaining on the sapphire surface, but atomic scale terraces are again visible. This etch showed a fair amount of progress, and the next images will show the results of leaving a sample in the BOE bath for 60 seconds.

I actually performed the oxygen ash and 60 second BOE etch on two pieces of wafer e542. In both cases, the final film surface shows signs of insufficient rinsing. On the first piece, the niobium surface has etch

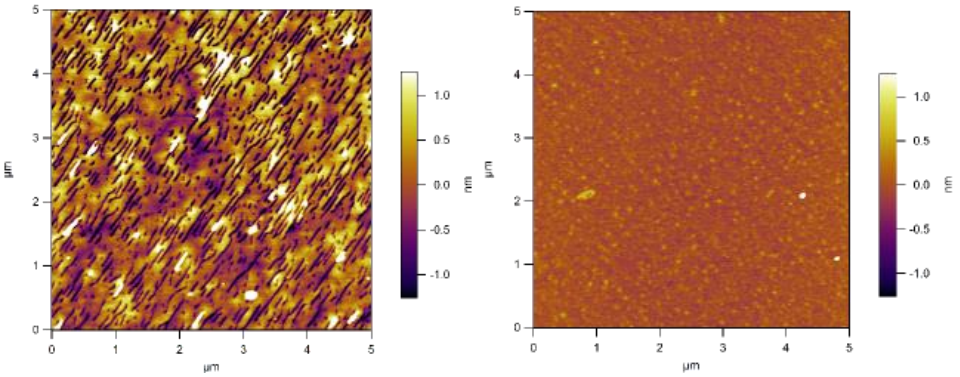

Figure 64 (Left) AFM of the niobium surface after the oxygen ash and 60 second BOE etch showing etch streaks. (Right) AFM of the sapphire surface after the same etch; some atomic step-edges are visible.

streaks across the film. A close analysis of the microscopy indicates that without the etch streaks the sample would be in pristine condition. The sapphire surface is even more promising; with the exception of a few spots, the surface is flat with possible terracing along the diagonal. A second sample had the same etch process done to try to resolve the issue with the etch streaks. In this case (Figure 65), I observe large, round spots covering the niobium surface. These spots are 40 – 80 nm tall, and appear to be spherical in shape. I believe these are "water spots" and are due to incomplete rinsing and drying of the sample during the cleaning process. A roughness analysis of the entire scanned image results in a RMS roughness value around 4000 pm. Removing the water spots from the analysis produces a RMS roughness value around 225 pm – a lower value than the original unprocessed film. The water spots do not appear on the sapphire surface, which clearly shows terracing.

The combination of an oxygen ash and BOE etch is clearly a potential candidate for cleaning niobium and sapphire surfaces. However, it is also evident that the procedure must be carefully carried out to ensure

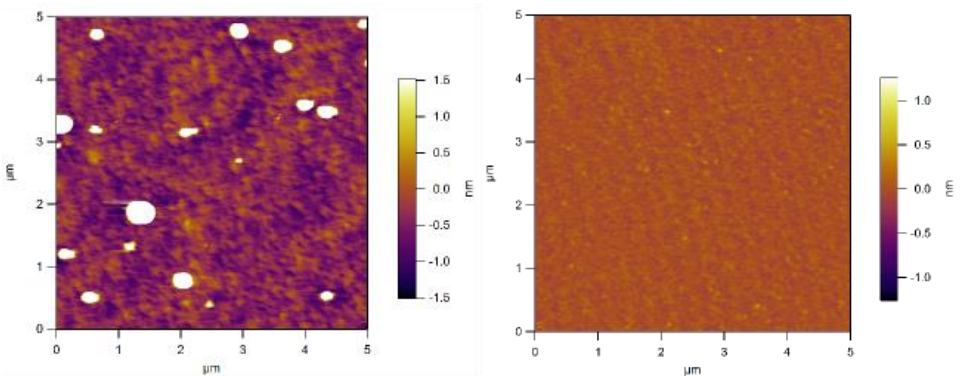

Figure 65 (Left) AFM of the niobium surface after the oxygen ash and 60 second BOE etch for a second sample. Water spots dot the surface of the sample. Note that the height scale for this image is different from all other images in this chapter. (Right) AFM of the sapphire surface after the same etch. Atomic step-edges are clear here; no water spots are visible on the sapphire.

success. Most importantly, the images above indicate that very thorough rinsing is required to provide a pristine niobium surface.

The BOE etch was used to remove any oxide formed on the niobium film by the oxygen ash. Another way to remove this oxide was to use a SF_6 reactive ion etch on the wafer. Since the sapphire surface should already be fully oxidized, the only oxide to be removed would be on the niobium surface, so using the niobium RIE etch recipe was a reasonable test. As before, the oxygen ash used a pressure of 200mTorr, 175 Watts of RF power, and lasted for 20 seconds. The SF_6 etch also used a pressure of

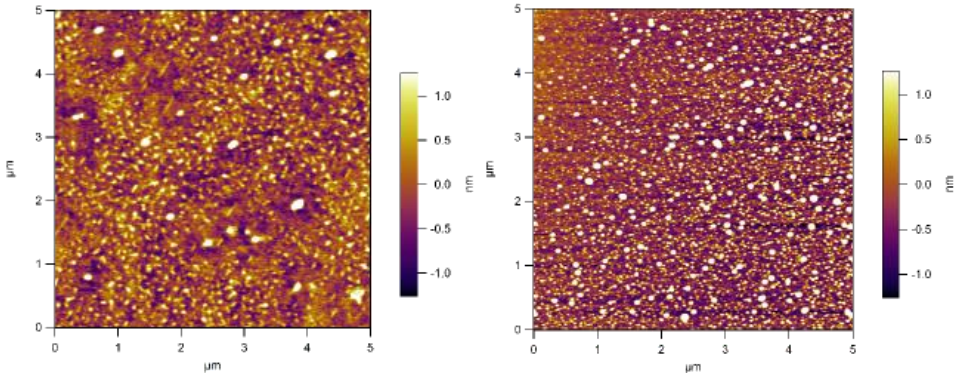

Figure 66 (Left) AFM of the niobium surface after the oxygen ash and SF_6 RIE etch. (Right) AFM of the sapphire surface after the oxygen ash and SF_6 RIE etch.

200mTorr, but here had an RF power of 35 Watts and lasted for 30 seconds. The film height after the RIE etch was 250 nm; 50 nm had been removed from the initial film. However, neither the niobium nor the sapphire surface appear to be in pristine condition. Using a SF_6 RIE etch clearly will not result in clean resonant surfaces.

