Rapid Fabrication of Deep-Submicron Josephson Junctions

by

Bryan M. Cord

Submitted to the Department of Electrical Engineering and Computer Science in Partial Fulfillment of the Requirements for the Degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2004
[June 2004]

Author: _____
Department of Electrical Engineering and Computer Science
May 27, 2004

Certified By: _____
Terry P. Orlando
Professor of Electrical Engineering
Thesis Supervisor

Approved By: _____
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Rapid Fabrication of Deep-Submicron Josephson Junctions

By

Bryan M. Cord

Submitted to the Department of Electrical Engineering and Computer Science on May 27, 2004 in Partial Fulfillment of the Requirements for the Degree of Master of Science in Electrical Engineering

ABSTRACT

A process was designed to fabricate Nb-AlO_x-Nb Josephson junctions for quantum computing applications, with the goal of fabricating junctions as small as 0.2 μm in diameter in a time span of roughly two weeks. The process was based on the doubly-planarized all-refractory technology for superconductive circuits (DPARTS) process currently used to fabricate these devices at MIT Lincoln Laboratory and streamlined by removing or replacing the most time-consuming steps and several optical layers. In addition, 248-nm (deep-UV) photolithography was employed for the first time in a Nb-based process, with the goal of improving resolution past that achievable by standard i-line lithography. The process has five optical layers, two wiring layers, and a via layer, and is intended to be used for rapid-turnaround evaluation of simple circuits requiring only two wiring layers.

Anodization was used to produce a 50 nm film of NbO_x to isolate the wiring layers, replacing the time-consuming oxide deposition and planarization used in the DPARTS process. A novel metallization and liftoff process employing surface-poisoning of chemically-amplified resist was used to deposit the contact layer. Several new plasma etching techniques were developed to selectively etch the various materials present in the process as well.

The full process flow is briefly described in order of fabrication, followed by a detailed discussion of major process steps, issues, and results. Finally, testing results from devices fabricated using the new process are presented. An appendix detailing the design-of-experiment (DOE) approach used to characterize several process tools is also provided.

Thesis Supervisor: Terry P. Orlando
Title: Professor of Electrical Engineering

Acknowledgements

The majority of the work described here was performed at MIT Lincoln Laboratory in the Analog Device Technology group (86), and I would like to first thank the many people there who contributed in some way to the TD1 project. I am particularly grateful to Vladimir Bolkhovsky, who headed the project, for his vision, insight, support, and tenacity. Earle Macedo was also instrumental in making the project a success by contributing ideas and base material, as well as training me to use several processing tools. Terry Weir and Will Oliver were invaluable during device testing, as were Jay Sage and Dan Santiago during the CAD design of the TD1 reticle set. Mark Gouker, the leader of the Analog Device Technology group, was kind enough to allow me to work on a major government project, and the group secretary, Joanne Rantz, was infinitely helpful in more ways than can be easily listed.

The entire staff of the Microelectronics Laboratory (MEL), Lincoln Laboratory's clean room facility, were in some way involved in the fabrication of the TD1 wafers, and thanks go out to everyone on both shifts there. Special thanks are in order for Marshall Kaplan, Kyle Keenan, Denise Holohan, Jinete Buckley, Marge Wood, Sandi Roy, Carl Creel, Dave Astolfi, and Renee Lambert, who helped train me on various processing and metrology tools, process TD1 wafers on time, and were generally incredibly good sports about having a graduate student around the clean room.

At MIT, I would like to thank my thesis supervisor Terry Orlando for being endlessly supportive and understanding, as well as successfully facilitating my working position at Lincoln Laboratory. Karl Berggren was in charge of the TD1 project for the first three months that I worked on it, but continued to be a major source of ideas and insight after leaving Lincoln Laboratory to become a professor at MIT in September. Daniel Nakada, who was involved in fabrication on the original DPARTS process, graciously contributed his knowledge and expertise whenever it was needed. Finally, Henry I. Smith was not affiliated with the TD1 project, but his class on fabrication technology provided me with expertise that directly translated into some of the better SEM images in this thesis, among other things.

Table of Contents

3. Process Design

1 Introduction

1.1 Background

1.1.1 Quantum Computing

A quantum computer is a device that stores information in quantum variables and processes that information by causing those variables to interact in a quantum-mechanically coherent manner[1-5]. Typically, a quantum computer is made up of many two-state quantum systems known as "quantum bits" or qubits[6], which are then coupled together to perform computations. The actual quantum variables used can vary; photons, the spins of particles, and atomic orbital levels are some of the systems that have been examined for possible quantum-computation applications.

In order to perform a computation on a hypothetical quantum computer, one must be able to prepare the qubits in a known initial state, coherently manipulate the superpositions of the two states in a qubit, couple qubits together dynamically, and measure the state of a given qubit, all the while keeping the system free of noise and other sources of decoherence for the length of time it takes to perform the computation [7, 8].

Qubits have been physically implemented in several physical systems, such as ion traps[9], nuclear spins[10], and photon cavities[11], but some of the most promising designs to date exploit the quantum behavior of electrons in a superconductor[12].

A major advantage of superconductor-based circuits over other types of qubits, such as ion traps or nuclear-spin devices, is that they can be fabricated in large numbers with high reliability using techniques similar to those developed by the semiconductor industry over the past 50 years. The two materials most commonly used are aluminum (Al) and niobium (Nb) [12], both superconductors that are compatible with standard silicon-based semiconductor processes and thus available in many experimental fabrication facilities.

Quantum-mechanical superposition of states in superconducting devices has been demonstrated only in Al devices to date[13-15], but Nb is an attractive alternative for several reasons. Its critical temperature is 9.2K, the highest of all the elements. This is well above 4.2K, the temperature of liquid helium, and allows basic testing and process qualification to be done relatively cheaply and easily in liquid-helium systems. By contrast, the T_C of Al is 1.18K, which means Al-based devices must be cooled using an expensive and complex He^3 system (~0. 3K) or dilution refrigerator (~0.01K) to examine their superconductive behavior. Nb is also similar to silicon in many respects from a fabrication standpoint; fluorine-based plasma etches and other well-known processes used for Si fabrication can be readily adapted to Nb[16].

1.1.2 Superconductivity and Josephson Junctions[17]

The phenomenon of superconductivity occurs when certain conductors are cooled below a material-dependent critical temperature (T_c) and begin to exhibit electronic conduction

without resistance. From a quantum-computing standpoint, this effect is relevant because the Cooper pairs (pairs of electrons formed inside a superconductor) in superconducting materials exhibit behavior that can be described by a quantum-mechanical wave function, one of the few instances in nature of quantum behavior on a macroscopic scale. As a result, quantum systems can be designed using superconductors that exhibit much wider degrees of design freedom than, for example, the energy levels of an atom, which are fixed by the laws of nature.

The key element in any superconductive quantum circuit is the Josephson junction, which is made up of two layers of superconductive material separated by a thin insulating barrier. Cooper pairs can tunnel through the barrier with a probability dependent only on the phase difference between the macroscopic quantum wave functions on the two sides of the junction, according to the equation:

$$J = J_c \sin \phi \qquad (1)$$

Where J is the current density in the junction, J_c is the critical current density of the junction, and ϕ is the gauge-invariant difference in phases across the junction barrier. The critical current density is dependent on, among other things, the thickness of the barrier, and this is the parameter generally used to engineer a specific J_c when fabricating Josephson junctions.

Josephson junctions have many interesting circuit applications. Some of the most useful of these include superconducting quantum interference devices (SQUIDs), which can measure magnetic fields as small as 10^{-15} Tesla[17], rapid single flux quantum (RSFQ) logic circuits, which allow digital logic to be implemented in superconducting materials, and, most importantly for the purposes of this discussion, several types of qubits.

1.1.3 Persistent-Current Qubit[12]

From a circuit design standpoint, a Josephson junction is characterized by its critical current (I_c) and intrinsic capacitance (C). When a junction is placed in a loop of superconducting material, a Josephson inductance L_J can be defined as well, where:

$$L_J = \Phi_0 / 2\pi I_c \qquad (2)$$

Where the fundamental constant $\Phi_0 = h/2e$, the superconducting flux quantum. The ratio of the Josephson inductance L_J to the loop inductance L_S is used to characterize Josephson circuits via the parameter $\Lambda = L_J / L_S$.

In circuits where $\Lambda >> 1$, the loop inductance can be ignored. Quantum mechanically, these types of circuits are characterized by two energy scales, the Josephson coupling energy E_j and the Coulombic energy for single charges E_c, where:

$$E_J = I_0 \Phi_0 / 2\pi \qquad (3)$$

$$E_c = e^2 / 2C \qquad (4)$$

When $\Lambda >> 1$ and $E_J > E_C$, the phase of the macroscopic wave function in the circuit is well-defined and the charge at a given node of the wave function fluctuates. When $E_J < E_C$, the charge at the nodes is well-defined and the phase fluctuates. When E_J and E_C are within a few orders of magnitude of one another, the tunneling states in the system become nontrivial and fluctuations in both charge and phase can be observed[18]. While quantum-mechanical behavior has been observed in circuits of this type in all three energy regimes[13-15, 18-22], fluctuating charges in the substrate generally represent a strong source of decoherence and make design of a controllable qubit difficult in this type of circuit[8, 23, 24].

When $\Lambda << 1$, the magnetic flux in the circuit becomes significant and the circuit's quantum variables can be related to the flux in the loop and its time derivative. The simplest circuit of this type consists of a superconducting loop with a single Josephson junction and is known as an rf SQUID. Macroscopic quantum states have been observed in these types of circuits[24-28], and superposition of states has been demonstrated as well[29]. A major advantage of this second type of circuit in quantum computing is that the two quantum mechanical states are represented as circulating currents of opposite signs, which generate a flux that is easily measured via a SQUID. These types of circuits are much more easily decoupled from their environment than charge-type devices as well, improving coherence times.

The persistent-current (PC) qubit is a compromise between the two approaches, attempting to combine the superior quantum-mechanical behavior of the first type of circuit with the circulating-current nature and improved coherence times of the second type[12]. The qubit consists of a loop of superconducting material with three Josephson junctions. The combined Josephson inductance of the junctions places the circuit in the $\Lambda >> 1$ regime, but when the respective areas of the three junctions are properly engineered, the two quantum-mechanical states in the circuit are represented by circulating currents of opposite signs, as in the rf SQUID.

The PC qubit satisfies all requirements for a functional quantum bit[12]. At a sufficiently low temperature, the persistent-current states are easily prepared in their ground state, as well as easily and precisely manipulated with magnetic fields using a SQUID or other device. Two qubits can be inductively coupled together, and the degree of inductive coupling can be varied. The state of the PC qubit is readily measured and manipulated using a SQUID. The circulating-current nature of the PC states makes decoupling the qubit from its environment much more straightforward than in charge-based Josephson circuits.

1.2 PC Qubit Fabrication

A major challenge in developing any process to fabricate Josephson junctions is the fact that the junctions do not tolerate high processing temperatures. The exact mechanism of this degradation is unknown, but is thought to be a combination of Nb-Al interdiffusion

at the barrier, oxygen diffusion at the surface, and changes in the composition of the barrier layers will begin to degrade the critical current density (J_C) and quality of the junctions at temperatures above about 150°C, and the degradation becomes more pronounced with increased temperature and heating time[30-32]. Semiconductor-based processes typically employ extremely high temperatures (400°C and above) to anneal defects out of materials and improve the quality and uniformity of materials, so alternatives to common silicon-processing techniques such as thermal oxidation, annealing, and high-temperature plasma etching must be used when designing a Josephson circuit fabrication process.

The "QC" process currently used to fabricate superconductive quantum circuits in the MIT Lincoln Laboratory Microelectronics Laboratory (MEL) is based on the doubly-planarized all-refractory technology for superconductive electronics (DPARTS) process developed at MIT-LL, based on the similar PARTS process from IBM[33]. It is Nb-based, with a thin layer of aluminum oxide (AlO_x) providing the barrier for Josephson junctions (niobium oxide can also be used as a tunnel barrier material, but its high magnetic permeability is thought to make it unsuitable for quantum-computing applications[34]). The process employs a junction-definition layer, three wiring layers, a platinum resistor layer, two via layers, and a gold contact layer, for a total of eight optical layers. The layers are defined using i-line (365 nm) projection lithography, with a resolution limit of roughly 0.75 μm. Thick layers of SiO_2, deposited at low temperature via plasma-enhanced chemical vapor deposition (PECVD) and planarized via chemical-mechanical

polishing (CMP), are used to isolate the various wiring layers from one another and provide a flat topography for each lithographic step.

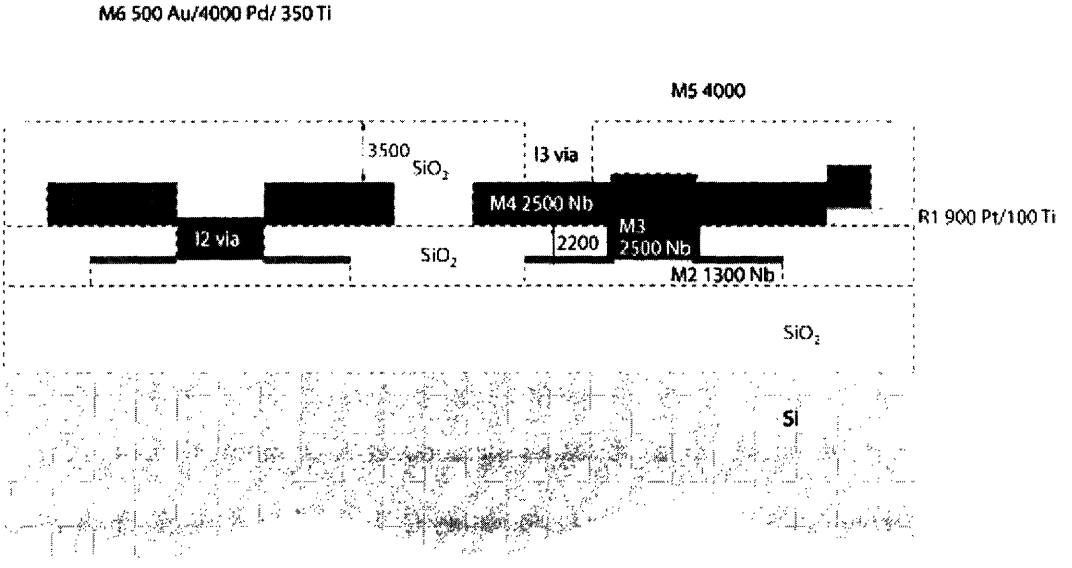

Fig. 1: Schematic cross-section of the Lincoln Laboratory "QC" Nb-based superconductive process, showing all layers.

Using devices fabricated with the QC process, our research group has demonstrated two-state quantum mechanical behavior in a three-junction persistent-current qubit circuit[35, 36]. The three wiring layers and platinum resistor layer in the QC process allow for implementation of a wide variety of circuit designs, such as RSFQ control circuitry and complex multi-qubit coupling experiments.

1.3 Motivation for New Process

While robust and well-developed, the QC process suffers from two major limitations. The turnaround time to fabricate a run of devices in the LL MEL is approximately four months, making short loops for process-qualification structures and other simple devices impractical.

Additionally, the smallest junctions that can be fabricated using the DPARTS-based process are roughly 0.75 μm in diameter. While this is acceptable for quantum-computing applications, the ability to fabricate smaller junctions would allow the E_J/E_C ratio discussed in Chapter 1.1 to be decreased, allowing much more latitude in the qubit experiments that can be performed using circuits fabricated at Lincoln Laboratory.

With the goal of addressing these two shortcomings simultaneously, a new process for fabricating Nb-based Josephson circuits at Lincoln Laboratory was proposed. By eliminating the third wiring layer and time-consuming CMP processing, the approximate fabrication time was reduced from four months to ten days (hence the name of the new process, TD). Additionally, the i-line lithography was replaced with deep-UV (248 nm), with the goal of improving resolution to 0.2-0.3 μm.

Fig. 2: Schematic chart of QC streamlining into TD process, illustrating removal of the most time-consuming steps and R1/M5/I3 layers.

The removal of the CMP steps meant that a new dielectric material had to be used to insulate the wiring layers from one another. Since it has been shown that NbO_x grown via anodization is sufficient for this insulative layer[37], the oxide deposition and planarization steps were replaced by anodization steps. The extensive in-process

metrology of the QC process was also removed, as troubleshooting a process with only a ten-day turnaround is most easily accomplished by simply beginning a new run.

The loss of the upper wiring (M5) and resistor (R1) layers reduced the number of optical layers in the TD process to five, but also limited the complexity of circuits that could be produced with the process relative to the QC process. Since the goal of the TD process was to rapidly produce simple qubit circuits for qualification and evaluation, this was an acceptable tradeoff; if necessary, additional wiring and resistor layers can be added to the process at in the future.

The key challenges faced in developing the TD process were as follows:

- Design a stable process for performing high-resolution 248 nm photolithography on Nb.
- Design a selective anodization process capable of resolving dimensions as small as 0.3 μm
- Improve selectivity and directionality of all etch steps used in the QC process to allow for smaller feature resolution
- Fully isolate the M2 layer from the M4 layer without using the PECVD/CMP approach in the QC process

Under these new constraints, a process was developed allowing both rapid turnaround time and resolution of roughly 0.3 μm, with further resolution improvements possible.