Another etchant that is frequently used in semiconductor work to remove organic residues is “Piranha”. This name generally refers to mixtures of peroxide and sulfuric acid; in my case, I used a 10:1 mixture of H_2SO_4 and H_2O_2. Mixing the chemicals produces a highly exothermic reaction – after allowing the

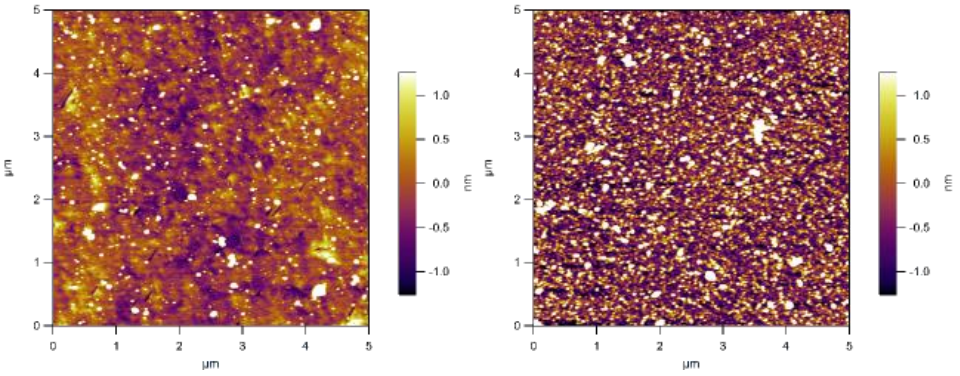

Figure 67 (Left) AFM of the niobium surface after the Piranha etch. (Right) AFM of the sapphire surface after the same etch.

solution to cool, I immersed my sample in the etchant for 60 seconds. The resulting surfaces have the largest roughness values of all the tested films, with the exception of the water spots in Figure 65. The sapphire surface, in particular, is quite rough, suggesting that the Piranha etch is possibly damaging the underlying sapphire. As a result, while Piranha may be an excellent for silicon wafers, it is not appropriate to clean sapphire and niobium films.

Another technique used to remove organics in the semiconductor industry is known as the "RCA clean", as it was initially developed at the Radio Corporation of America. The process uses a mixture of 1:1:5

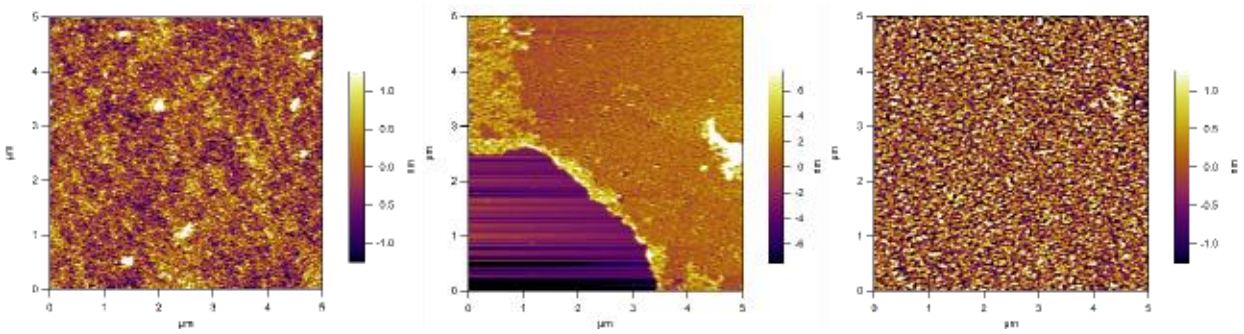

Figure 68 (Left) AFM of the niobium surface after the RCA clean. (Middle) AFM of the residue on the sapphire substrate. Height scale is +/- 7 nm for this image. (Right) AFM of the residue atop the sapphire substrate.

ammonium hydroxide (NH_4OH), hydrogen peroxide (H_2O_2) and water to remove organic contaminants. On silicon wafers, this solution leaves behind a thin oxide layer, which is removed by BOE. For my sample, it was immersed in the ammonium hydroxide solution for five minutes.

After thorough rinsing, it was etched in BOE for 35 seconds. The cleaning process did not leave the niobium in pristine condition, but the process did not seem to damage the film. Unfortunately, this process appeared to leave behind a layer on top of the sapphire. An image of this layer is shown in Figure 68. The height scale shows that this layer is clearly atop the substrate. A scan of just this layer is also shown; this is the image analyzed to provide the RMS roughness. As with the Piranha etch, the RCA cleaning process does not appear to be suitable for cleaning sapphire and niobium films.

Another etch that I used was a variation on a wet etch that I commonly use to fabricate trilayer junctions. I use Transene Aluminum Etchant Type D to remove aluminum when creating Josephson junctions. This solution is a mixture of phosphoric and acetic acids, as well as a proprietary nitrogenizing component. To remove the photoresist, I mixed a solution of 10:1 Transene and hydrogen peroxide and immersed the sample in the etchant for 60 seconds. The resulting niobium surface is rougher than the

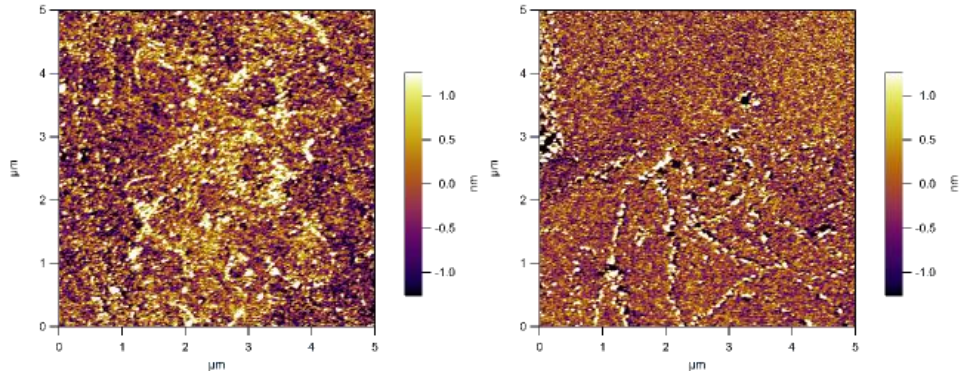

Figure 69 (Left) AFM of the niobium surface after the Transene etch. (Right) AFM of the sapphire surface after the same etch.

acetone-only cleaned sample. The sapphire surface has some type of debris decorating the surface. Since the Transene etchant is intended to remove aluminum and aluminum oxide, it is perhaps not surprising that this mixture has damaged the substrate surface. Both the niobium and sapphire surfaces indicate that this Transene-peroxide mixture is not an effective cleaning tool for these samples.