1.4 Thesis Outline

An overview of the process flow used to do this, including an outline, diagrams, and brief notes on the evolution of key steps, appears in chapter 2 and is organized by masking level. More details on specific processing steps and the challenges and issues they presented are organized by process step (photolithography, etch, etc.) in chapter 3. In chapter 4, the primary results of the work are summarized and discussed; namely, the success of fabrication and test data obtained from devices produced using the TD1 process. Chapter 5 details further work to be done and addresses the problems still present in the TD1 process. An appendix with an overview of the response surface design (RSD) method used to optimize several of the plasma etch steps used in TD1 is also provided.

2 Process Flow

2.1 Overview

2.1.1 Process Outline

The TD1 process flow is the result of research, experimentation, and refinement, and as a result has undergone many incarnations as new data were gathered and analyzed. In brief, its final form is presented in list format below for ease of reference; each step is discussed in detail in section 2.2:

1. Generate Nb-Al/AlO_x-Nb trilayer with Al under-layer via sputter deposition
2. Pattern M3 (junction definition) layer via photolithography
3. Etch M3 layer and Al/AlO_x barrier via plasma etch
4. Anodize M3 layer
5. Remove M3 mask
6. Pattern M2 (lower wiring) layer
7. Etch M2 layer down to Al under-layer
8. Anodize M2 sidewalls and Al under-layer
9. Remove M2 mask
10. Pattern I2 (via) layer
11. Etch through NbO_x layer, stopping on M2 Nb
12. Remove I2 mask

13. Deposit M4 (upper wiring) layer via sputter deposition

14. Pattern M4 layer

15. Etch M4 layer, stopping on NbO_x and AlO_x

16. Remove M4 mask

17. Pattern M6 (contact) layer

18. Deposit Ti-Au contacts via metallization and liftoff

2.1.2 Process Description

The first step in the process is the buildup of the trilayer [33]. This is done using an MRC dc magnetron sputtering system to deposit successive layers of Al and Nb. The Al "barrier" layer is slightly oxidized *in situ* to produce a thin AlO_x Josephson tunnel barrier. The extent of this oxidation determines the current density, which can be varied between 100 and 500 A/cm^2.

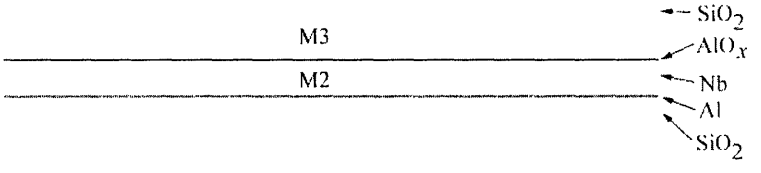

Fig. 3: Initial TD1 material stack, with all film types labeled.

The first lithographic layer (M3) is the most critical, since it defines the junction areas. The M3 layer must be patterned, etched, and anodized at the highest possible resolution. Selective anodization at high resolutions is difficult, so a thin SiO_2 adhesion layer is

deposited prior to M3 processing (see Chapter 3.3). The wafer is patterned with DUV resist and the SiO_2, Nb, and Al/AlO_x barrier layer are etched away using various wet and plasma etches. The exposed Nb is then anodized to produce a 50 nm film of NbO_x before the mask and adhesion layer are removed.

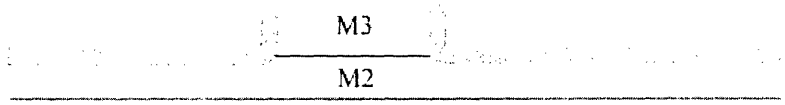

Fig. 4: Schematic cross-section of TD1 after completion of M3 processing.

The resolution requirements for the M2 layer are less severe, but an adhesion layer is still necessary to retain the resist pattern during anodization. The adhesion layer is deposited as before, and the M2 layer is patterned and etched down to the Al under-layer. The M2 sidewalls and Al under-layer are then anodized to completely isolate M2 with a dielectric film, and the mask and adhesion layer are removed.

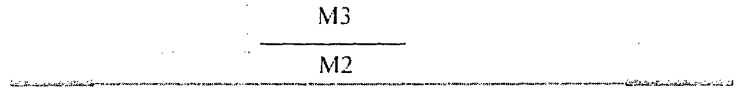

Fig. 5: Schematic cross-section of TD1 after completion of M2 processing.

In order to form contacts between the upper (M4) and lower (M2) wiring layers, vias through the insulating NbO_x layer are required. These are contained within the I2 layer, which is patterned and etched to remove the dielectric film in desired regions. The vias are generally large (~1 μm or greater), so resolution is not important in this step. No adhesion layer is necessary, and the resist mask is removed as usual after completion of the etch.

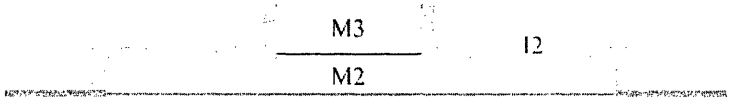

Fig. 6: Schematic cross-section of TD1 after completion of I2 processing

The M4 (upper wiring) layer is then deposited across the entire wafer using the same dc magnetron sputtering system used to build up the initial material stack. After deposition, the M4 layer is patterned with resist and removed selectively down to the NbO_x and AlO_x dielectric layers. Again, no adhesion layer is required since no anodization is used in this step.

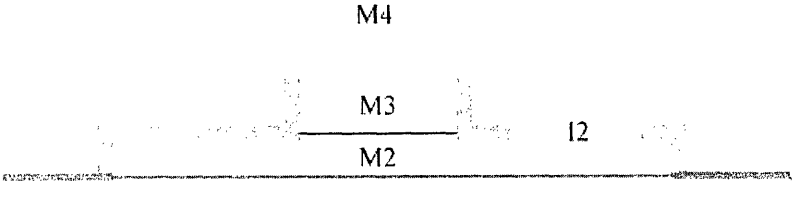

Fig. 7: Schematic cross-section of TD1 after completion of M4 processing

The final lithographic step (M6) is used to deposit gold contacts on top of M4. The wafer is patterned with resist and subjected to a surface-poisoning procedure to enhance sidewall overhang (see Chapter 3.5). Gold contacts are deposited on the resist-patterned wafer via electron-beam evaporation and the pattern is transferred via liftoff in an ultrasonic acetone bath.

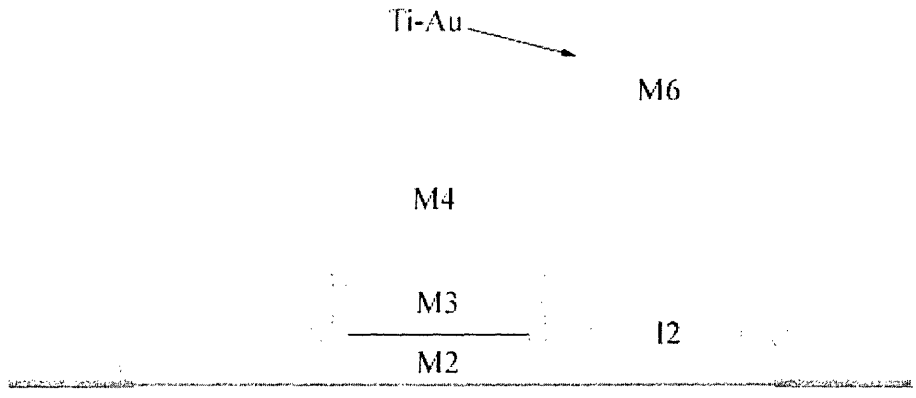

Fig. 8: Schematic cross-section of TD1 after all processing is complete

2.2 Process Details

2.2.1 Initial Material Generation

1. Generate Nb-Al/AlO_x-Nb trilayer with Al under-layer and LTO adhesion film
   - 1.1. Base material: p-type Si wafer
   - 1.2. Oxidize Si wafer to form a 200 nm substrate film of high-quality thermal SiO_2

NOTE: Steps 1.3 through 1.6 occur together in situ, using the same dc magnetron sputter-deposition tool

1.3. Deposit a 15 nm layer of Al to act as an etch stop for M2 and a conductor for anodization

1.4. Deposit 150 nm of Nb for the M2 layer

1.5. Deposit 6 nm Al and form a very thin oxide film on its surface via *in situ* oxidation

1.6. Deposit 150 nm of Nb for the M3 layer

1.7. Deposit 30 nm low-temperature SiO_2 (LTO) adhesion layer film via PECVD at 150°C

The adhesion layer was not originally used, but proved necessary for selective anodization at the resolutions required by TD1 (see chapter 3.3).

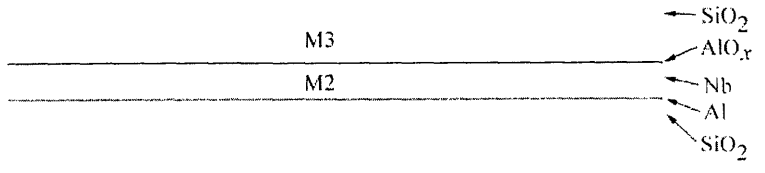

Fig. 9: Initial TD1 material stack, with all film types labeled.

2.2.2 M3 Processing

2. Pattern the M3 layer using 395 nm of Shipley UV5 248 nm photoresist

A layer of anti-reflective coating (ARC) was originally used here as well to improve the photolithographic resolution, but was abandoned due to issues with removing it (see chapters 3.1 and 3.4).

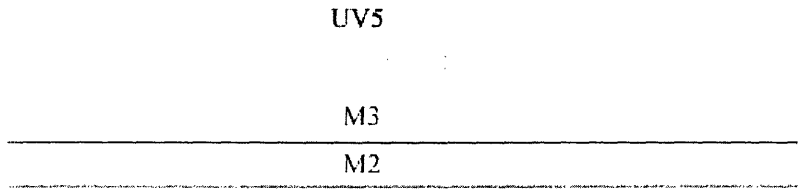

Fig. 10: Schematic cross-section of TD1 after M3 patterning

3. Etch the M3 layer and Al/AlO_x barrier

   3.1. Etch the LTO adhesion layer down to the underlying M3 Nb using a timed CF_4 plasma etch.

   3.2. Etch the M3 Nb layer using a SF_6 plasma etch, stopping on the AlO_x barrier.

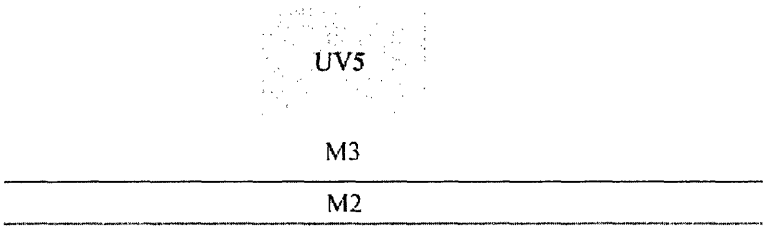

Fig. 11: Schematic cross-section of TD1 after M3 LTO and Nb etch

3.3. Etch the Al/AlO_x barrier layer using a 60 sec. automated immersion in TMAH-based photoresist developer

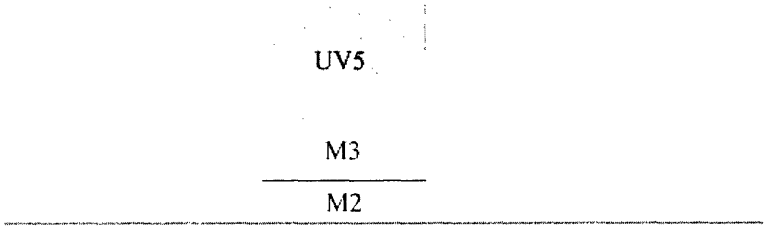

Fig. 12: Schematic cross-section of TD1 after M3 barrier etch

4. Anodize the exposed Nb to form a film of NbO_x approximately 50 nm thick.

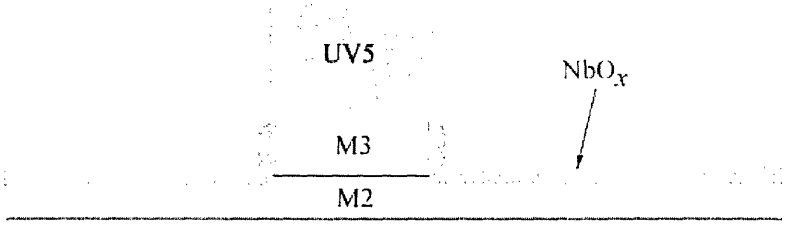

Fig. 13: Schematic cross-section of TD1 after M3 anodization

5. Remove the M3 mask

5.1. Remove the resist using a combination of mild plasma stripping and ALEG310 solvent.

A much harsher solvent was originally used for resist stripping in TD1, but was discovered to be attacking the Al in the barrier and under-layer and was replaced. See chapter 3.4.

5.2. Remove the LTO adhesion layer using a short silox (NH_4F-CH_3COOH) etch

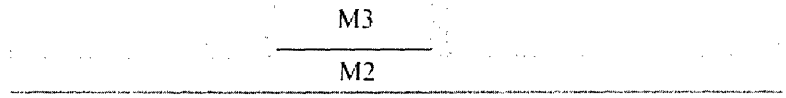

Fig. 14: Schematic cross-section of TD1 after completion of M3 processing.

2.2.3 M2 Processing

6. Pattern lower wiring (M2) layer

6.1. Deposit 30 nm LTO adhesion layer via PECVD, identical to that used in M3

6.2. Pattern the M2 layer using 395 nm of UV5 resist

7. Etch the M2 layer down to the Al under-layer

7.1. Etch the LTO layer down to the underlying NbO_x using a timed CF_4 plasma etch

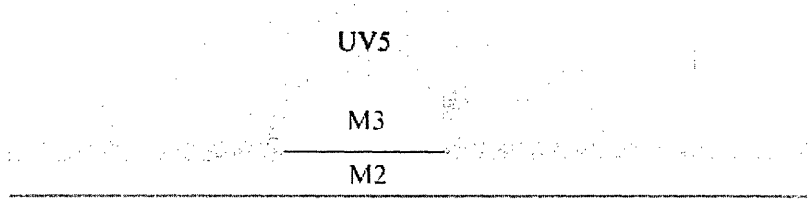

Fig. 15: Schematic cross-section of TD1 after M2 patterning and LTO etch

7.2. Etch through the NbO_x and Nb layers using a high-temperature SF_6 plasma etch, stopping on the underlying Al layer.

The etch recipe used here is identical to that used in the M3 and M4 Nb etches, with the temperature increased from 50°C to 80°C. See Chapter 3.2 for more details on this.

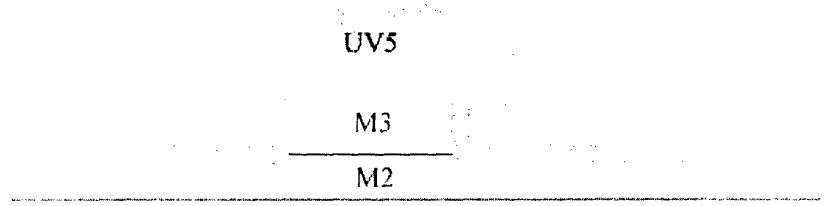

Fig. 16: Schematic cross-section of TD1 after M2 Nb etch

Originally, isolation of the M2 layer was done by selectively anodizing it down to the substrate, rather than etching it away. This process avoided the necessity of both the Al under-layer and M3 barrier etch (which had not been developed at the time), but was abandoned due to the instability of resist when anodized at voltages higher than about 75V. Even with an adhesion layer present, nearly all resist lifted (see chapter 3.3).

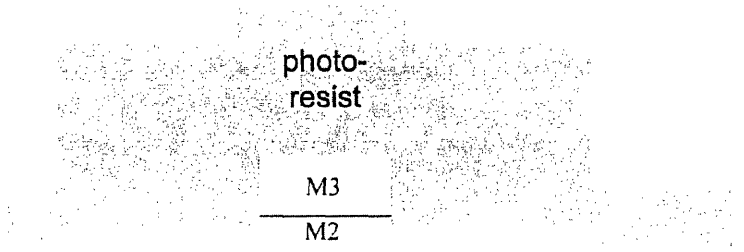

Fig. 17: Schematic of abandoned plan to isolate M2 features via full anodization of the layer.

8. Anodize the exposed Nb sidewalls and Al under-layer, producing an AlO_x substrate and isolating the M2 sidewalls with an insulating dielectric.

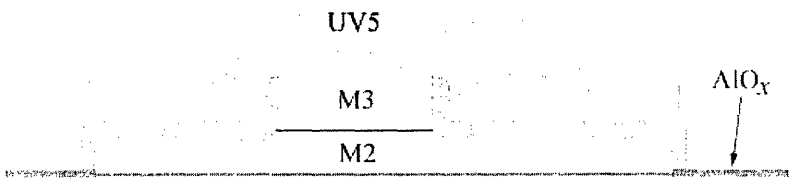

Fig. 18: Schematic cross-section of TD1 after M3 anodization

9. Remove the M2 mask

   9.1. Strip the resist using the wet/dry combination recipe used in M3 (see chapter 3.4 for more details)

   9.2. Strip the LTO adhesion layer using a short CF_4 plasma etch step

Silox cannot be used here, as it causes severe undercut and liftoff of all M2 features smaller than approximately 1 μm. See chapter 3.4 for more details on this.

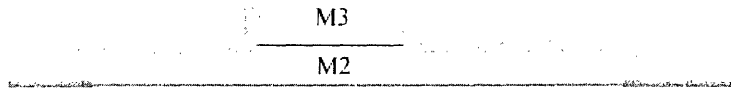

Fig. 19: Schematic cross-section of TD1 after completion of M2 processing.