The final solution that I used to clean the surface was not an acid, but rather the solvent trichloroethylene (TCE). I regularly used a TCE bath and sonication to remove organics from my samples before loading them into the MBE vacuum system. For these tests, I immersed the sample in 30mL of TCE and applied sonication for 60 seconds. The niobium surface is cleaner than the acetone-only cleaned sample. The sapphire surface, too, has some of the organics removed, though there aren't any visible atomic-scale terraces. This cleaning process shows the best results, next to the oxygen ash and

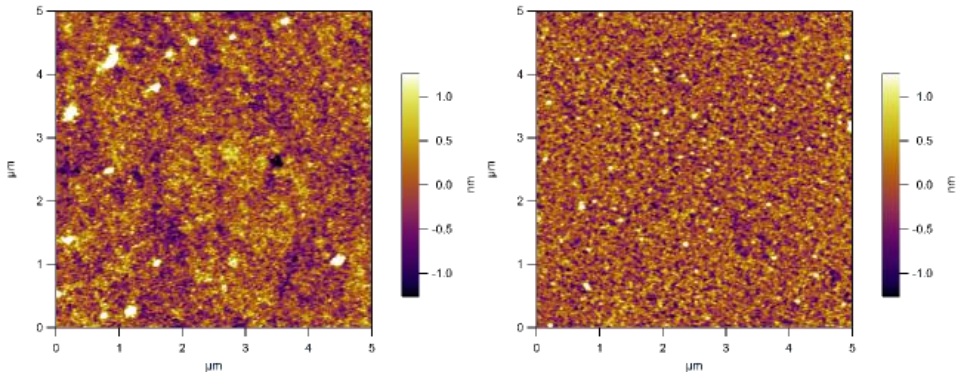

Figure 70 (Left) AFM of the niobium surface after the TCE solvent clean. (Right) AFM of the sapphire surface after the TCE solvent clean.

BOE combination. It is possible that a longer sonication time would improve the results, but it is also possible that sonication can damage the devices on the wafer.

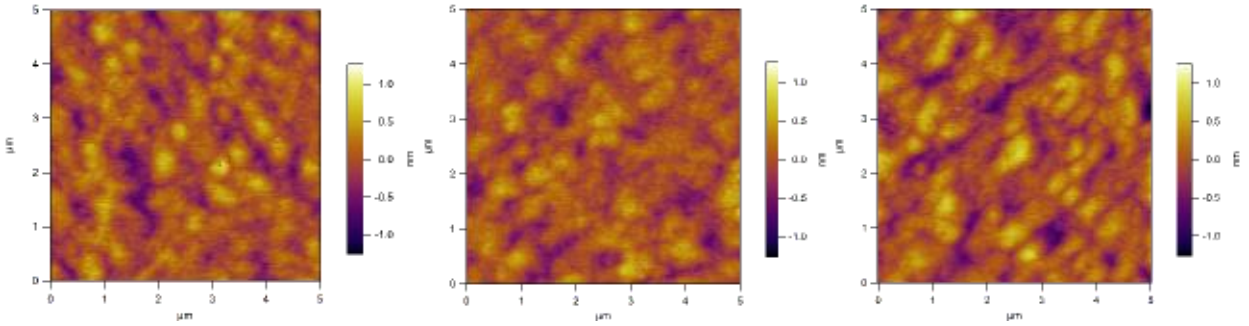

Figure 71 (Left) AFM of the niobium surface before any processing occurs. (Middle) AFM of the niobium surface after an oxygen ash and 10 second BOE etch. (Right) AFM of the niobium surface after an oxygen ash and 60 second BOE etch.

The cleaning procedure that I conclude is the best for niobium films and sapphire substrates is an oxygen plasma ash and a BOE etch. At this point, however, I have not provided any good images of a 60 second BOE etch. I suspected that a 60 second etch could be a viable option, if the rinsing procedure could be done correctly. To investigate this, I grew another niobium film, wafer e545 and performed microscopy on the unprocessed surface. As with wafer e542, I fabricated resonators and diced the sample – in this case into 15 pieces. All of the pieces underwent the oxygen plasma ash, then half the pieces were etched for 10 seconds in BOE and half the pieces were etched for 60 seconds in BOE. The roughness of the unprocessed film is 235pm; the RMS roughness values of the cleaned wafers are 215pm and 265pm for the 10 second and 60 second etches respectively. A visual inspection of the microscopy also indicates that the two etches are essentially equivalent. The sapphire surfaces are also very similar; the RMS roughness of the 10 second etch is 100pm, and the RMS roughness of the 60 second etch is 75pm. This wafer doesn't show the atomic-scale terracing that were observed on sample

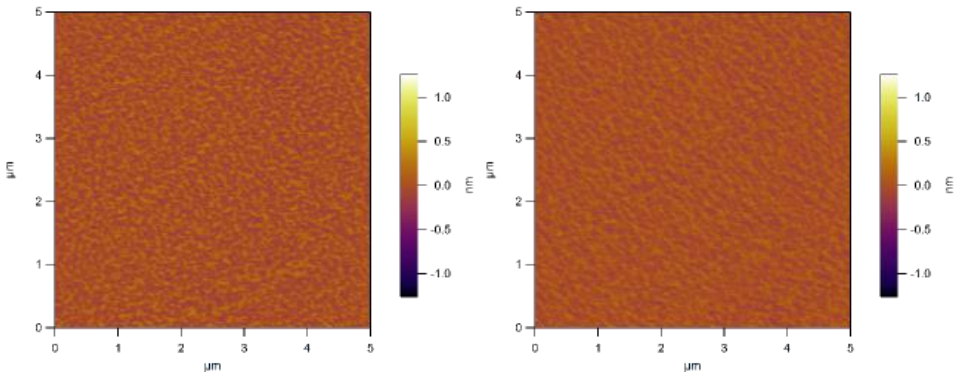

Figure 72 (Left) AFM of the sapphire surface after an oxygen ash and 10 second BOE etch. (Right) AFM of the sapphire surface after an oxygen ash and 60 second BOE etch.

e542, but the RMS roughness is equivalent between the two samples. As with wafer e542, I observed more debris and defects on samples that were (accidentally) dropped into the rinse water beaker, and thus not so thoroughly rinsed as the examples shown above. I conclude, then that an oxygen plasma ash when combined with a 10-60 second BOE etch is sufficient to clean the sample surface, provided that the sample is thoroughly rinsed following the acid bath.

7.3 Internal-Q and Coherence Measurements of Cleaned Samples

While the microscopy measurements show that the niobium and sapphire surface can be returned to pristine condition after an oxygen ash and BOE clean, these measurements cannot provide any information about the effectiveness of the cleaning procedure on quantum coherence in actual devices. Moreover, there is always a concern that the cleaning procedure might do invisible damage to the qubit, which would only be apparent upon measuring its coherence time. In this section, I shall report measurements made of both the internal-Q values for cleaned resonators, and T_1 and T_2^* values for epitaxial transmons that were also subjected to the oxygen ash and BOE clean.

The resonator internal-Q values were previously reported in Table 2, but the relevant values are repeated here. Wafer e545, which has AFM shown in Figure 71 and Figure 72, had an internal-Q of 130,000. Wafer e547 (no AFM) had an internal-Q of 250,000, among the highest of all the measured resonators. It should be noted that wafer e545 had a spurious resonance which artificially lowered the internal-Q of the chip. This spurious resonance is thought to be due to incomplete removal of the titanium from the back of the chip. Unfortunately, sample cleaning does not appear to improve the internal-Q by a large margin, if at all. Neither e545 nor e547 had the highest internal-Q for single-crystal niobium films, whether cleaned or not. At best, these measurements indicate that cleaning does not appear to reduce the internal-Q.

Coherence measurements made on transmons tell a different story. The transmons were fabricated from an all-epitaxial trilayer structure using the self-aligned structure described in Section 5.3; the fabrication process was identical to the Josephson junction, the only difference in fabrication was the geometry of the device. After fabrication, the transmons were sent to UCB for measurement. They were then returned to me for cleaning, then sent back to UCB for a second coherence time measurement. The results are shown in Table 5.