2.2.4 I2 Processing

10. Pattern the I2 layer with 395 nm of UV5 resist

Since this is a via layer, the mask is reverse-tone; that is, the features are clear and the field is dark.

11. Etch through the NbO_x dielectric layer using a high-temperature CF_4 plasma etch

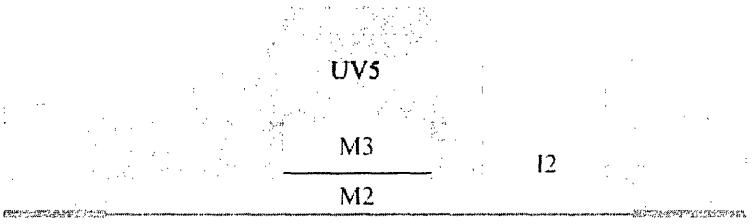

Fig. 20: Schematic cross-section of TD1 after I2 patterning and etch

12. Strip the resist as before, using the mild wet/dry recipe developed for the TD1 process.

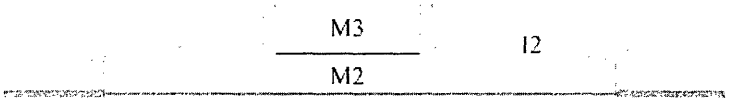

Fig. 21: Schematic cross-section of TD1 after completion of I2 processing

2.2.5 M4 Processing

13. Deposit 200 nm of Nb via dc magnetron sputtering

The M4 layer MUST be thicker than both the M2 and M3 layers, as it has to remain continuous when stepping over both of them.

14. Pattern the M4 layer with 395 nm of UV5 resist

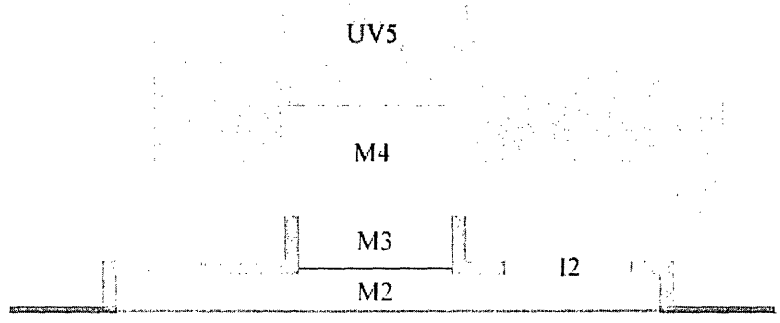

Fig. 22: Schematic cross-section of TD1 after M4 deposition and patterning. Note the self-aligned contact formed between M3 and M4.

15. Etch the M4 layer using a SF_6 plasma etch, stopping on the AlO_x and NbO_x films.

16. Strip the photoresist using the TD1 wet/dry/wet recipe.

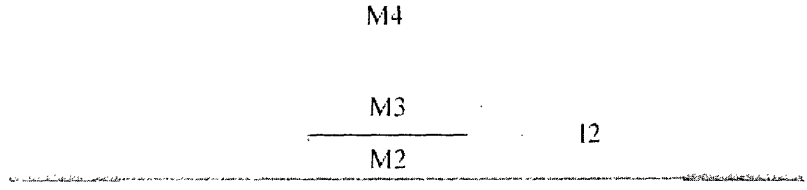

Fig. 23: Schematic cross-section of TD1 after completion of M4 processing

2.2.6 M6 Processing

For more details on the surface-poisoning procedure used to engineer the resist profile at this step, see chapter 3.5.

17. Pattern the M6 contact layer

    17.1. Coat the wafer with 1.5 μm of Shipley UV26 248 nm resist.

    17.2. Immerse the coated wafer in a solution of 1% developer for 1 min.

    17.3. Expose and develop the M6 pattern

18. Deposit contacts via metallization/liftoff

    18.1. Deposit ~50 nm of Ti and ~400 nm of Au on the wafer via electron-beam metal evaporation.

    18.2. Immerse the wafer in an ultrasonic acetone bath to dissolve the resist and complete contact pattern liftoff.

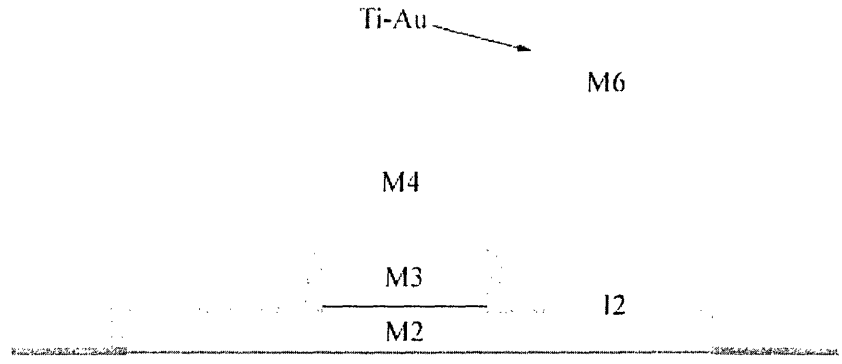

Fig. 24: Schematic cross-section of TD1 after all processing is complete

3 Process Design

3.1 Photolithography

3.1.1 Overview

The TD1 process is the first superconductive device process to incorporate 248 nm (deep-UV) photolithography to define patterns. 248-nm lithography is similar to the more familiar i- and g-line lithographic processes, in that it uses light and a film of photoresist to transfer a pattern from a photomask to the wafer, but differs in several key respects[38].

The light source used in 248-nm lithography is typically a KrF excimer laser, as the Hg lamps used in more traditional lithography setups exhibit generally low power transmission in the DUV regime. The KrF source is typically of much lower intensity than conventional lamp sources, so the photoresist used in 248-nm lithography must be correspondingly more sensitive than standard i-line resists. A class of photosensitive compounds known as chemically-amplified (or acid-catalyzed) resists exhibit both high sensitivity and high transmittance in the 248 nm regime, and so comprise the vast majority of commercially-available DUV resists, including the Shipley Microelectronics UV5 resist used in the TD1 photolithography[39].

The basic mechanism of chemically-amplified resists is comprised of two major steps[39-41]. In the first step, an incident photon strikes an acid-generator molecule that produces

an acid upon exposure to radiation. In the second, the acid catalyzes reactions with acid-labile polymers in the resist to change their solubility in the developer solution. Since a single photogenerated acid molecule can catalyze a large number of development reactions, the sensitivity of chemically-amplified resist is extremely high. Moreover, since acid generation is the only photochemical event occurring in the resist, the acid-labile base resist material can be designed with high transmittance at 248 nm, ensuring uniform exposure throughout the film thickness. The specific makeup of chemically-amplified resists varies and is usually proprietary, but the many possible combinations of acid generator molecules and acid-labile polymers produce generally similar effects.

While the photochemical acid-generation step occurs readily during exposure, the secondary acid-catalyzed reactions that differentiate exposed regions from unexposed regions in chemically-amplified resist generally do not occur to a significant degree at room temperature. As a result, the critical dimensions (CDs) of the photomask pattern are highly dependent on a post-exposure/pre-development bake. An insufficiently long or hot bake can result in irregular sidewall patterns, while overbaking causes severe critical-dimension reduction and may completely remove smaller features.

In order to further stabilize the resist after development, a post-development UV flood exposure may be used[42]. This flood exposure is performed simply by exposing the wafer to UV light at post-exposure bake temperatures, forming crosslinks in the polymers of the resist film and generally making it less susceptible to swelling, flowing, and liftoff during high-temperature plasma etching or harsh chemical processing.

3.1.2 Implementation

Nb devices having never before been fabricated using 248 nm photolithography, it was necessary to develop a process to optimize yield and allow resolution of features as small as possible, as well as determine the feasibility of the deep-UV/Nb combination itself.

The MIT Lincoln Laboratory Microelectronics Laboratory (MEL) uses a Canon FPA-3000 DUV stepper for 248-nm lithography work. The tool has a practical feature resolution limit of roughly 200 nm on silicon, but the fact that Nb is significantly (~25%) more reflective than silicon means that the minimum achievable printed feature size of the TD1 process will probably be significantly larger due to reflection at the substrate[41] (fig. 25).

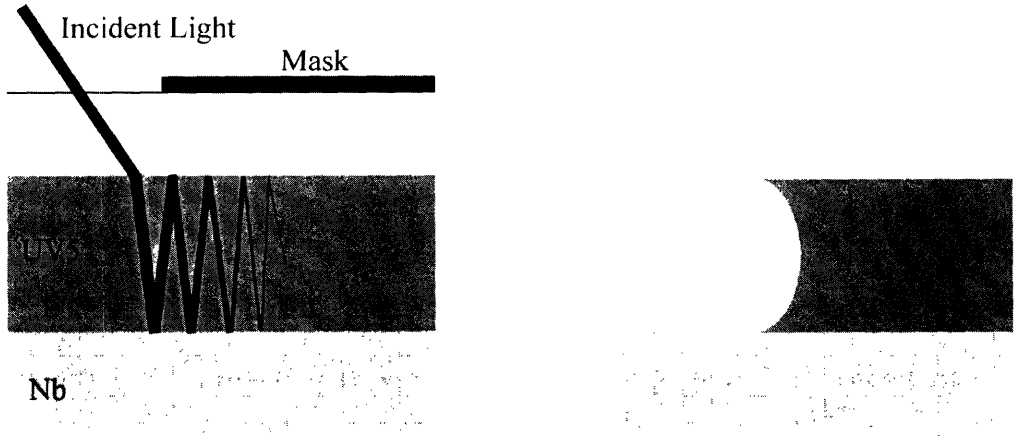

Fig. 25: Effect of reflection from substrate on photolithography. A standing wave is set up as incident light reflects between the resist-Nb and resist-air interfaces during exposure (left), resulting in an unpredictable resist sidewall pattern on development (right).

3.1.3 Characterization and Optimization

When optimizing optical projection photolithography (where the mask pattern is projected through a system of lenses onto the substrate, rather than placed directly on it as in contact lithography) with a 248-nm resist system, the three most important parameters are post-exposure bake, exposure dose, and focal plane location (focus). The post-exposure bake is the most critical, as it largely determines the extent to which the resist-exposure reaction takes place and thus has a marked impact on feature sizes and sidewall profiles. The total amount of radiation the wafer is exposed to (known as exposure dose and typically specified in J/m^2) is also important, as it controls the amount of acid catalyst produced. For optimum resolution, it is also necessary for the mask image to be in focus on the wafer plane. Variations in the distance between the focal plane of the lens system and the wafer surface can also be used to control sidewall angle, which can help compensate for effects such as substrate reflection.

The post-exposure bake is a well-established parameter, with the Lincoln Laboratory DUV process utilizing a bake of 130°C for 90 seconds. No attempt was made to alter this part of the process, as it is not easily modified and the standard values give a good compromise between sidewall uniformity and critical dimension size. The focal length and exposure energy need to be concurrently optimized in order to yield the best combination of sidewall profile and feature size for the critical feature of the mask layer being exposed, however. This could in principle be done by computer simulation, but

given the number of real-world variables involved is most readily accomplished via direct experimentation.

The experiment performed to optimize the TD1 photolithography step was a simple, standard characterization procedure known as a focus/exposure array. A blank silicon wafer was coated with Nb on one side, covered in 395 nm of Shipley UV5 248-nm photoresist, and placed in the stepper. The stepper was then configured to pattern 49 reticles on the wafer in a 7x7 grid, varying the focus in one direction and the exposure energy in the other. The limits used for exposure energy were 50 J/m^2 - 140 J/m^2 (in increments of 15 J/m^2); the focus was varied from +0.3 μm to -0.2 μm. The wafer was then developed and examined with a scanning electron microscope (SEM). The result of each combination is shown in the table below (table 1). Since the most critical photolithography step in the process is the M3 exposure (as it defines the junction area), the lithography was optimized for small, widely-spaced posts similar to the features present on the actual M3 mask.

		Exposure (J/m^2)						
Structure	**Focus (μm)**	**50**	**65**	**80**	**95**	**110**	**125**	**140**
0.4 μm post	**0.3**							
0.3 μm post	**0.3**							
0.4 μm post	**0.2**		0.68		0.5		0.44	
0.3 μm post	**0.2**		0.52		0.36		0.31	
0.4 μm post	**0**		0.63		0.47		0.41	0.38
0.3 μm post	**0**		0.49		0.35		0.28	0.22
0.4 μm post	**-0.2**		0.69		0.47		0.4	
0.3 μm post	**-0.2**		0.51		0.33		0.28	

Table 1: Actual yielded feature sizes of 0.3 μm and 0.4 μm posts at various focus/exposure settings for 395 nm of UV5 resist on bare Nb. A blank field indicates that the feature did not yield or was not examined (all dimensions in μm).

According to the table, the printed dimensions are closest to the mask dimensions at –0.2 μm focus and 125 J/m^2 exposure energy. However, the table fails to take sidewall angles into account; a sloped sidewall profile can create serious problems in plasma etching and other process steps. Characterizing sidewall angle is somewhat more difficult; the easiest method is to simply tilt the sample in the SEM and gauge the profile "by eye". SEM images for several different focus/exposure values are shown below (fig. 26).

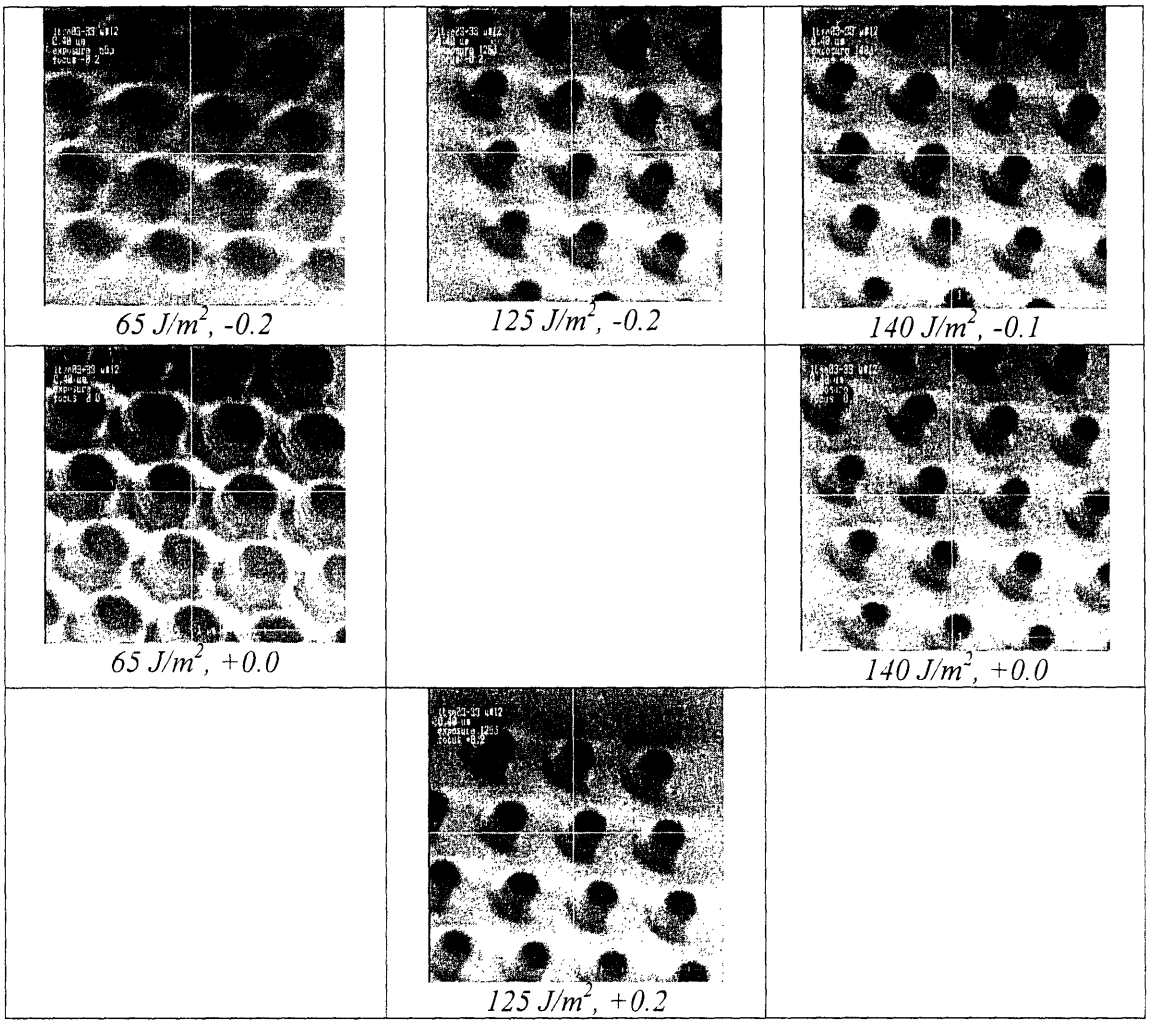

Fig. 26: Scanning electron micrographs (tilted 45 degrees) of 0.4 μm post sidewall profiles at various focus/exposure settings. Note the poor resolution at 65 J/m² (left), slanted profile at 125 J/m²/-0.2 (top center), and severe size reduction at 140 J/m² (right). Unfortunately, the 125J/m²/+0.0 image was not available; sidewalls at this setting were unacceptably slanted, however. Also note the standing-wave pattern due to the previously discussed substrate reflection, easily visible at 65J/m²/+0.0.