	T_1	T_2^*
BEFORE CLEANING	540 nanoseconds	80 nanoseconds
AFTER CLEANING	1200 nanoseconds	850 nanoseconds
IMPROVEMENT FACTOR	2.2	10.6

Table 5 Coherence times for an epitaxial transmon device before and after the cleaning procedure.

There is a two-fold improvement in the decoherence time T_1 and more than a 10-fold improvement in the dephasing time T_2^*. Both measurements indicate that fabrication residue is a source of noise in superconducting qubits, and removing this residue will be a key part of producing long coherence time qubits.

The fact that $T_2^* \ll 2\,T_1$ means that extraneous dephasing is present in these devices. The improvement in T_1 indicates that at least some of the effect of cleaning was to remove lossy material in the vicinity of the devices, specifically where the electromagnetic field from the device operation is non-zero. Additionally, the larger improvement in T_2^* shows that some of the defects present that were cleaned away lead to dephasing, perhaps caused by two level systems. Such organic debris is clearly a problem for qubits and methods to clean the processed surface need to be developed.

Chapter 8 Summary and Future Directions

The goal of this project was to understand and reduce sources of noise in niobium-based superconducting qubits, with a particular focus on noise sources which affect flux qubits. This was done by using a UHV MBE chamber to precisely control the growth parameters of the films used in the qubits. DC-SQUIDs, individual Josephson junctions, resonators, and transmons were all fabricated to evaluate different noise sources. The films themselves were analyzed using AFM, RHEED, and TEM to study the crystal structure and topography of the niobium.

By using an epitaxial niobium film on sapphire, I was able to determine that localized MIGS are, in fact, a source of flux noise. In the paper by Choi, et al [25], it was theorized that disorder will induce localization of evanescent waves in the insulating substrate resulting in states with single occupancy due to Coulomb blockade. These localized states are magnetically active, modulating the external flux threading the qubit. TEM measurements of niobium films grown at 610°C, 675°C, and 750°C show excellent order at the interface. DC-SQUIDs made from these films were sensitive to such magnetic fluctuations allowing for measurement of the flux spectral noise density. Measurements made on SQUIDs made from all-epitaxial niobium trilayers (complete with MBE junctions) show a small decrease in flux noise, but may not conclusively infer the existence of the localized MIGS. Comparing polycrystalline reference devices made by NIST to my devices showed the single-crystal SQUIDs had on average a lower flux noise at 1Hz by approximately 1.5 dB. This ambiguity is likely due to flux noise from the Josephson junction itself, which may be large enough to be a significant contributor to noise in the SQUID circuit. I found that, smaller junctions, deposited using shadow-evaporation, are not a significant contributor to flux noise. Devices made using epitaxial loops and shadow-evaporated junctions show the lowest yet-reported flux noise in SQUIDs. Compared to the flux noise observed in polycrystalline samples (NIST), I measured a lower flux spectral noise density by more than 11 dB. This substantial reduction in flux noise is due to the use of single crystal materials and is experimental validation that disorder-induced localized MIGS leads to flux noise.

In the course of analyzing the epitaxy of the niobium at the substrate/metal interface, I discovered that oxygen was depleted from the sapphire due to the high temperature of the sapphire substrate during the niobium growth. This surprising result had not previously been anticipated, and further measurements may provide a better estimate of the magnitude of noise from oxygen reduction of the sapphire substrate. The oxygen out-diffusion I measured followed a thermally activated temperature dependence, and I estimated that diffusion would be completely quenched at temperatures in the range

of 430°C to 465°C. I propose that reduction of oxygen from sapphire is a problem because the oxygen vacancies in aluminum oxide form charge traps. When a single electron occupies one of these traps, its spin contributes to flux noise in the system. These charge traps are almost certainly inhomogeneous in structure and energy spectrum. The emission and trapping rates will depend on the structure and spectrum, resulting in an ensemble of defects that have a 1/f noise spectrum.

In light of the previous result, which shows that disorder leads to localized MIGS, it is necessary to strike a balance between epitaxy and low oxygen reduction. STEM shows an epitaxial interface is obtained at growth temperatures as low as 610°C, although low magnification TEM shows a high density of dislocations occurring away from the substrate interface. AFM measurements show smooth films as low as 637°C, so growth between 640-650°C should give good epitaxy. It is expected that there will be oxygen vacancies in the substrate for growths in this range – further experimentation will be required to determine if these defects will significantly reduce coherence times; see my suggested solution to this problem below.

Another type of oxygen vacancy defect was studied in the Josephson junction barrier composition. This thin insulator was grown under different oxygen pressures and with different aluminum fluxes. Analysis of the R_nA products suggests that low oxidation dose leaves oxygen vacancies in the grown alumina barrier. Stronger evidence that low oxygen dose is detrimental for these films is found in the temperature dependence of the critical current. For a sufficiently high oxidation, the critical current temperature dependence is virtually ideal, but for low doses, the temperature dependence is approximately linear with temperature. This linear behavior is similar to the critical current temperature dependence observed in aluminum oxide barriers intentionally doped with magnetic defects [63]. My junctions have no explicitly added magnetic impurities; in my case the oxygen-poor barriers exhibit magnetic impurities that can be attributed to singly-occupied electron traps due to oxygen vacancies in the alumina. They lead to flux noise in the same manner that oxygen vacancies lead to flux noise in the substrate.

Finally, the effect of chemical contaminants on both T_1 and T_2 was explored. The main contaminant is incompletely removed photoresist (MMA/PMMA) residue that is strongly bonded to either the metal surface or the substrate. Nine different chemical and plasma etches were tested to determine the best method to remove photoresist residue from the surfaces of the devices after fabrication. Based on AFM images, I found that the best method used an oxygen plasma ash combined with a 10 – 60 second buffered hydrofluoric acid etch (BOE) to return the sample to a pristine state. This process was

subsequently tested on a transmon qubit and it resulted in a 2.2 times improvement on T_1 and a 10.6 times improvement on T_2^*. Clearly, residual material from fabrication processes leads to environmental dephasing and loss; cleaning processes must be incorporated into all fabrication processes to maximize coherence times.

Based on these results, the ideal flux qubit made using niobium would have a base electrode grown between at lower temperatures that permit an epitaxial interface while limiting oxygen out-diffusion from the substrate. The buffer aluminum should be thin, on the order of 20-50Å. The aluminum flux should be as slow as possible while still being able to make reliable RHEED oscillation measurements – the best results here used a growth rate of 40 seconds/monolayer, and 50 seconds/monolayer might be achievable. The slow aluminum flux rate should be paired with a high oxygen deposition pressure; here the highest used was 5×10^{-6} but higher pressures may be feasible. If a different pumping scheme were used (i.e. one that didn't use a cryo pump), it might be possible to use ozone instead of oxygen and produce a truly perfect homogenous aluminum oxide barrier. Finally, sample cleaning must be incorporated into the qubit fabrication. In this work, sample cleaning was only done after device wiring, immediately before testing – it may be better/necessary to incorporate cleaning into all processing steps, so that dirt is not buried beneath the wiring or under dielectric layers. Resonators made from this structure would be expected to have Q-factors on the order of 700,000 based on the Q's reported in this work. A rough estimate of the coherence time from this Q and the average qubit frequency of 5 GHz follows: this frequency corresponds to a period of 2×10^{-10} seconds and multiplying by the Q results in a time of 140 microseconds. It is reasonable to conclude that films with Q's of this magnitude can be fabricated into qubits with coherence times greater than 100 microseconds.