The reticle exposed using a dose of 140 J/m² and a focus of +0.0 μm showed the best combination of sidewall profile and critical-dimension size for 0.4 μm posts.

3.1.4 Further Optimization - ARC

Unfortunately, yielding 0.3 μm posts over Nb proved problematic, as the increasing effects of substrate reflection made producing acceptable sidewall profiles impossible. To work around this issue and improve our minimum junction size, an anti-reflective coating (ARC) material was used, in this case Shipley AR3.

Back-side anti-reflective coating (BARC) is a material deposited under a photoresist layer to improve resolution by drastically limiting substrate reflection[43]. It has three key features: a refractive index matched to that of compatible photoresists to provide near-total transmission of 248 nm light at the ARC-resist interface; a high extinction coefficient for 248 nm photons; and a thickness matched to $n\lambda/4$, where n is the refractive index and λ the wavelength of incident light, to ensure that waves trapped between the substrate-ARC and ARC-resist interfaces interfere destructively. The result of this is that any photons reaching the ARC-resist interface are transmitted into the ARC material and quickly absorbed, rather than reflecting back from the substrate and creating a standing wave in the resist pattern (fig. 27).

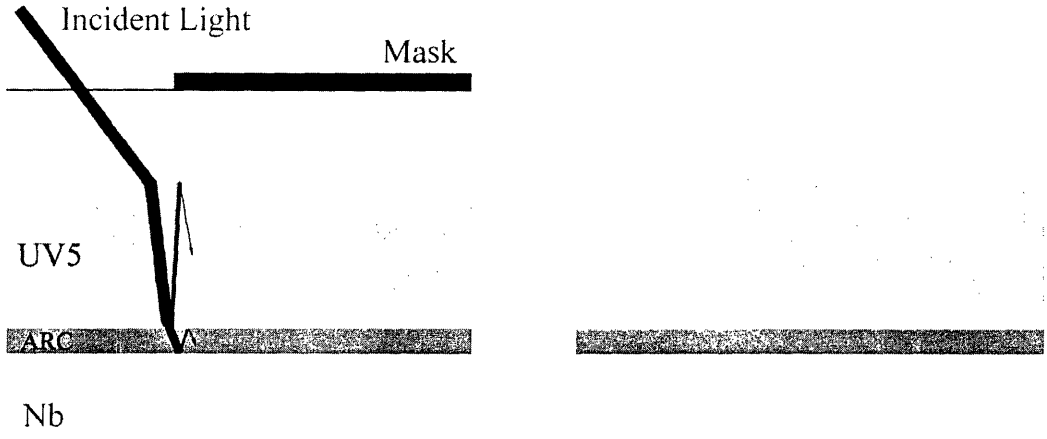

Fig. 27: Effect of ARC in reducing substrate reflection. Light striking the resist-ARC interface during exposure (left) is drastically attenuated by the properties of the ARC material, greatly reducing undercut due to reflection (right).

ARC application in the TD1 process posed some difficulties, as the standard application procedure requires it to be baked at a temperature of 215°C for 60 seconds after application to drive out solvents and stabilize the film. The trilayers used in the TD1 process exhibit damage at temperatures higher than about 150°C, however, so the standard ARC application procedure was incompatible with TD1. Fortunately, application with a lower bake temperature and a longer bake time (160°C for 180 seconds) caused no major problems with either the ARC stability. Past experiments at MIT-LL suggest that the trilayer damage from exposure to such temperatures would be minimal, despite its technical violation of the 150°C processing limit.

The addition of ARC to the process also required an extra etch step to be added after the photolithography to pattern the ARC. This was accomplished using a 15-second CF_4/O_2 plasma etch in the MEL's Trikon etch tool and caused no additional complications in the process.

With the addition of ARC to the process, lithographic resolutions greatly improved. Another focus/exposure matrix was performed on a wafer coated with 62 nm of ARC and 395 nm of UV5 resist, and 0.3 μm posts yielded easily (table 2).

		Exposure (J/m^2)						
Structure	**Focus (μm)**	**60**	**80**	**100**	**120**	**140**	**160**	**180**
0.4 μm post	**0.3**				0.33			
0.3 μm post	**0.3**							
0.4 μm post	**0.2**			0.52	0.44	0.43	0.37	
0.3 μm post	**0.2**			0.42	0.34	0.31	0.29	
0.4 μm post	**0**		0.56	0.51	0.42	0.37	0.35	
0.3 μm post	**0**		0.43	0.37	0.28	0.22	0.21	
0.4 μm post	**-0.2**						0.28	
0.3 μm post	**-0.2**							

Table 2: Actual yielded feature sizes at various focus and exposure values when a62nm ARC layer is used between the Nb and UV5 resist. 0.3 μm posts yielded at a good range of values, and even 0.25 μm is achievable. Blank fields indicate that the given feature did not yield or was not examined.

The printed dimensions at a given focus/exposure are slightly higher with ARC than without, which is expected as the ARC is absorbing light that would otherwise expose more of the resist. The best agreement with mask dimensions here occurs at a slightly higher energy than before as a result; the closest match is at $140J/m^2$ and +0.2 μm.

Sidewall profiles were similarly improved, as the footing and undercut seen on the resist-only wafers were reduced or eliminated with the addition of ARC (fig. 28).

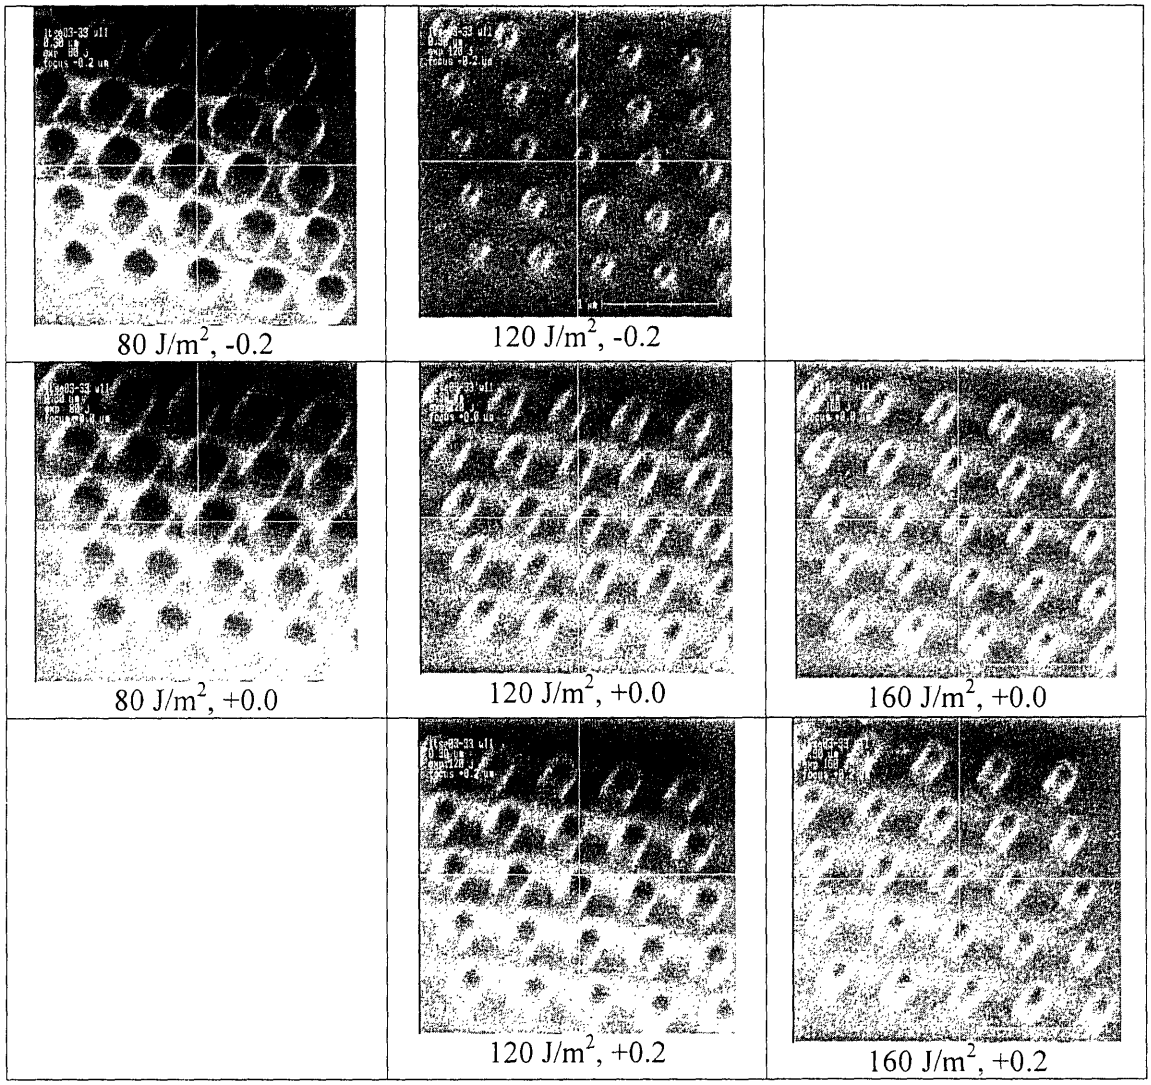

Fig. 28: Scanning electron micrographs (tilted 45 degrees) of 0.3 μm post sidewall profiles at various focus/exposure settings. Note that the sidewalls are nearly vertical in all cases, though feature size is still dependent on exposure energy to a large extent.

The optimal focus/exposure combination with ARC is clearly 120J/m^2 and +0.2; the posts on this reticle exhibited almost perfect sidewall profiles in addition to good agreement with the mask dimensions.

Unfortunately, issues with its removal (see Chapter 3.4) prevented the use of ARC in the current incarnation of the TD1 process, effectively limiting printed dimensions to 400 nm. Work is being done to alleviate this, with experiments focusing on using a high-pressure CF_4/O_2 plasma recipe to strip away the ARC, but no solution has yet been found.

3.2 Etch

3.2.1 Overview

With the exception of the gold contact layer (discussed in chapter 3.5), the patterned images of the different layers in the TD1 process are transferred to the material via selective etching. In selective etching, the photoresist-patterned wafer is exposed to an agent that attacks the material not covered by resist while (ideally) preserving the covered regions. The agent can be either a chemical bath or a plasma, though most modern microelectronics processes use exclusively plasma etching[16]. In this way, a photoresist pattern can be translated into physical topography on a wafer.

In designing an etch process, there are several important parameters to consider. Etch resolution, obviously, is key; the process must be able to resolve the smallest printed lithographic features. Resolution was the main motivation for the industry-wide switch from wet to "dry" (plasma) etching; while wet etches are isotropic (they remove material in all directions at an equal rate), plasma etches can be designed with a high degree of

directionality, reducing the lateral etching that can destroy small features and introduce undesirable size reduction on larger ones.

Fig. 29: Schematic diagram of isotropic (left) vs. anisotropic etching. Note that the isotropic etch has undercut and distorted the feature much more than the directional etch.

Selectivity is also a major concern. All etch processes need to be stopped after a certain point, generally at an interface between two materials, and some method is needed to know when this transition occurs. This is most simply done by choosing an etch chemistry that etches the top material much more quickly than the underlying material, measuring the etch rate of the top material, and calculating the time it will take to etch to the interface.

This process is known as a timed etch, and is widely used but suffers from some disadvantages. No plasma-etching tool or wet etching bath is completely uniform, and some regions of the wafer will be etched to the interface sooner than others. As a result, an additional "overetch time" on top of the calculated etch time is always required to ensure that the etched material is cleared all over the wafer; this time is generally tool-

dependent, but a widely-used rule-of-thumb is that 25% of the calculated etch time is a suitable overetch. The amount of etching of the underlying material due to this overetch is dependent on the etch selectivity; that is, the ratio of the etch rates of the top and bottom materials. A chemistry with a high selectivity between the two materials (at least a 3:1 ratio is a practical value) is thus required to minimize substrate damage. In addition, a timed-etch recipe assumes that the etch rate of a material in a given etch tool remains constant over time. In a high-traffic research laboratory this can be a dangerous assumption, and periodic re-characterization is always necessary.

More sophisticated plasma-etching tools can also use endpoint detection, which uses a spectrometer, interferometer, or other tools to examine the system and look for evidence that the etch has reached the material interface. This is preferable to a timed etch because it does not rely on a known etch rate and is less dependent on a high selectivity; the operator or an automated program can simply stop the etch when the endpoint is detected, though inherent nonuniformities in the tool still make a high selectivity desirable.

3.2.2 TD1 Etch Summary

The TD1 process contains seven selective etch steps (eight if ARC is used), each of which presents its own unique challenges. These are as follows, organized by lithographic layer (see Chapter 2 for details of process flow):

- M3 layer
    - ARC etch (if used)
    - LTO adhesion layer etch
    - Nb etch
    - Al/AlO_x barrier etch
- M2 layer
    - LTO adhesion layer etch
    - Nb/NbO_x etch
- I2 layer
    - NbO_x etch
- M4 layer
    - Nb etch

The requirements, chemistries, etch rates, and selectivities for the etch steps in each layer follow.

3.2.3 M3 Etches

The M3 stack was the most complex to etch, with three separate materials to be etched and an additional layer of ARC to be removed if it was used to improve the photolithography. Unfortunately, ARC has proven incompatible with the TD1 process so far (see Chapter 3.4) and so was not used in any of the "production" runs of TD1.

Fig. 30: Schematic of TD1 before (left) and after all M3 etching. The LTO, Nb, and Al/AlO_x layers have been removed. ARC is not present in this schematic.

LTO adhesion layer

The LTO adhesion layer is an extremely thin (30 nm), low-density layer of SiO_2 designed to improve resist adhesion during anodization (see Chapter 3.3). It is easily etched in fluorine-based plasma chemistries, but without good selectivity to the underlying Nb. Since the underlying Nb is being etched in the next step (and, as will be seen, stops cleanly on the AlO_x barrier), this is not a major issue. The standard CF_4 etch in Lincoln Laboratory's Lam Rainbow plasma etch tool was found to etch LTO at approximately 10 nm/sec. Taking plasma-stabilization time

and necessary overetch into account, a 5-second plasma etch was used to successfully remove the LTO prior to the Nb etch.

Nb counter-electrode

To define the M3 features, the 150 nm Nb film on top of the Al/AlO_x barrier needs to be selectively etched away. Nb is readily etched in fluorine chemistry as well, and an SF_6-based recipe has been developed[33] (prior to TD1) on Lincoln Laboratory's Plasma-therm tool to etch Nb with near-vertical profiles via formation of sidewall polymers. This etch is simplified even further because Al and its oxides are not etched in fluorine-based chemistries at all, causing the etch to stop when the Al/AlO_x barrier is reached. A spectrometer in the plasma-therm tool monitors the fluorine radiative line in the plasma, which jumps sharply when the Al is reached and alerts the operator to stop the etch manually. The etch rate of Nb in this recipe is about 1 nm/sec, so the entire step takes approximately 3 min.

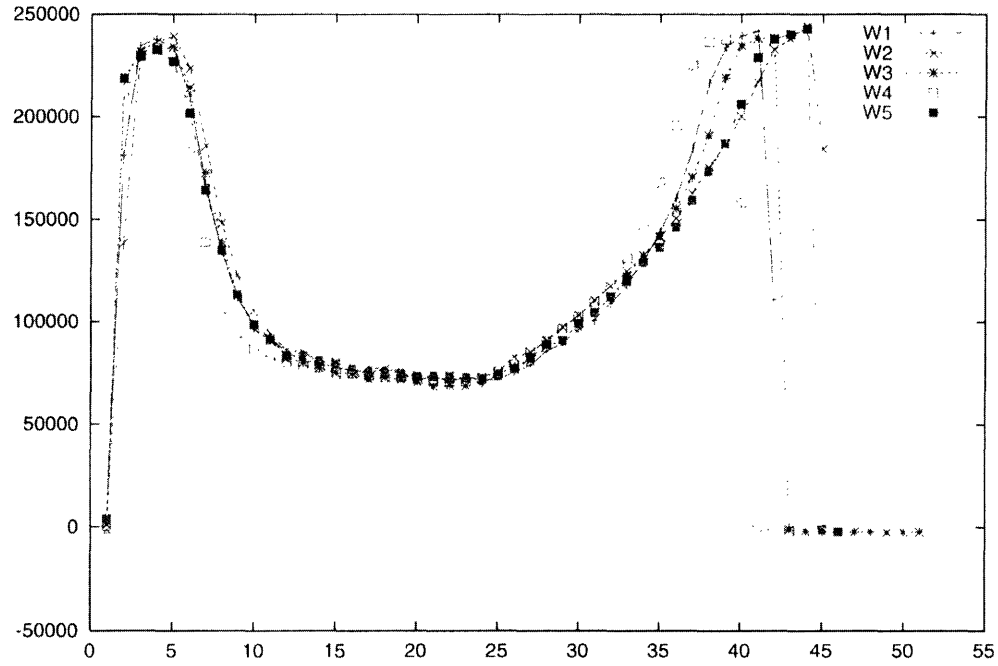

Fig 31: Data output from Procvis® spectrometer monitor on Plasma-Therm etch tool for several wafers. The fluorine line is monitored at set intervals during the process; the rapid rise at the end of the etch corresponds to the endpoint.

Al/AlO_x barrier

While the fact that fluorine etches will remove Nb and not Al is useful in many parts of the TD1 process, it can also cause problems. If the barrier is left in place during the M3 anodization step, the dielectric formed will be an composite of AlO_x and NbO_x, rather than simply NbO_x. While NbO_x is easily etched (see M2 etch section), this alloy is not, and so the Al must be removed before the Nb underlying it is anodized. The barrier layer is very thin (~5 nm) and can be easily removed in a number of ways, such as with phosphoric acid or a chlorine-based plasma etch. The method used in the TD1 process, however, exploits the fact that tetramethyl ammonium hydroxide (TMAH), the active ingredient in most

photoresist developers, attacks thin films of Al and AlO_x. Simply running the wafers through a standard 60-second automated develop-track process removed the barrier without any need to deal with acid baths or development of a plasma-etch recipe.

A problem with the developer etch is that, like all wet etches, it is isotropic[16]. As a result, it slightly undercut the junctions by etching the barrier under the edges of M3. While this was not a significant problem with the standard 60-second developer recipe, longer developer immersions used on experimental wafers with thicker Al barriers had smaller features completely undercut by the developer (see fig. 32). Since the TD1 process may eventually evolve towards a thicker barrier layer, a dry-etch process (Cl-based) should eventually be designed to replace the developer etch and reduce undercut.

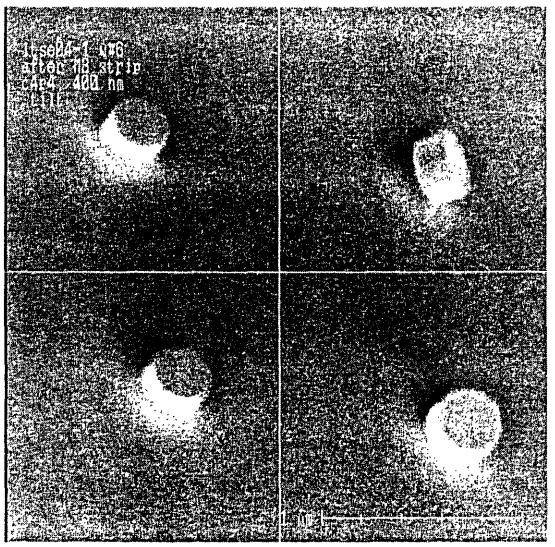

Fig 32: Scanning electron micrograph showing undercut and liftoff of 400 nm M3 post by TMAH developer etch of Al/AlO_x barrier.

3.2.4 M2 Etches

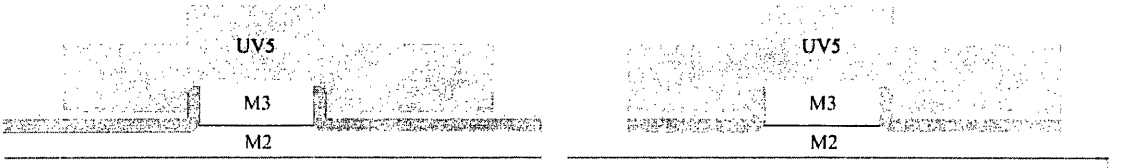

Fig. 33: Schematic of TD1 before (left) and after all M2 etching. The LTO adhesion layer, anodized NbO_x dielectric, and lower Nb layer have been removed.

LTO adhesion layer

The etch process for the LTO adhesion layer at M2 is identical to that at M3. The materials involved are all exactly the same, and the CF_4-based timed etch used before works equally well here.

NbO_x/Nb base electrode

The etch recipe used to etch the M3 Nb layer will also etch NbO_x, but the etch rate is roughly 1/10 that of Nb, making it somewhat impractical to etch both layers in one step using this recipe. The NbO_x etch rate in the standard TD1 Nb etch recipe can be greatly increased by increasing the temperature, however[44-46], and a variant of the recipe with the electrode temperature increased from 50°C to

80°C is used here to etch both films simultaneously. The tradeoff of the temperature increase is a minimal loss of sidewall verticality, but this is acceptable since the resolution requirements of the M2 layer are much more relaxed than the M3 layer.

Again, the etch can be stopped on an Al layer (in this case, the Al under-layer) with an easily detectable endpoint, so it was never necessary to characterize the respective etch rates of Nb and NbO_x for this recipe. The process generally takes between three and four minutes to complete.

The sidewalls in this process are still nearly vertical despite the high temperature, but this is actually disadvantageous. An outwardly sloped sidewall would both increase the surface area available for the M2 anodization and make M4-over-M2 step coverage less problematic later in the process. Experiments are currently in progress to engineer a recipe with a sloped profile (see fig. 34).

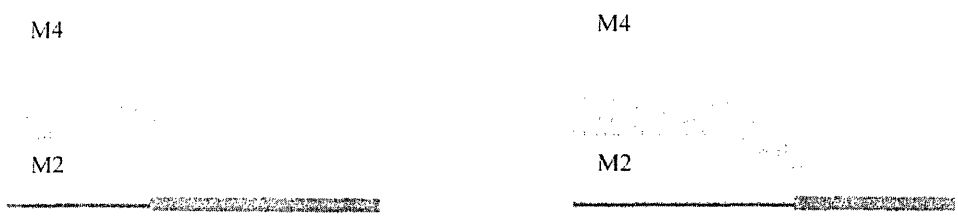

Fig. 34: Step coverage of straight vs. sloped M4-over-M2 profiles. Note the thin coverage at the edge of the straight profile on the left.

3.2.5 I2 Etch

In order to form vias between the M2 and M4 layers, the insulating layer of NbO_x between the two must be etched away prior to M4 deposition.

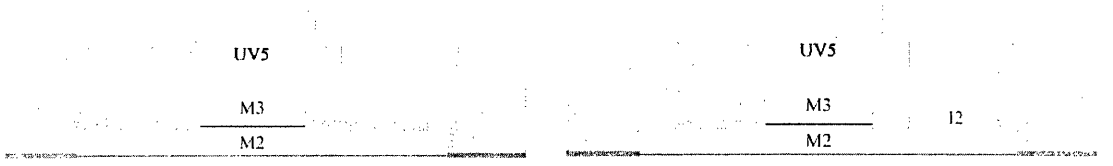

Fig. 35: Schematic of TD1 before (left) and after I2 etch.

In this case, the etch must remove NbO_x and stop on the underlying Nb. All known variations on the SF_6-based recipe used on the M3 and M2 layers have been experimentally determined etch NbO_x more slowly than Nb and so cannot be used. High-temperature CF_4 plasma etches, however, can exhibit the etch selectivity needed here, if only slightly[46]. The etch rates of Nb and NbO_x were measured in a 60°C CF_4-based recipe in the Lam Rainbow tool and a selectivity of nearly 2:1 was observed; not optimal but certainly usable (see fig. 36). Unfortunately, the Rainbow tool is not equipped with endpoint detection hardware, so a timed etch is the only solution at this step. Since the M2 layer is superconducting, overetching the I2 vias into the underlying Nb should not affect device performance unless the entire layer is etched through, though overetch should still be kept to a minimum to ensure good via filling during M4 deposition.

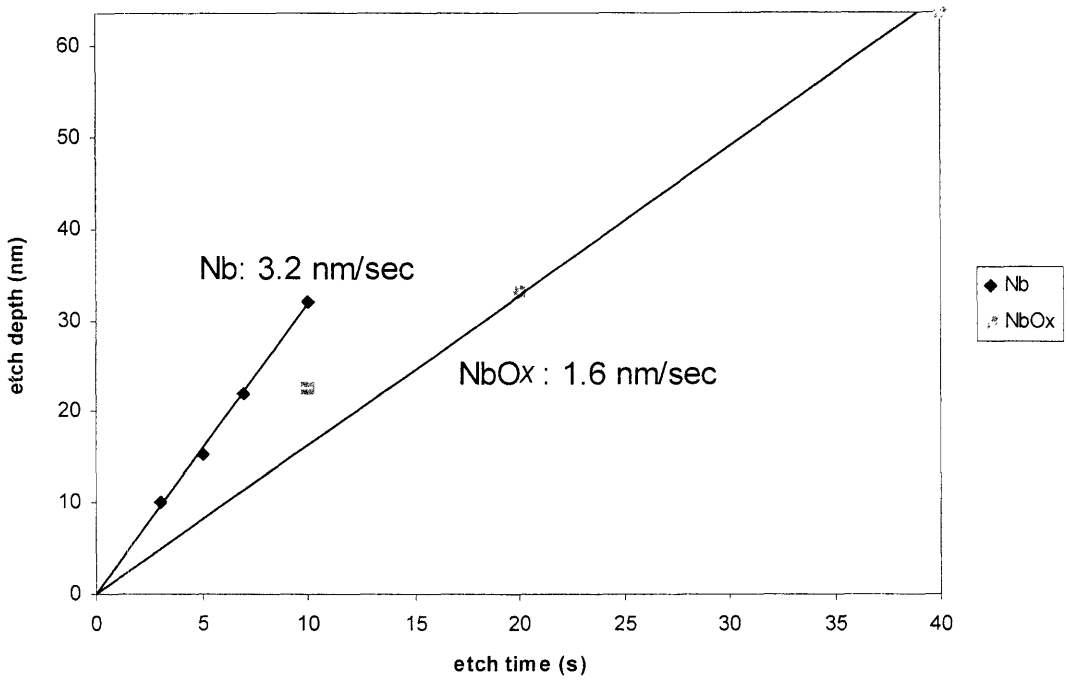

Fig. 36: Empirically-determined etch rates of Nb and NbO_x in the Lam Rainbow tool using a 60°C CF_4-based recipe. Unpatterned Nb and NbO_x wafers were used for the experiment, and measurements of initial and final thicknesses were taken using a spectrometer (for the dielectric) and a four-point probe (for the metal). Selectivity is roughly 2:1.

3.2.6 M4 Etch

The M4 etch is similar to the M3 etch, but somewhat complicated because it needs to stop on two different materials (NbO_x and AlO_x), as well as avoid damaging existing features on the wafer.

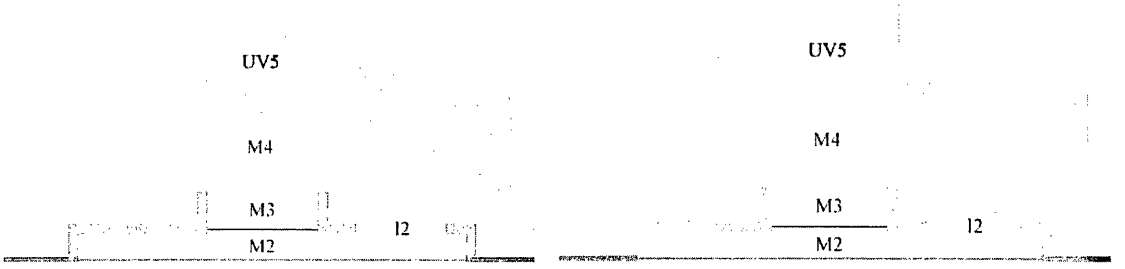

Fig. 37: Schematic of TD1 before (left) and after M4 etch. Note the different endpoint materials, depending on the topography of the wafer.

As demonstrated in the previous etch steps, stopping a Nb etch on an Al-based layer is straightforward. The NbO_x regions are slightly more difficult, although the high Nb/NbO_x selectivity that made the 50°C SF_6 recipe unsuitable for the M2 etch is advantageous here, as the etch rate of NbO_x in that recipe is only 0.2 nm/sec, allowing it to essentially function as an etch stop layer.

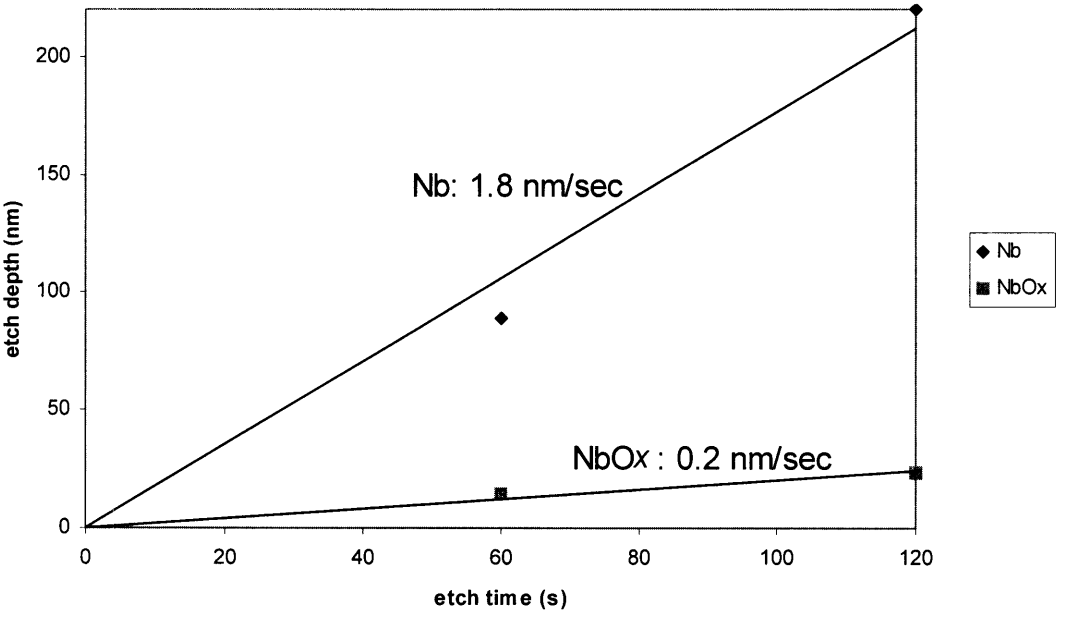

Fig. 38: Empirically-determined etch rates of Nb and NbO_x in the standard TD1 SF_6-based recipe at 50°C in the Plasma-therm etch tool.

Since the height of the fluorine peak on the spectrometer is inversely proportional to the etch rate, the method of endpoint detection used in the M3/M2 etches works equally well here, since the etch slows down by over a factor of ten when it hits the two oxide layers.

The verticality of the SF_6 etch worked against the process again here, as a major problem in the TD1 process was the presence of residual M4 "stringers" on the M2 sidewalls that the etch failed to remove (see fig. 39). These produced shorts and, in extreme cases, caused entire wafers of devices to fail. The "quick fix" to the problem was a severe overetch (1 min. or greater) after the endpoint was detected on the M4 etch, but a more

isotropic etch would be a much more elegant solution to this; experiments are currently being performed in this area.

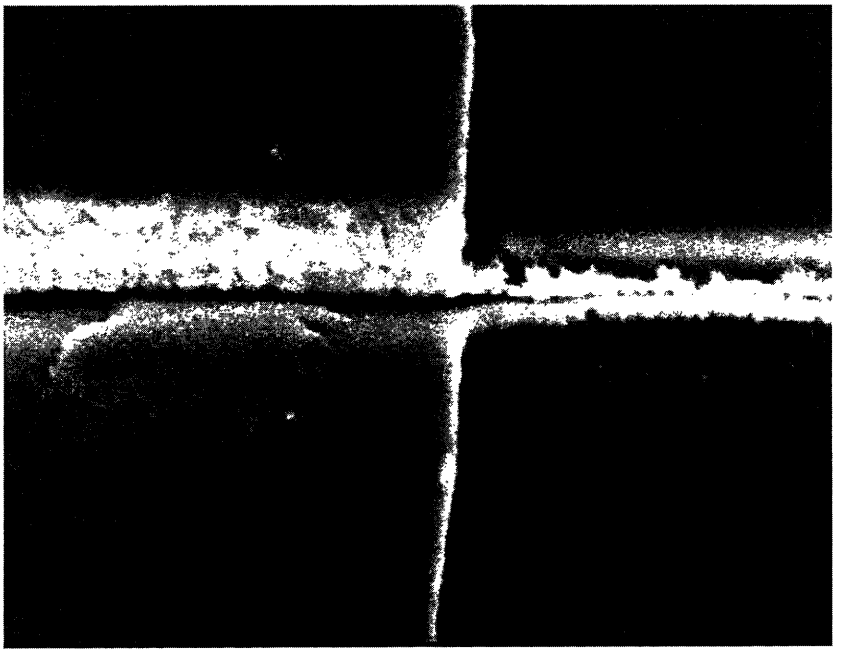

Fig. 39: Tilted scanning electron micrograph of residual M4 along the M2 sidewall (white residue in middle right of image). The M4 step over M2 is on the left side of the image. (Image courtesy MIT NanoStructures Laboratory)

3.2.7 Additional Investigation

While acceptable etch steps exist for each of the TD1 layers, more characterization work (particularly on the SF_6 etch) would greatly improve the robustness and stability of the process. A Design of Experiment (DOE) matrix (see Appendix A) is currently being analyzed to study the effects of temperature, pressure, and power on the isotropy and sidewall profile of the Plasma-therm SF_6 etch.

3.3 Anodiz7ation

3.3.1 Motivation

In order to isolate the upper and lower wiring layers and avoid short-circuiting the junction, a dielectric film between the two is necessary. The two most common ways to achieve this are thermal oxidation and chemical vapor deposition (CVD). The high temperatures required for thermal oxidation of Nb render it impractical here, as the mechanical properties of the trilayer impose a roughly 150°C limit on all processing temperatures[30]. Silicon dioxide (SiO_2) can be deposited at much lower temperatures via plasma-enhanced chemical vapor deposition (PECVD), and indeed is used elsewhere in the TD1 process, but the oxide film produced is generally porous and low-quality. In addition, both of these processes lack selectivity; oxide must be grown or deposited everywhere on the sample and then removed in the areas where it is not needed. Previous LTSE processes used a thick PECVD SiO_2 film as an isolation layer and removed it from the top of the M3 contact via chemical-mechanical planarization (CMP), but the significant processing overhead this represented made it undesirable for TD1, which emphasizes rapid turnaround time. An alternate (and somewhat unorthodox) process known as anodization was used to selectively grow a ~50 nm niobium oxide (NbO_x) film on the wafer at room temperature instead[47].