There are several experiments that would be useful extensions of this work. All of the trilayer Josephson junctions used in this work were relatively large (> 1 micron x 1 micron) and may have sources of noise (such as strain) that are not present in their smaller shadow-evaporated aluminum junction counterparts [18]. Continued improvements in the self-aligned junction technology might allow for small niobium junctions on the same scale as the shadow-evaporated junctions, allowing for a true side-by-side comparison. Alternatively, fabrication of large aluminum junctions might reveal whether these devices have similar noise levels as the niobium junctions.

Based on the results of the cleaning process, it would be prudent to repeat the flux noise measurements with SQUIDs that had been subjected to a cleaning process, in case any excess noise was due to

processing residue. It would also be interesting to see if cleaning the photoresist residue might allow for a measureable difference between SQUIDs with cladding layers and SQUIDs without them.

Finally, it would be useful to continue the TEM study of the niobium/sapphire interface. Measurements of samples grown at additional temperatures would help determine the rate of adsorption of oxygen from the sapphire. In particular, measurements in the 610°C - 675°C range, which appears to be the ideal growth range, would help determine the quantity of oxygen vacancies in these films. Of specific interest would be whether time is a significant contributor to oxygen reduction and whether oxygen vacancies can be reduced by heating the substrate for a shorter period of time. In particular, I would suggest growing the first few layers of the base electrode at a temperature high enough to get good long range crystalline order. Then the substrate temperature would be reduced to a temperature thought significantly low enough to quench diffusion. By reducing the time the samples are hot, it may be possible to obtain both a good interface and nearly zero oxygen vacancies in the substrate. Flux noise measurements of films would determine the effect of oxygen depletion on noise in superconducting qubits.

Appendix A The Shadow-Evaporation Technique

In this section I shall describe the shadow-evaporation process (also called double-angle evaporation) used to create small aluminum junctions. This process was used to create the junctions for the SQUIDs described in Figure 56 and Figure 57, but is otherwise ancillary to the bulk of this thesis.

The shadow-evaporation technique is commonly used by researchers in the superconducting qubit field [8, 9, 10, 17, 67] and was originally developed by Dolan in 1977 [68]. It has the benefit of allowing for two layers to be grown with only a single lithography step and uses lift-off rather than etching. Two

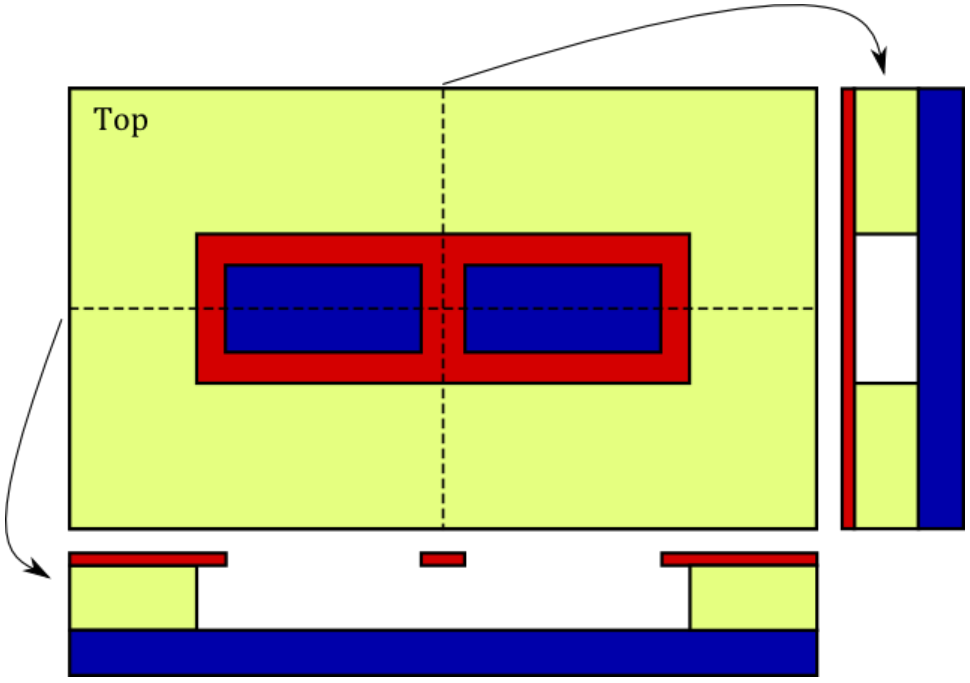

Figure A-1 Top view and cross section views of the photoresist pattern used for shadow-evaporation. Red is the topmost photoresist, yellow is the underlying photoresist, and blue is the substrate. In the top view, the red layer should actually cover the entire surface, but as this obscures the yellow, the red is only shown where there is no yellow underneath it. The cross section views are cuts at the two dashed lines. These clearly show the undercut of the yellow layer and the suspended bridge of the red layer.

layers of photoresist are spun, the bottom of which must be able to have a substantial undercut relative to the top layer. The key to the shadow-evaporation process is a suspended bridge, where the top photoresist is still extant as a mask, but the lower photoresist has been washed away in the developer solution. Deposition at an angle relative to this surface allows for overlap beneath this bridge structure as seen in Figure A-2.

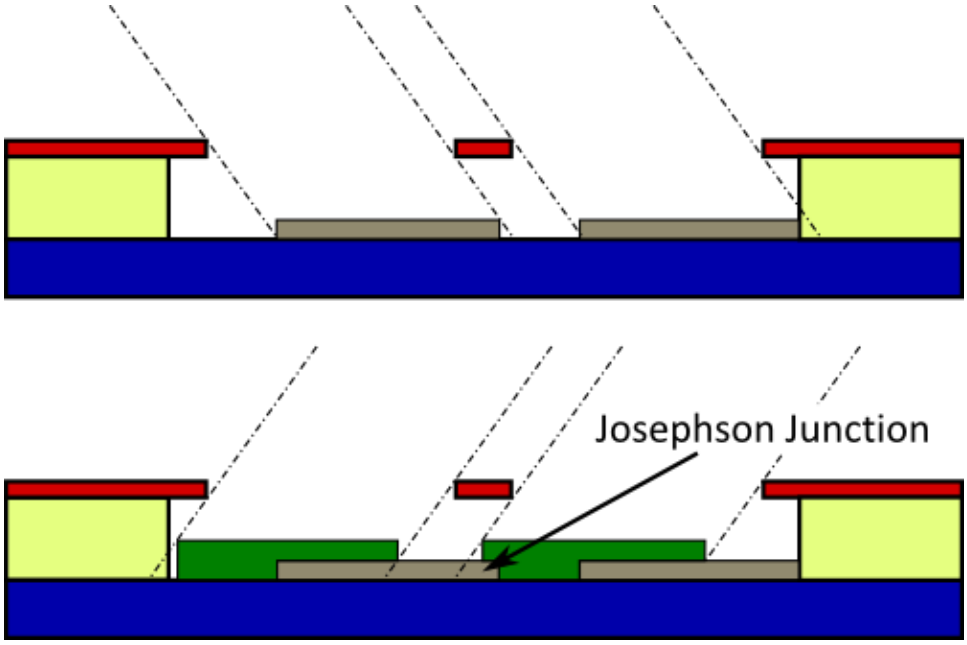

Figure A-2 Shadow evaporation deposition steps. (Top) First deposition comes in from the left. (Bottom) Second deposition comes in from the right and has a small overlap with the first layer.