3.3.2 Theory[48]

Anodization is a process most commonly used in metallurgical applications to grow a passivating oxide film on a metal surface. It employs a cathode of nonreactive metal, the sample to be anodized, and an electrolytic bath; when a potential greater than some material-dependent critical potential is produced between the cathode and the sample to be anodized in the solution, oxidation will take place at the surface of the material.

The electrochemistry involved in anodization is nearly identical to that of electroplating and essentially involves ionic transport at high electric fields. Ions and electrons move through three distinct regions during the process: the metal, the oxide film already grown, and the electrolytic solution. The specifics of the process can become very complex, as some materials can form more than one oxide, different phases of oxide, or oxides with significant defects that inhibit ionic conduction. Fortunately, TD1 only requires anodization of Nb and Al, both "valve" metals that exhibit simple exponential current-field dependence (similar to a diode) when polarized anodically

The anodization of Nb by the method used for TD1 forms niobium pentoxide (Nb_2O_5), as this is the most stable of the niobium oxides[49]. The half-cell reaction taking place at the anode when Nb is anodized is:

$$2Nb + 5H_2O \rightarrow Nb_2O_5 + 10H^+ + 10e^-$$

Note that the half-reaction is independent of the actual choice of electrolyte, which will be discussed further. Similarly, the half-reaction for anodizing Al is:

$$2Al + 3H_2O \rightarrow Al_2O_3 + 6H^+ + 6e^-$$

The stoichiometry of these reactions indicates the amount of charge required to oxidize (anodize) a given amount of metal; 2 moles of Nb and 10 faradays of charge will form 1 mol of Nb_2O_5, while 2 mol of Al and 6F of charge are required to produce 1 mol of Al_2O_3.

Constant-current anodization of valve metals can be modeled as a problem of charge carriers and potential barriers. When a specific current is applied to the system, cations in the metal react with oxygen in the water to form a metal oxide film at the sample-electrolyte interface. As the surface oxide film forms, it creates a potential barrier to further anodization-- a higher voltage is required to "drive" the carriers through the film to the interaction point and maintain a constant ionic current. The potential increase per unit thickness of the oxide film is known as the specific ionic resistivity, and varies by material. By continually ramping the applied voltage to maintain a set constant current until a specified "breakpoint voltage", an oxide film with a thickness dependent on the specific ionic resistivity and the maximum voltage can be grown[49]. For a given oxide film, the final thickness is linearly related to the breakpoint voltage, with a slope inversely proportional to the specific ionic resistivity of the oxide. This material-

dependent parameter is usually given in nm/V; values for several valve metal oxides are given in the chart below.

Metal	Oxide Formed	nm/V
Al	Al_2O_3	1.29
Ti	TiO_2	2
Si	SiO_2	0.39
Ge	GeO_2	4.7
Sn	SnO_2	1.9
Nb	Nb_2O_5	2.25
Ta	Ta_2O_5	1.62

Table 3: Anodic oxides and thickness-vs.-voltage values for several common "valve" metals.[48]

There are two types of carrier movements to consider during valve metal anodization; cations from the metal and oxygenated ions from the electrolyte. As a result, determination of the ratio of oxide formed on the surface of the substrate to oxide formed by actually oxidizing the substrate is nontrivial. The exact amount of substrate consumption taking place is dependent on the carriers' relative mobilities through the various materials (electrolyte, oxide, and metal) at given fields, and a detailed analysis is beyond the scope of an engineering discussion. In general, it has been empirically observed in the Lincoln system for both Al and Nb that producing an anodic oxide of thickness x will consume roughly a thickness $x/2$ of the substrate; that is, a 50 nm film

will produce approximately 25 nm of oxide on the surface and oxidize 25 nm of the substrate[49]. This could be varied somewhat by using different electrolyte solutions, but was perfectly satisfactory for the purposes of TD1.

The current used in constant-current anodization is in principle not related to the film thickness and only affects processing speed. However, it has been observed that both the oxide density and oxide uniformity are inversely related to the current[50]. This observation is unsurprising, as rapid oxidation in any oxidation process will generally cause voids to form in the film and result in a lower density, and anodization is no exception when the amount of carriers simultaneously participating in oxidation is increased. Higher working currents also cause the process to be more susceptible to local variations in the resistivity of the electrolyte, which in turn can cause variations of the potential at the surface of the sample and nonuniform film growth. An empirical compromise of 225 mA was used in TD1, as it gave reasonable balance between film quality and processing time.

3.3.3 Implementation

The anodization apparatus at MIT Lincoln Laboratory consists of a bath of ammonium tartarate solution (buffered with NH_4OH at pH 5.1), a platinum-coated cathode wafer, and a computer-controlled power supply attached to both wafers via alligator clips. The wafer to be anodized and the platinum cathode are immersed in the bath and the cathode is clipped to the positive terminal and the anode to the negative. A computer program is

then used to control the power supply and drive a constant, set current through the system while allowing the voltage to increase to a specified breakpoint. As discussed, the oxide thickness and breakpoint voltage were found to have a stable, linear relationship when anodizing Nb and Al, which held true to very high voltages (see fig. 40). After anodization, a standard dump rinse and spin-dry cycle were used to remove excess electrolyte.

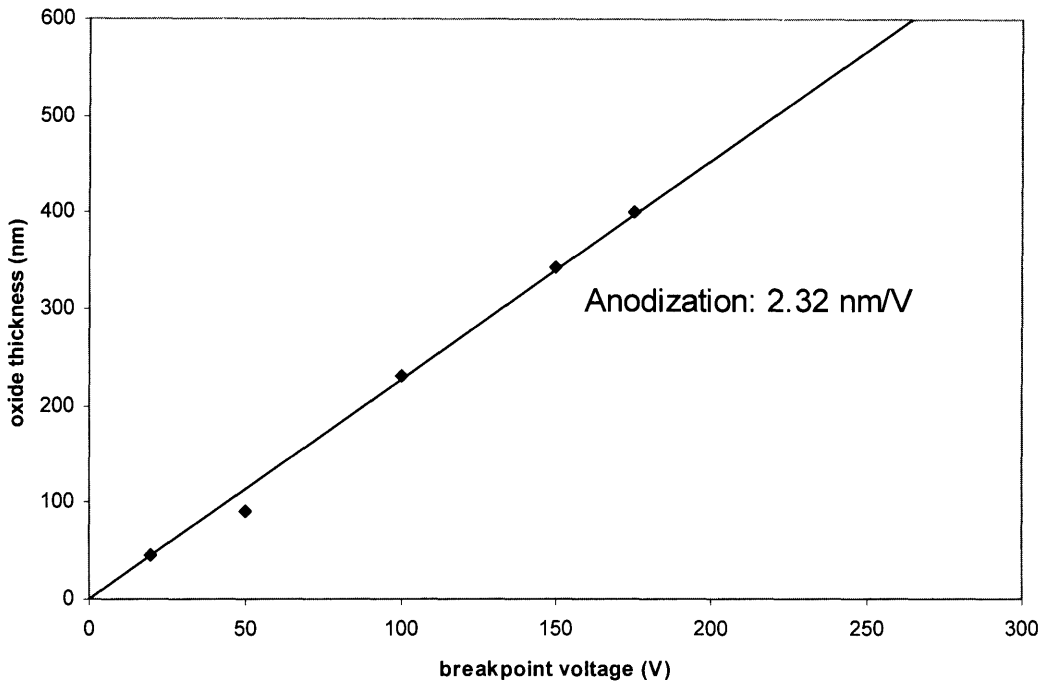

Fig. 40: Experimental data for thickness of NbO_x vs. anodization breakpoint voltage using the MIT-LL setup. Note that the slope is 2.32 nm/V, slightly higher than the theoretical value given in the table above.

The TD1 process contains two anodization steps, which differ slightly. In the first, the M3 sidewalls and M2 surface are oxidized selectively to the top M3 contact. A

breakpoint voltage of 20V is used in this step, resulting in an oxide film roughly 46 nm thick and insulating the M2 surface and M3 sidewalls from the future deposition of M4.

Fig. 41: Isolating the M3 layer and M2 surface via anodization. Cross-section of the process before (left) and after (right) the M3 anodization step.

In the second anodization step, the M2 sidewalls need to be passivated after M2 has been etched down to the SiO_2 substrate. Since SiO_2 is an insulator, the only way to make contact with M2 and allow current to flow is by depositing a thin conductive layer under the trilayer. 15nm of Al is used for this purpose, as it also functions as an ideal etch stop for the M2 etch. Rather than a set breakpoint voltage, this anodization step relies on an endpoint--when the Al layer has completely oxidized, current can no longer flow and the anodization abruptly stops. For a 15 nm Al layer, the anodization stops at 17-18V, resulting in a ~40 nm oxide film on the M2 sidewalls and a ~20 nm AlO_x film in the field.

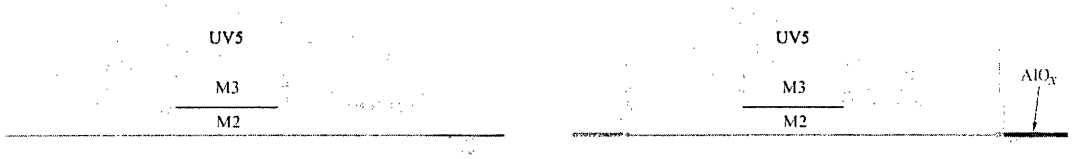

Fig. 42: Isolating the M2 sidewalls via anodization. Cross-section of process both before (left) and after (right) the anodization of the M2 sidewalls and thin Al layer in the field.

3.3.4 Selective Anodization at High Resolutions

As discussed, one of the key advantages of anodization as an oxidation process is selectivity; if photoresist is present on a sample to be anodized, ions cannot flow through the resist and the metal under it will be protected. This fact is the basis of the selective Nb anodization process (SNAP) for fabricating Josephson junctions[51]. However, poor adhesion of photoresist to metals during anodization has generally imposed limits on the minimum feature size realizable with selective anodization[37].

The application of large electric fields to photoresist films is thought to cause them to warp and contract, leading to problems with adhesion[52]. While the exact mechanism of the warping is unknown (chemical compositions of photoresists are generally proprietary), it is thought to be the consequence of interactions between the electric field and polar regions of the polymeric molecules of the resist. The result is that, when immersed in electrolyte during anodization, resist tends to lift up around the edges and allow electrolyte to "creep" underneath. For larger features this simply results in some thin oxide around the edges and is unimportant, but features smaller than 1 μm will be

entirely undercut by the electrolyte and lift off the wafer completely (see fig. 43). As the TD1 process is concerned with features as small as 0.2 μm, this is a significant concern.

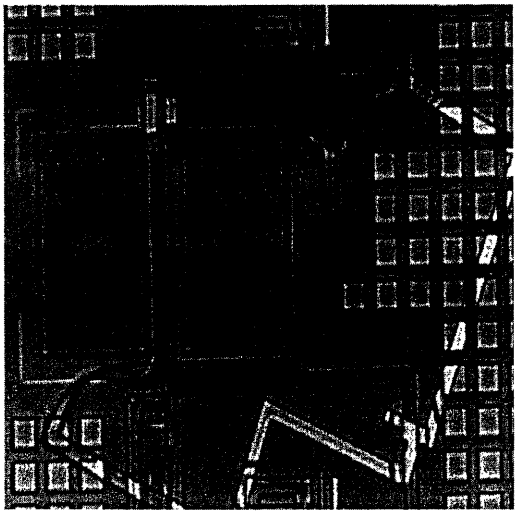

Fig. 43: Optical micrograph of resist lifting after anodization. The submicron lines have clearly lifted off the substrate (blue) and the underlying metal has been anodized.

The resist liftoff during anodization was exhaustively characterized via direct experimentation. It was originally thought to be an attack on the photoresist by the electrolyte (slightly acidic with a pH of 5.1), but the fact that liftoff was never observed without an applied voltage present suggested otherwise, as did the tendency of the resist to retain its shape even after liftoff (see fig. 44). Varying the current setpoint had little or no effect, which was unsurprising as the current has nothing to do with the maximum electric field the sample is exposed to. Moreover, the liftoff became proportionately more pronounced at higher voltages; above 100V, nearly all resist disappeared. Several types of photoresist (both standard i-line and DUV) exhibited identical behavior, but since most photoresist is chemically similar this was expected. The conclusion was that a

photoresist mask on metal was not feasible for selective anodization at the resolutions required by TD1.

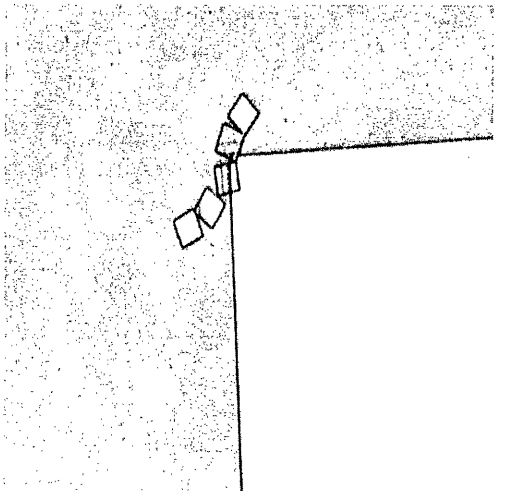

Fig. 44: Optical micrograph of small resist features that have lifted from the substrate and become stuck to a larger, still-attached feature after anodization. Note that the structure of the resist (small squares) remains even though the adhesion has failed.

The obvious alternative was a "hard mask" of some other material[37]. Unfortunately, the low maximum processing temperature of the TD1 process drastically limited choices. a PECVD-deposited SiO_2 film turned out to be the only feasible hard mask option, as it could be deposited at low temperatures, patterned very accurately with a CF_4 plasma etch, and easily stripped using a silox (NH_4F-CH_3COOH) wet etch with negligible attack of the underlying Nb and NbO_x. The low-temperature film (LTO) was blanket-deposited (several thicknesses, from 30 nm to 150 nm, were experimented with), patterned with UV5 resist, and etched down to the Nb using a timed plasma etch. The resist was then stripped, leaving an oxide film mirroring the photoresist pattern on the surface of the wafer.

The use of an oxide mask eliminated the liftoff problem, but introduced an equally severe problem: the "fluffy" oxide grown at 150°C was extremely porous and nonuniform. Even at the highest thicknesses used, the oxide did not fully prevent ionic conduction and anodization of the underlying metal. In principle a thicker film might have fixed this, but the increased undercut involved in pattern-etching a thicker film made this an unsatisfactory solution. The improved adhesion was useful, however, and eventually led to a compromise between the resist-mask and oxide-mask approaches.

While the porosity and nonuniformity of the available LTO made it a poor anodization mask, these same features (as well as its excellent adhesion to Nb) made it ideal as an adhesion layer for photoresist (see fig. 45). A very thin (30 nm) layer of LTO was deposited prior to the photolithography step and pattern-etched as before. Now, however, the resist was not stripped from the samples prior to anodization. The "rough" surface of the oxide served to hold the resist in place, despite stresses from the electric field, and once the liftoff problem was solved the photoresist served as an ideal anodization barrier. After anodization, the entire mask stack was removed by stripping the photoresist (see chapter 3.4) and then the oxide via a short silox etch. Using this masking method, the resolution of the TD1 selective anodization process improved to the point where it was on the same order as the photolithographic resolution; that is, any feature definable lithographically could also be isolated anodically with an exposed upper contact. In practice this translated to features about 0.2 μm in diameter (see fig. 46), a factor of five improvement over the standard resist-over-metal anodization process.

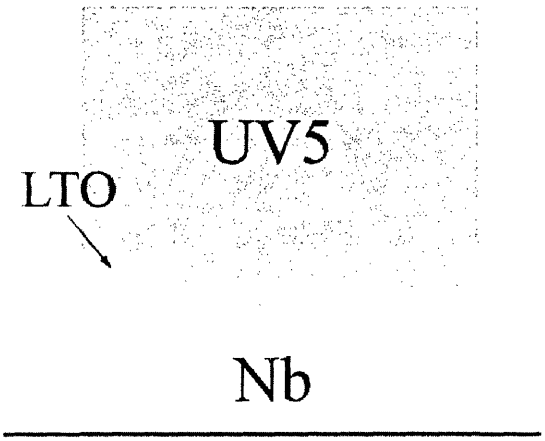

Fig. 45: Cross-section schematic illustrating how the seemingly low-quality thin LTO film actually improves resist adhesion.

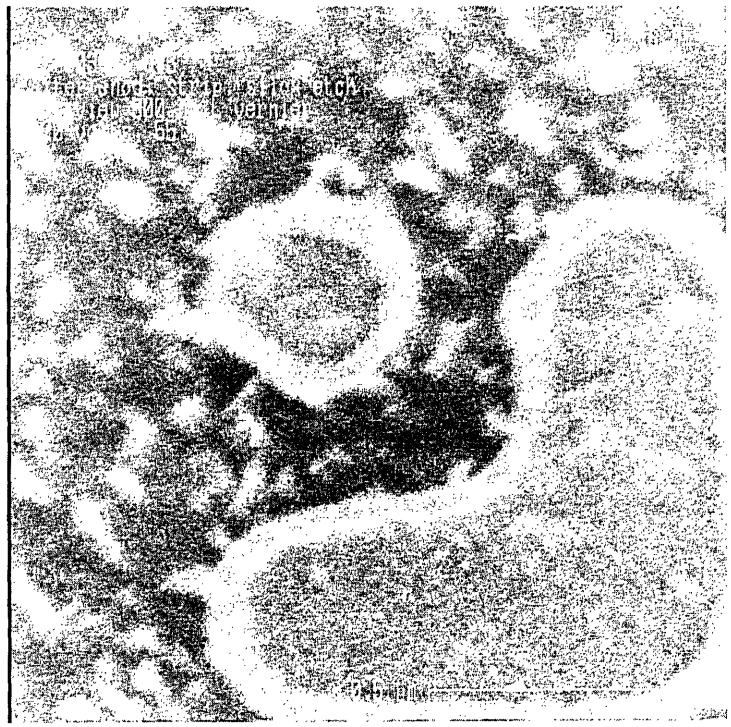

Fig. 46: Scanning electron micrograph of ~0.25 μm post resolved with the TD1 anodization process. The oxide sidewalls and metal contact are clearly visible. Note that the field surface roughness is an artifact of the experiment this image was taken from and not present in the actual process.

3.4 Mask Removal

3.4.1 Overview

Once the various lithographic patterns have been etched into the wafer itself, it is necessary to remove the photoresist and other masking materials (ARC, SiO_2) in order to continue processing. In general microelectronics processing, resist stripping is accomplished using either a harsh solvent with a high pH, a high-pressure oxygen plasma, or some combination of the two. However, as is the case in many of the other steps in the TD1 process, the strict temperature limit on processing enforced by the Nb-AlO_x-Nb trilayer system, as well as concerns that an oxygen plasma would oxidize and damage the underlying Nb, complicates the problem significantly.

There are three major mask material types used in TD1, as follows (see Chapter 2 for process flow details). Each requires its own specific strip process to fully remove the mask material and avoid substrate damage.

- The M3 (junction definition) layer strip (fig. 47). The masking material here is a thin (~30 nm) adhesion layer of LTO (low-temperature PECVD SiO_2) under 395 nm of Shipley UV5 DUV resist. The exposed substrate here is the Nb contact between the M3 and M4 layers and the NbO_x barrier between the M2 and M4 layers. Depending on the resolution limits required, a layer of Shipley AR3 ARC may also be present between the LTO and resist.

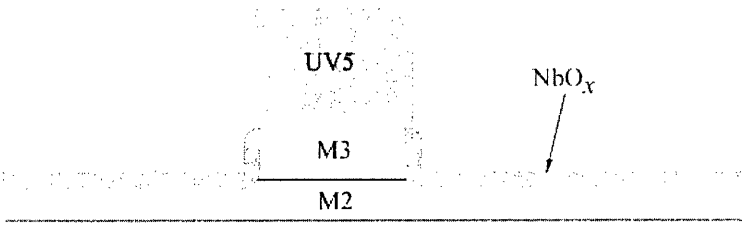

Fig 47: Cross-sectional schematic of TD1 process just before M3 strip.

- The M2 (lower wiring) layer strip (fig. 48). The masking material is the same as the M3 step (with the exception of the ARC, as the resolution requirements for this layer are much less stringent), but the AlO_x under-layer is now exposed, limiting the choice of solvents drastically.

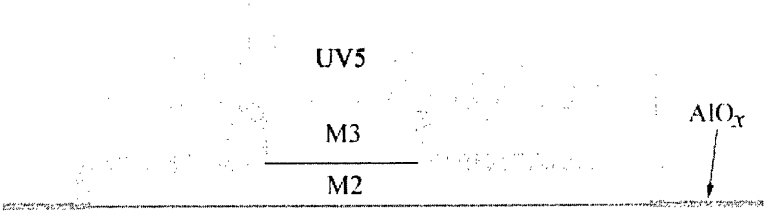

Fig 48: Cross-sectional schematic of TD1 process just before M2 strip.

- The strip steps for the via and upper wiring (I2, M4) mask layers. Since the masks in these steps do not have to resist anodization damage, the material used is simple UV5 without an adhesion layer.

3.4.2 M3 Strip Issues

The solvent used to remove photoresist in previous Nb-based superconducting electronics processes at Lincoln Laboratory was Air Products ACT935, a harsh strip agent with a pH of nearly 14. While relatively harmless to Nb, however, ACT935 severely etches thin layers of Al and AlO_x. In older processes, this resulted in a reduction in effective junction area due to solvent attack of the Al-AlO_x barrier between M3 and M2. The extent of this area reduction was approximately 150 nm per side, which was a nonissue with large junctions but became very significant at the deep-submicron dimensions used in TD1.

In light of this, ACT935 was abandoned in favor of Air Products ALEG310, a much gentler solvent with a pH only slightly above 7.0. This new agent did not attack Al, but the tradeoff was a marked decrease in stripping effectiveness. To counter this, a process was developed in Lincoln Laboratory's Gasonics oxygen-plasma resist strip system to be used in conjunction with the wet strip.

While stripping resist with oxygen plasmas is standard practice in the semiconductor industry, it is generally not compatible with Nb-based processes due to the high diffusivity of oxygen through the grain boundaries of sputtered Nb[31]. Additionally, almost any plasma-strip process will heat the substrate material well past our 150°C processing limit. While these limitations make designing a full, robust strip recipe for TD1 in a plasma system impractical, it was possible to modify an existing "descum" (no

heat used) recipe to slightly ash some of the more stubborn areas of resist and polymeric etch byproducts without heating the wafer past about 100°C (which also kept oxygen diffusion to a minimum). The final M3 resist strip recipe was as follows:

1. 30 min. ALEG310 immersion
2. 140 sec. Gasonics descum without heat
3. 30 min. ALEG310 immersion

Following the resist strip, a short (5 sec.) silox (NH_4F-CH_3COOH) wet etch was used to remove the thin adhesion layer, with no observable adverse effects on the exposed Nb or junction area. We refer to this process as the TD1 "wet/dry/wet" strip process.

The ALEG310-based strip recipe was successful in removing UV- and plasma-hardened UV5 resist without observable undercut of the Al barrier layer. Unfortunately, it was less than consistent in removing polymers formed at the edges of features

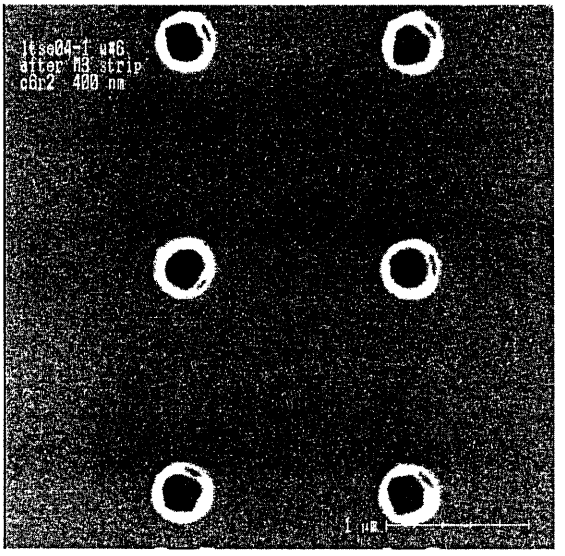

Fig. 49: Scanning electron micrograph of 400 nm M3 posts after mask strip. Note the polymer still present around the edges of the mesas.

This polymer-removal problem was not improved by increasing the immersion time, adding additional wet and dry strip steps, or rearranging the order of the strips, and the conclusion was that the SF_6 Nb etch produces some type of polymer that cannot easily be stripped with ALEG310. Fortunately, fewer than 25% of the small features on a typical TD1 wafer exhibited this residue, so it was possible to fabricate working devices using the wet/dry/wet strip process until more research could be done on an alternative strip.

Removing AR3 anti-reflective coating proved much more problematic with the new strip recipe. Wafers with ARC as part of the mask stack were left with severe residues after strip and were not improved by subsequent re-strips. The conclusion was that ARC could not be used in the TD1 process without a major breakthrough in the strip procedure. This effectively limited lithographic resolution of the process to about 400 nm. Analysis is

currently being done by Air Products to design a strip recipe compatible with both AR3 and Al, but results from the study are still pending.

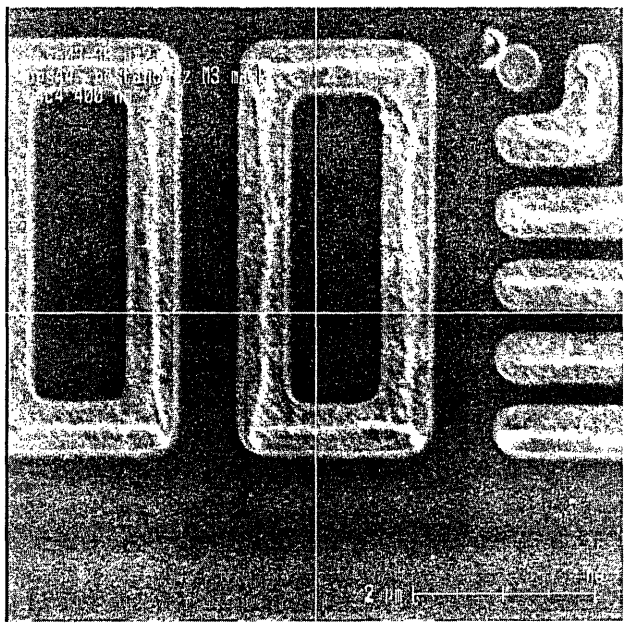

Fig. 50: Scanning electron micrograph of 400 nm M3 lines and posts after attempted strip of resist and ARC. Note residues covering nearly all of the feature surfaces.

Alternative methods of reaching sub-400 nm resolutions with the TD1 process, such as using electron-beam lithography to define the M3 layer or simply biasing the process with etches or other steps, are currently being investigated, but 400 nm remains the process resolution limit (from a photolithography standpoint; biasing during etch and other steps can reduce dimensions from those printed) at this time.

3.4.3 M2 Strip Issues

With the exception of ARC, the M2 mask material is identical to that of M3. Based solely on that consideration, removal of the M2 mask would appear to be trivial since the main stumbling block of the M3 strip is no longer present. Indeed, stripping the resist from M2 proved straightforward with the wet/dry/wet (WDW) process developed for M3.

However, problems arose when the LTO adhesion layer was stripped using silox immersion. After even a brief (5 sec.) silox etch, all M2 features less than 1 μm wide were observed to lift off the wafer entirely.

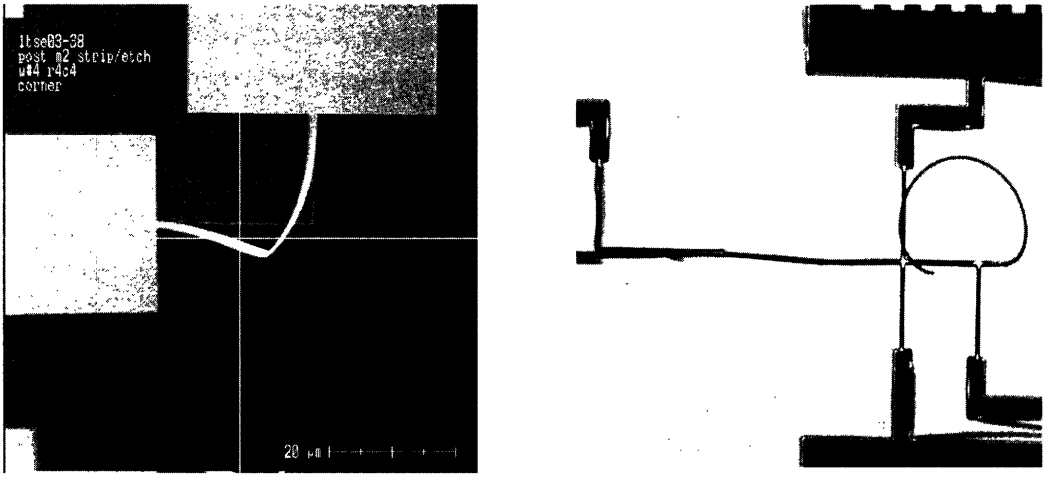

Fig. 51: Scanning electron (left) and optical micrographs of M2 feature lifting after silox etch.

Silox is known to be benign to AlO_x and NbO_x, in principle the only two materials exposed to it during M2 strip[16], so the cause of the liftoff was not immediately obvious. Extensive analysis of the lifted wafers using both atomic-force (AFM) and scanning-

electron (SEM) microscopy was done to determine the extent of the liftoff, and measurements taken via AFM (fig. 52) indicated that the M2 features were undercut by over 1 μm laterally.

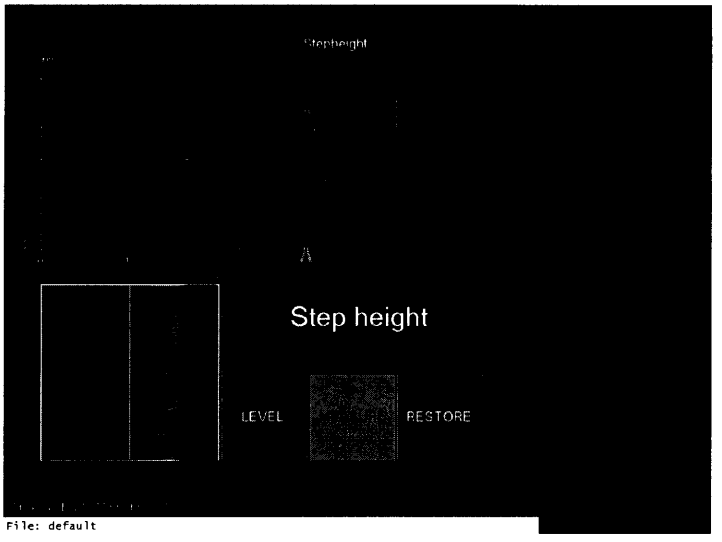

Fig. 52: Atomic-force microscope cross-sectional plot of lifted M2 edge. Note that the feature is almost twice as tall at the edge as in the bulk, and that the lifted region extends well over 1 μm into the bulk.

In addition to the obvious problem of small wires being completely removed, the lifting of the edges on the larger M2 features caused problems during the sputter deposition of the upper (M4) wiring layer. M4-over-M2 step coverage was impossible (fig. 53), as the edge of the M2 features was now almost 400 nm high and the M4 film was only 200 nm thick, and the vacancy under the lifted M2 allowed sputtered Nb to get underneath it and cause shorts between features. An alternative oxide strip method, as well as the cause of the liftoff, had to be found.

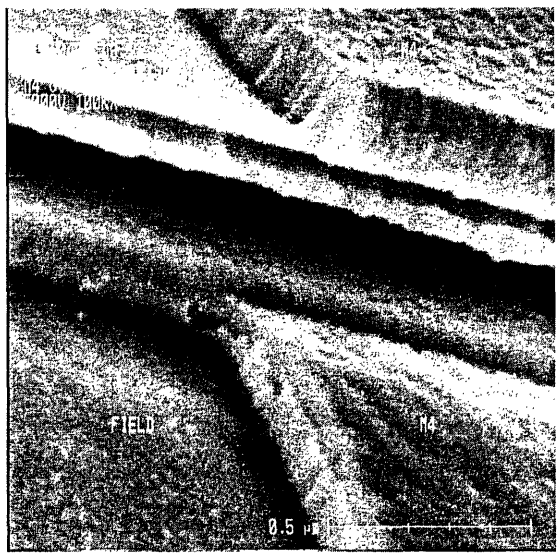

Fig. 53: Scanning electron micrograph showing the effect of M2 lifting on the M4 layer. Note the discontinuity between the M4 in the field and the M4 over M2 due to lifting, as well as the M4 underneath the lifted M2.

Sidewall polymers generated during the M2 etch seemed like the most likely cause of the lifting. While silox does not attack AlO_x, it will attack Al, and to a certain extent the high-quality thermal SiO_2 substrate. Polymer on the M2 sidewalls during anodization could keep the Al under-layer from fully anodizing near M2, leaving a thin "seam" of Al at the feature edges. Resist stripping removes the polymer and exposes the Al seam, which is then attacked by silox and allows the silox access to the SiO_2-Al and Al-Nb interfaces under the trilayer, which it quickly attacks (see fig. 54). The liftoff is greatly speeded up by compressive stress on the M2 layer that causes it to "peel" once attack on the interface has begun.