After the first deposition, the sample is rotated so that it is at a different angle relative to the evaporation source (there is no requirement that the two angles be symmetric or opposite). The second deposition results in overlap with the initial layer. As seen in Figure A-2, there are three overlapping regions, two large and one small, each of which is a Josephson junction. With a thoughtful design, the two large junctions can be sufficiently large to always stay in the zero-voltage state and do not contribute to the dynamics of the system. An example of two shadow-evaporated junctions are shown in Figure A-3.

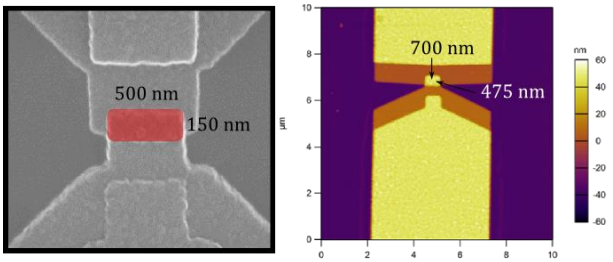

Figure A-3 Examples of shadow-evaporated junctions. (Left) SEM micrograph of a junction. The junction is highlighted red for clarity. Image courtesy D. van Harlingen, UIUC. (Right) AFM of an aluminum shadow-evaporated junction.

The shadow-evaporation process works well with materials like aluminum that can be grown smoothly at room temperature. Elevated substrate temperatures are not possible with this process because even as little as 100°C may cause the photoresist to flow and destroy the pattern. Furthermore, shadow-evaporation with multiple materials (such as niobium-aluminum oxide-niobium) poses technical challenges that have yet to be overcome.

References

[1] D. Deutsch and R. Jozsa, "Rapid Solution of Problems by Quantum Computation," *Proceedings of the Royal Society of London Series A,* vol. 439, no. 1907, pp. 553-558, 1992.

[2] L. K. Grover, "Quantum mechanics helps in searching for a needle in a haystack," *Physical Review Letters,* vol. 79, no. 2, pp. 325-328, 1997.

[3] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *35th Annual Symposium on Foundations of Computer Science*, Los Alamitos, CA, 1994.

[4] M. A. Nielsen and I. L. Chuang, Quantum Compuation and Quantum Information, Cambridge: Cambridge University Press, 2010.

[5] D. P. DiVincenzo, "The Physical Implementation of Quantum Computation," *Fortschr Phys,* vol. 48, no. 9, pp. 771-783, 2000.

[6] J. M. Martinis, S. Nam, J. Aumentado and C. Urbina, "Rabi Oscillations in a Large Josephson-Junction Qubit," *Physical Review Letters,* vol. 89, no. 11, 2002.

[7] V. Bouchiat, D. Vion, P. Joyez, D. Esteve and M. H. Devoret, "Quantum coherence with a single cooper pair," *Physica Scripta,* vol. T76, pp. 165-170, 1998.

[8] J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal and S. Lloyd, "Josephson Persistent-Current Qubit," *Science,* vol. 285, pp. 1036-1039, 1999.

[9] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin and R. J. Schoelkopf, "Charge-insensitive qubit design derived from the Cooper pair box," *Physical Review A,* vol. 76, no. 4, 2007.

[10] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve and M. H. Devoret, "Manipulating the Quantum State of an Electrical Circuit," *Science,* vol. 296, pp. 886-889, 2002.

[11] D. J. Van Harlingen, T. L. Robertson, B. L. T. Plourde, P. A. Reichardt, T. A. Crane and J. Clarke, "Decoherence in Josephson-junction qubits due to critical-current fluctuations," *Physical Review B,* vol. 70, no. 064517, 2004.

[12] J. R. Friedman, V. Patel, W. Chen, S. K. Tolpygo and J. E. Lukens, "Quantum superposition of distinct macroscopic states," *Nature,* vol. 406, pp. 43-46, 2000.

[13] T. P. Orlando, S. Lloyd, L. S. Levitov, K. K. Bergren, M. J. Feldman, M. F. Bocko, J. E. Mooij, C. J. P. Harmans and C. H. van der Wal, "Flux-based superconducting qubits for quantum computation," *Physica C,* pp. 194-200, 2002.

[14] Y. Makhlin, G. Schon and A. Shnirman, "Quantum-state engineering with Josephson-junction devices," *Reviews of Modern Physics,* vol. 73, pp. 357-400, 2001.

[15] T. P. Orlando, J. E. Mooij, L. Tian, C. H. van der Wal, L. S. Levitov, S. Lloyd and J. J. Mazo, "Superconducting Persistent-Current Qubit," *Physical Review B,* vol. 60, no. 22, pp. 15398-15413, 1999.

[16] J. Clarke and F. K. Wilhelm, "Superconducting Quantum Bits," *Nature,* vol. 453, pp. 1031-1042, 2008.

[17] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret and R. J. Schoelkopf, "Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture," *Physical Review Letters,* vol. 107, no. 24, 2011.

[18] W. D. Oliver and P. B. Welander, "Materials in superconducting quantum bits," *MRS Bulletin,* vol. 38, pp. 816-825, 2013.

[19] F. C. Wellstood, C. Urbina and J. Clarke, "Low-frequency noise in DC superconducting quantum interference devices below 1K," *Applied Physics Letters,* vol. 50, no. 12, pp. 772-774, 1987.

[20] R. Koch, J. Clarke, W. M. Goubau, J. M. Martinis, C. M. Pegrum and D. J. Van Harlingen, "Flicker (1/f) Noise in Tunnel Junction DC-SQUIDs," *Journal of Low Temperature Physics,* vol. 51, no. 1/2, 1983.

[21] R. H. Koch, D. P. DiVincenzo and J. Clarke, "Model fo 1/f Flux Noise in SQUIDs and Qubits," *Physical Review Letters,* vol. 98, no. 26, 2007.

[22] S. Sendelbach, D. Hover, A. Kittel, M. Muck, J. M. Martinis and R. McDermott, "Magnetism in SQUIDs at millikelvin temperatures," *Physical Review Letters,* vol. 100, no. 22, 2008.

[23] H. Bluhm, J. A. Bert, N. C. Koshnick, M. E. Huber and K. A. Moler, "Spinlike Susceptibility of Metallic and Insulating Thin Films at Low Temperatures," *Physical Review Letters,* vol. 103, no. 2, 2009.

[24] P. G. Bjornsson, B. W. Gardner, J. R. Kirtley and K. A. Moler, "Scanning superconducting quantum interference device microscope in a dilution refrigerator," *Reivew of Scientific Instruments,* vol. 72, no. 11, 2001.

[25] S. Choi, D. H. Lee, S. G. Louie and J. Clarke, "Localization of metal-induced gap states at the metal-insulator interface: origin of flux noise in SQUIDs and superconducting qubits," *Physical Review Letters,* vol. 103, no. 19, 2009.

[26] P. W. Anderson, "Absence of Diffusion in Certain Random Lattices," *Phyiscal Review,* vol. 109, no. 5, pp. 1492-1505, 1958.

[27] P. B. Welander and J. N. Eckstein, "Strained single-crystal Al2O3 grown layer by layer on Nb (110) thin films," *Applied Physics Letters,* vol. 90, no. 24, 2007.

[28] A. R. Wildes, J. Mayer and K. Theis-Brohl, "The growth and structure of epitaxial niobium on sapphire," *Thin Solid Films,* pp. 7-34, 2001.

[29] C. P. Flynn, "Constrains on the Growth of Metallic Superlattices," *Journal of Physics F: Metal Physics,* vol. 18, pp. L195-L200, 1988.

[30] P. B. Welander, T. J. McArdle and J. N. Eckstein, "Reduced leakage current in Josephson tunnel junctions with codeposited barriers," *Applied Physics Letters,* vol. 97, p. 233510, 2010.

[31] N. W. Ashcroft and N. D. Mermin, Solid State Physics, Thomson Learning, 1976.

[32] S. A. Wolf, S. B. Qadri, J. H. Claassen, T. L. Francavilla and B. J. Dalrymple, "Epitaxial Growth of Superconducting Niobium Thin Films by Ultrahigh Vacuum Evaporation," *Journal of Vacuum Science and Technology A,* vol. 4, pp. 524-527, 1986.

[33] G. Oya, M. Koishi and Y. Sawada, "Highquality singlecrystal Nb films and influences of substrates on the epitaxial growth," *Journal of Applied Physics,* vol. 60, pp. 1440-1446, 1986.

[34] G. L. Zhou and C. P. Flynn, "Fingered Morphology of Niobium (110) Grown by Molecular-Beam Epitaxy," *Physical Review B,* vol. 59, no. 12, pp. 7860-7867, 1999.

[35] P. Welander, *Epitaxial Aluminum Oxide Thin Films on Niobium (110): A Study of Their Growth and Their Use in Superconducting Tunnel-Junctions,* University of Illinois, 2007.

[36] W. De Sorbo, "Effect of Dissolved Gases on Some Superconducting Properties of Niobium*," *Physical Review,* vol. 132, no. 1, pp. 107-121, 1963.

[37] Y. V. Shvyd'ko, M. Lucht, E. Gerdau, M. Lerche, E. E. Alp, W. Sturhahn, J. Sutter and T. S. Toellner, "Measuring Wavelengths and Lattice Constants with the Mössbauer Wavelength Standard," *Journal of Synchrotron Radiation,* vol. 9, pp. 17-23, 2002.

[38] G. J. Gutekunst, J. Mayer and M. Rühle, "Atomic Structure of Epitaxial Nb-Al2O3 Interfaces I. Coherent Regions," *Philosophical Magazine A,* vol. 75, no. 5, pp. 1329-1355, 1997.

[39] S. M. Durbin, J. E. Cunningham and C. P. Flynn, "Growth of Single-Crystal Metal Super-Lattices in Chosen Orientations," *Journal of Physics F: Metal Physics,* vol. 12, no. L75, 1982.

[40] Y. Igarashi and M. Kanayama, "Highquality Singlecrystal Nb and Ta Films Formed by an Ultrahigh Vacuum Arc Method," *Journal of Applied Physics,* vol. 57, pp. 849-854, 1985.

[41] C. Sürgers, M. Schöck and H. Löhneysen, "Oxygen-induced surface structure of Nb(110)," *Surface Science,* vol. 471, pp. 209-218, 2001.

[42] I. Arfaoui, J. Cousty and C. Guillot, "A model of the NbOx nanocyrstals tiling a Nb(110) surface annealed in UHV," *Surface Science,* vol. 557, pp. 119-128, 2004.

[43] M. Ondrejcek, R. S. Appleton, J. Swiech, V. L. Petrova and C. P. Flynn, "Thermally activated stripe reconstruction induced by O on Nb(011)," *Physical Review Letters,* vol. 87, no. 11, 2001.

[44] E. R. Dobrovinskaya, L. A. Lytvynov and V. Pishchik, Sapphire: Material, Manufacturing, Applications, New York: Springer, 2009.

[45] MicroChem Corp, *Nano PMMA and Copolymer Data Sheet,* Newton, MA: MicroChem Corp, 2001.

[46] J. Wenner, R. Barends, R. C. Bialczak, Y. Chen, J. Kelly, E. Lucero, M. Mariantoni, A. Megrant, P. J. J. O'Malley, D. Sank, A. Vainsencher, H. Want, T. C. White, Y. Yin, J. Zhao, A. N. Cleland and J. M. Martinis, "Surface Loss Simulations of Superconducting Coplanar Waveguide Resonators," *Applied Physics Letters,* vol. 99, p. 113513, 2011.

[47] A. Megrant, C. Neill, R. Barends, B. Chiaro, Y. Chen, L. Feigl, J. Kelly, E. Lucero, M. Mariantoni, P. J. J. O'Malley, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Y. Yin, J. Zhao, C. J. Palmstrøm, J. M. Martinis and A. N. Cleland, "Planar superconducting resonators with internal quality factors above one million," *Applied Physis Letters,* vol. 100, no. 11, 2012.

[48] C. D. Nugroho, V. Orlyanchik and D. J. Van Harlingen, "Low frequency resistance and critical current fluctuations in Al-based Josephson junctions," *Applied Physics Letters,* vol. 102, 2013.

[49] T. Van Duzer and C. W. Turner, Principles of Superconductive Devices and Circuits, 2nd ed., Upper Saddle River: Prentice Hall, 1999.

[50] M. Tinkham, Introduction to Superconductivity, Mineola, New York: Dover Publications, Inc., 1996.

[51] B. L. Blackford and R. H. March, "Temperature dependence of the energy gap in superconducting Al-Al2O3-Al tunnel junctions," *Canadian Journal of Physics,* vol. 46, no. 141, 1968.

[52] A. W. Lichtenberger, C. P. McClay, R. J. Mattauch, M. J. Feldman, S. K. Pan and A. R. Kerr, "Fabrication of Nb/Al-Al2O3/Nb Junctions with extremely low leakage currents," *IEEE Transactions on Magnetics,* vol. 25, no. 2, pp. 1247-1250, 1989.

[53] H. A. Huggins and M. Gurvitch, "Preparation and characterization of Nb/Al-oxide-Nb tunnel junctions," *Journal of Applied Physics,* vol. 57, no. 6, pp. 2103-2109, 1985.

[54] R. Monaco, R. Cristiano, L. Frunzio and C. Nappi, "Investigation of low-temperature I-V curves of high-quality Nb/Al-AlOx/Nb Josephson junctions," *Journal of Applied Physics,* vol. 71, no. 4, pp. 1888-1892, 1992.

[55] A. Hahn, S. Hofmann, A. Krause and P. Seidel, "Tunneling results on gap anisotropy in niobium," *Physica C,* vol. 296, pp. 103-118, 1998.

[56] A. Barone and G. Paterno, Physics and Applications of the Josephson Effect, New York: John Wiley and Sons, 1982.

[57] W. Simon, E. K. Liebemann, M. Simon and E. Bucher, "Nb-Al-AlOx-Al-Nb tunnel junctions using electron beam evaporation," *Journal of Applied Physics,* vol. 72, no. 9, pp. 4474-4476, 1992.

[58] A. Inoue, A. Nakayama and Y. Okabe, "Nb/AlOx/Nb Josephson Tunnel Junctions Using Electron-Beam Evaporation," *Japanese Journal of Applied Physics,* vol. 27, no. 7, pp. 1234-1238, 1988.

[59] J. K. Hulm, C. K. Jones, R. A. Hein and J. W. Gibson, "Superconductivity in the TiO and NbO systems," *Journal of Low Temperature Physics,* vol. 7, no. 3-4, pp. 291-307, 1972.

[60] J. Bardeen, L. N. Cooper and J. R. Schrieffer, "Theory of Superconductivity," *Phys. Rev.,* vol. 108, no. 5, pp. 1175--1204, 1957.

[61] V. Ambegaokar and A. Baratoff, "Tunneling between superconductors," *Physical Review Letters,* vol. 10, no. 11, pp. 486-489, 1963.

[62] V. Ambegaokar and A. Baratoff, "Errata: Tunneling Between Superconductors," *Physical Review Letters,* vol. 11, no. 2, p. 104, 1963.

[63] A. Bill, S. A. Wolf, Y. N. Ovchinnikov and V. Z. Kresin, "Effect of magnetic impurity correlations on Josephson tunneling," *Physica C,* vol. 298, pp. 231-239, 1998.

[64] B. D. Josephson, "Possible new effects in superconductive tunnelling," *Physics Letters,* vol. 1, no. 7, pp. 251-253, 1962.

[65] P. Seidel and H. G. Meyer, "Critical Josephson current of tunnel junctions," *Physica Status Solidi (b),* vol. 122, no. 1, pp. K105-K107, 1984.

[66] K. L. Chavez and D. W. Hess, "Removal of Resist Materials Using Acetic Acid," *Journal of the Electrochemical Society,* vol. 150, no. 4, pp. G284-G291, 2003.

[67] V. E. Manucharyan, J. Koch, L. I. Glazman and M. H. Devoret, "Fluxonium: single Cooper-pair circuit free of charge offsets," *Science,* vol. 326, pp. 113-116, 2009.

[68] G. J. Dolan, "Offset masks for lift-off photoprocessing," *Applied Physics Letters,* vol. 31, no. 5, pp. 337-339, 1977.

[69] Z. Yoscovits, *Materials sources of decoherence in superconducting qubits,* Urbana: University of Illinois at Urbana-Champaign, 2014.

[70] Wikipedia, "Atomic Force Microscopy --- Wikipedia, The Free Encyclopedia," 6 October 2009. [Online]. Available: http://en.wikipedia.org/wiki/Atomic_force_microscopy. [Accessed 7 August 2013].

[71] Z. Wang, A. Kawakami, Y. Uzawa and B. Komiyam, "High critical current density NbN/AlN/NbN tunnel junctions fabricated on ambient temperature MgO substrates," *Applied Physics Letters,* vol. 64, no. 15, pp. 2034-2036, 1994.

[72] Z. Wang, A. Saito, A. Kawakami and K. Hamasaki, "Characterization of low frequency noise in epitaxial NbN/AlN/NbN tunnel junctions," *IEEE Transactions on Applied Superconductivity,* vol. 13, no. 2, pp. 131-134, 2003.

[73] R. T. Wakai and D. J. van Harlingen, "Low-frequency nois and discrete charge trapping in small-area tunnel junction DC SQUID's," *Applied Physics Letters,* vol. 49, no. 10, pp. 593-595, 1986.

[74] C. H. van der Wal, A. C. J. ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. Harmans, T. P. Orlando, S. Lloyd and J. E. Mooij, "Quantum superposition of macroscopic persistent-current states," *Science,* vol. 290, pp. 773-777, 2000.

[75] W. Święch, R. S. Appleton, B. D. Wiemeyer and C. P. Flynn, "LEEM investigation of refractory metals grown on insulating substrates: Nb (110) on Sapphire (1120)," *Surface Review and Letters,* vol. 5, no. 6, pp. 1221-1231, 1998.

[76] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin and R. J. Schoelkopf, "Resolving photon number states in a superconducting circuit," *Nature,* vol. 445, 2007.

[77] B. Savo, F. C. Wellstood and J. Clarke, "Low-frequency excess noise in Nb-Al2O3-Nb Josephson tunnel junctions," *Applied Physics Letters,* vol. 50, no. 24, pp. 1757-1759, 1987.

[78] J. M. Sage, V. Bolkhovsky, W. D. Oliver, B. Turek and P. B. Welander, "Study of loss in superconducting coplanar waveguide resonators," *Journal of Applied Physics,* vol. 109, p. 063915, 2011.

[79] C. T. Rogers and R. A. Buhrman, "Composition of 1/f noise in metal-insulator-metal tunnel junctions," *Physical Review Letters,* vol. 53, no. 13, pp. 1272-1275, 1984.

[80] S. Ramo, J. R. Whinnery and T. Van Duzer, Fields and Waves in Communication Electronics, New York: John Wiley & Sons, Inc., 1994.

[81] R. Pantel, M. Bujor and J. Bardolle, "Continuous measurement of surface potential variations during oxygen adsorption on the (100), (110), and (111) faces of niobium using mirror electron microscope," *Surface Science,* vol. 62, pp. 589-609, 1977.

[82] R. McDermott, "Materials Origins of Decoherence in Superconducting Qubits," *IEEE Transactions on Applied Superconductivity,* vol. 19, no. 1, pp. 2-13, 2009.

[83] J. Leppäkangas and M. Marthaler, "Fragility of flux qubits against quasiparticle tunneling," *Physical Review B,* vol. 85, p. 144503, 2012.

[84] M. L. Kronberg, "Plastic Deformation of Single Crystals of Sapphire: Basal Slip and Twinning," *Acta Metallurgica,* vol. 5, no. 9, pp. 507-524, 1957.

[85] J. Clarke, *Personal Communication with the Author,* 2011.

[86] MicroChem Corp., "Nano PMMA and Copolymer," 2001. [Online]. Available: http://www.microchem.com/pdf/PMMA_Data_Sheet.pdf.