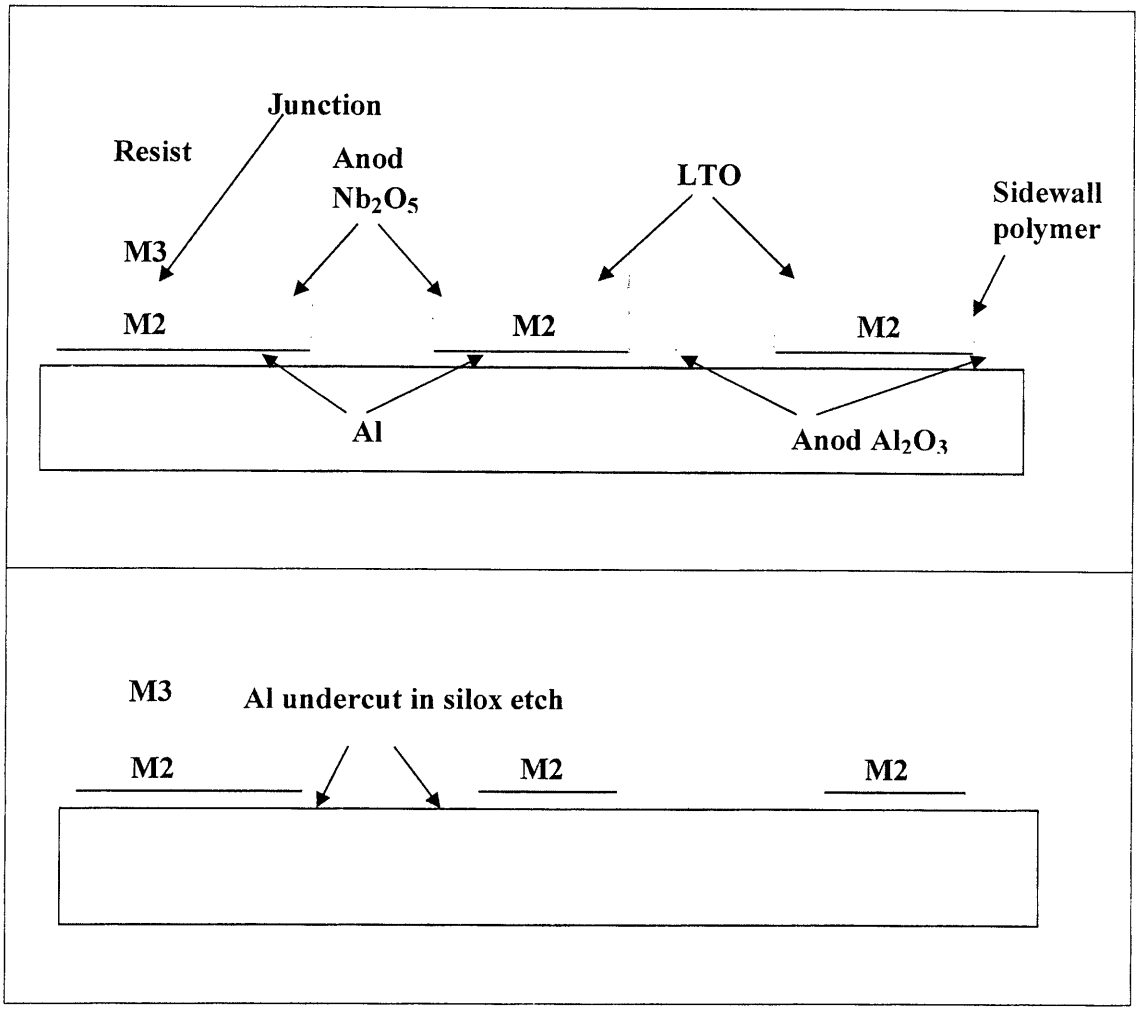

Fig 54: Possible mechanism for M2 liftoff. Sidewall polymers are formed during M2 etch (top) and block anodization of Al under-layer, which is attacked by the silox etch.

Several methods of stopping the M2 liftoff were attempted. Noting that the silox does not cause the M3 features to lift, the Al under-layer was slightly anodized so that M2 was sitting on AlO_x rather than Al. The wafers were annealed at 400°C after the Al layer deposition but before the trilayer deposition to attempt to improve SiO_2-Al adhesion. The deposition pressure of the M2 Nb layer was also varied, as the stress of the film is proportional to the deposition pressure. None of these process variations made any

significant improvement on the liftoff, and it was concluded that silox could not be used to strip the M2 adhesion layer.

As an alternative to silox, the CF_4 plasma etch used to pattern the LTO adhesion layer was also used to strip the adhesion layer at M2. Since the thickness of the layer and etch rate of the recipe on LTO were known, a simple timed etch was relatively straightforward. The significant tradeoff to using plasma to remove the LTO, however, was loss of selectivity. Silox attacked the NbO_x barrier layer negligibly, while CF_4 plasma etched it at 2-3 nm/sec. The Lam Rainbow plasma etcher used in this step is an inherently nonuniform tool, and an overetch of at least 25% is required to ensure that all the LTO on the wafer is removed. This overetch translated to a nonuniform loss of between 5 and 10 nm of dielectric between the M2 and M4 wiring layers, which degrades device performance but is still preferable to the liftoff seen when using silox.

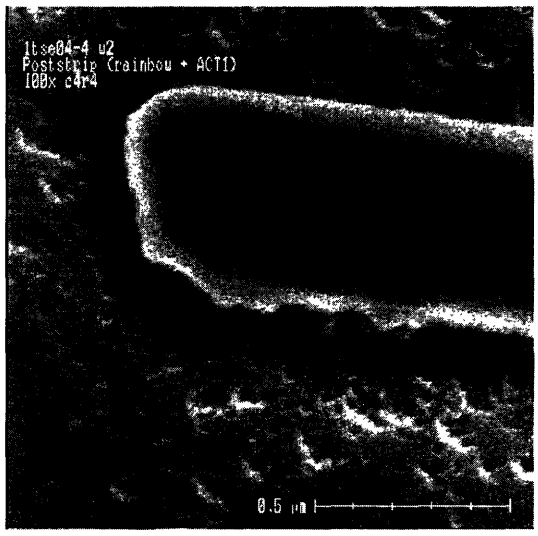

Fig. 55: Scanning electron micrograph of a 400 nm M2 line after plasma strip of the LTO. The rough edge is most likely etch polymers that were not fully removed.

3.4.4 I2/M4 Removal

The remaining two masks (M6 is removed via liftoff) are simply UV5, and their removal was much more straightforward. No problems beyond the occasional residues were encountered with using the wet/dry/wet strip to remove the I2 and M4 masks.

3.5 Metallization and Liftoff

3.5.1 Overview

In order to make contact to the upper wiring layer (M4) of the TD1 process, gold contact pads had to be deposited. Since alignment and resolution were not an issue here (the only features being patterned are very large contact pads) and pattern-etching gold is problematic at best, a simpler metal deposition and liftoff process is used rather than the pattern-and-etch method used to define previous metal layers in the process.

A liftoff process consists of three steps: patterning, metal deposition, and liftoff. In the patterning step, a negative image of the desired pattern is placed on the wafer with photoresist as usual. The areas where a gold contact is desired are clear, while the rest of the wafer is covered in resist. A thin (50 nm) layer of titanium was then deposited on the resist-coated wafer via electron-beam evaporation, followed by 400-500 nm of gold by the same method. A thin Ti layer improves adhesion between the Nb and the Au films and is generally used in all metal/liftoff processes. The wafer was then immersed in an

ultrasonic acetone bath, which dissolved the resist and caused the metal in the field to "lift off" the wafer, while the metal in the clear regions remains. Alignment and precise dimension control of metallization/liftoff processes is difficult, but the features being patterned were generally between 10 and 100 μm in diameter so this was unimportant for our purposes.

3.5.2 Processing Considerations

The most critical factor in a liftoff-based process is the engineering of a discontinuity between the metal on top of the resist and the metal on the substrate. This means that the resist layer must be thicker than the layer of metal film deposited; the empirical rule-of-thumb generally used is that the resist layer should be at least 1.5x as thick as the metal layer.

Though metal deposited in electron-beam evaporation is ideally generally deposited near-vertically, no process is perfect and the actual deposition in the Lincoln Laboratory evaporation system is far from vertical. In order to avoid sidewall coverage of the resist layer with metal, a resist profile wider at the top than at the bottom is desirable. An inwardly-sloped sidewall, "lip" at the edge of the resist, or both can greatly improve the robustness and reliability of a liftoff process.

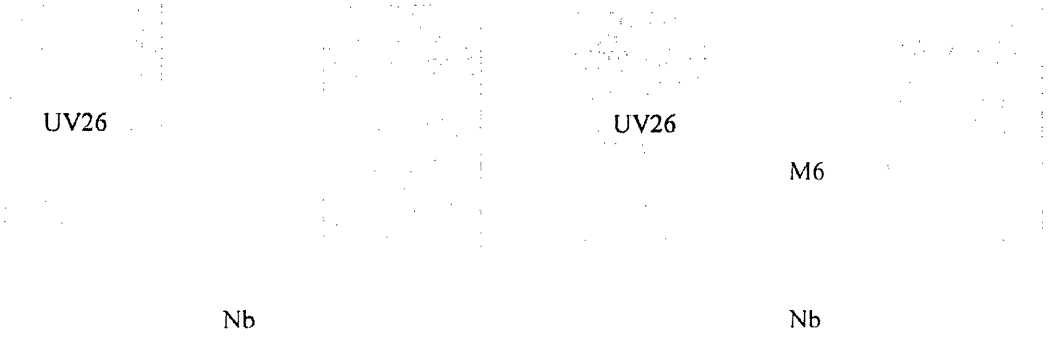

Fig. 56: Cross-sectional schematic of a vertical resist profile before (left) and after metal deposition. Note the thin layer of gold on the resist sidewall, even though the resist is significantly thicker than the metal layer.

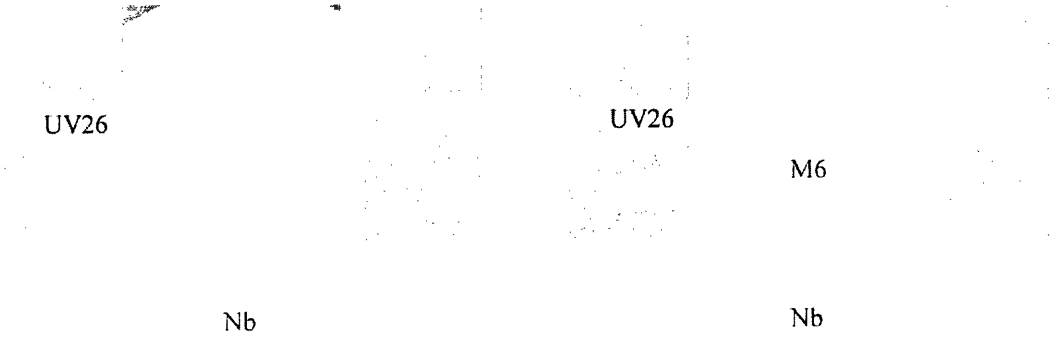

Fig. 57: Cross-sectional schematic of a resist profile with a "lip" at the edge of the resist region to reduce sidewall coverage and improve separation between metal regions.

3.5.3 Implementation

Since no process existed at Lincoln Laboratory for metallization and liftoff with deep-UV resist, a new process was engineered from the ground up. The first thing to consider in generating the process was the type of resist used. The Shipley UV5 used elsewhere in the process (and also the standard resist for DUV work at Lincoln Laboratory) was undesirable, as it could not be spin-coated at thicknesses much greater than about 750 nm. Shipley UV26 resist proved much more suitable, however, as it could be used at thicknesses up to 1500 nm and was similar to UV5 in most other respects, which reduced time needed to characterize it.

An inward-sloped resist profile can be obtained via controlling the focus during photolithography, but this is generally impractical on a production-grade stepper such as the one used at Lincoln Laboratory, where the focus is only controllable to within about 0.1 μm. Instead, a sidewall "lip" was formed by exploiting the properties of chemically-amplified photoresist.

One of the major drawbacks of chemically-amplified resist (CAR) is its susceptibility to surface contamination by bases[39]. Even low levels of contaminants in a film of CAR can neutralize the acid molecules used to catalyze the development reactions and cause odd sidewall profiles and incomplete surface development. The Lincoln Laboratory clean room filtered carefully, so contamination is generally a non-issue, but a deliberate,

controlled contamination of the surface of a UV26 film can be used to produce a liftoff-ready profile.

The active ingredient in almost all developer used in microelectronics is tetramethyl ammonium hydroxide (TMAH), a base. Immersing a wafer coated with chemically-amplified resist in a dilute TMAH solution for a set amount of time will allow hydroxide ions to diffuse into the resist and neutralize acid molecules near the surface. The thickness of this "poisoned" region is dependent on both the concentration of the solution and the immersion time.

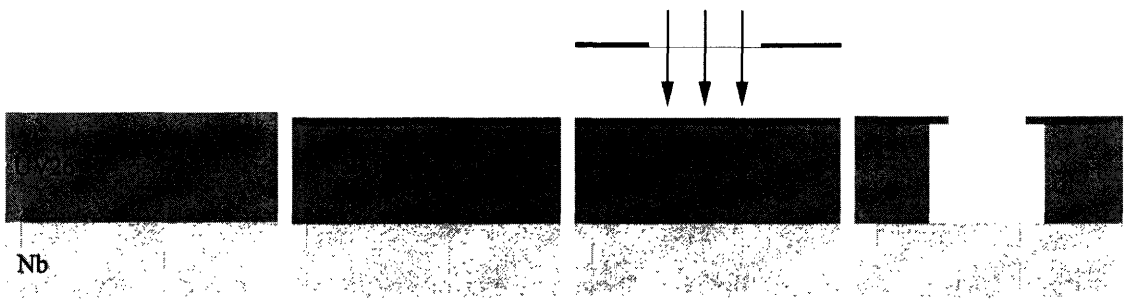

Fig 58: Schematic representation of surface-poisoning process. From left to right: resist coat, surface poisoning via TMAH immersion (poisoned region in green), exposure, and development. Note overhang in the post-development resist profile due to the poisoned surface region.

The exact process employed for the M6 layer of TD1 was as follows:

1. Spin-coat 1500 nm of UV26 resist on wafer
2. 1 min. immersion in 1% TMAH, followed by 5 min. rinse in DI water
3. Exposure (dose: ~200 J/cm^2)
4. Development in standard develop track

Some experimentation was done with the concentration/time of the TMAH immersion, as well as the effect of performing it before vs. after exposure, but since this is essentially a binary process (the overhang is either present or not), extensive characterization work was not performed once a working combination of parameters was hit upon.

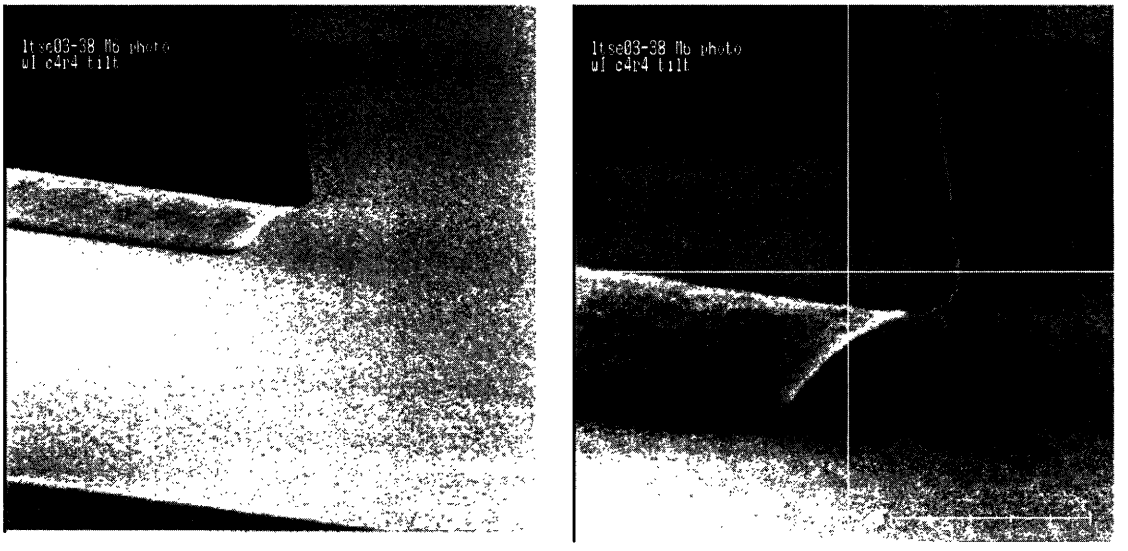

Fig. 59: Scanning electron micrographs of the resist profile generated using the described surface-poisoning process. The overhang is clearly visible in the right image.

Using the UV26-based surface-poisoning process developed in-house, deposition and liftoff of the M6 contact layer was successful on all TD1 wafers to date. If the TD1 process is ever expanded to use a platinum resistor layer, the M6 process can be easily adapted to that, although more characterization work would probably be necessary to ensure better critical dimension control.

4 Summary of Results and Discussion

4.1 Processing

To date, two TD1 lots, containing a total of 20 wafers, have been fabricated in the Lincoln Laboratory Microelectronics Laboratory (MEL), a class-10 clean room facility used for many types of experimental microelectronics processing.

As mentioned previously (see chapter 3.4), removal of the M2 adhesion layer via silox immersion proved fatal to the TD1 process, causing small M2 features and the edges of large M2 features to detach from the substrate and peel up. The first TD1 lot received a 30 sec. silox immersion to strip the M2 mask, and all submicron M2 features were destroyed on every wafer in the lot. Despite this catastrophe, processing continued and the wafers were fabricated to completion in the hope that larger junctions would yield.

The silox immersion time was shortened to 5 seconds for the second TD1 run, in the hopes that this would reduce the M2 lifting issue. It did not, however, and the lifting on the second run was essentially no better than the first. Following the failure of the second run, a process was developed to strip the LTO adhesion layer via plasma-etching, avoiding the silox immersion and eliminating the liftoff problem, but a full TD1 run incorporating this new process step has not yet been completed.

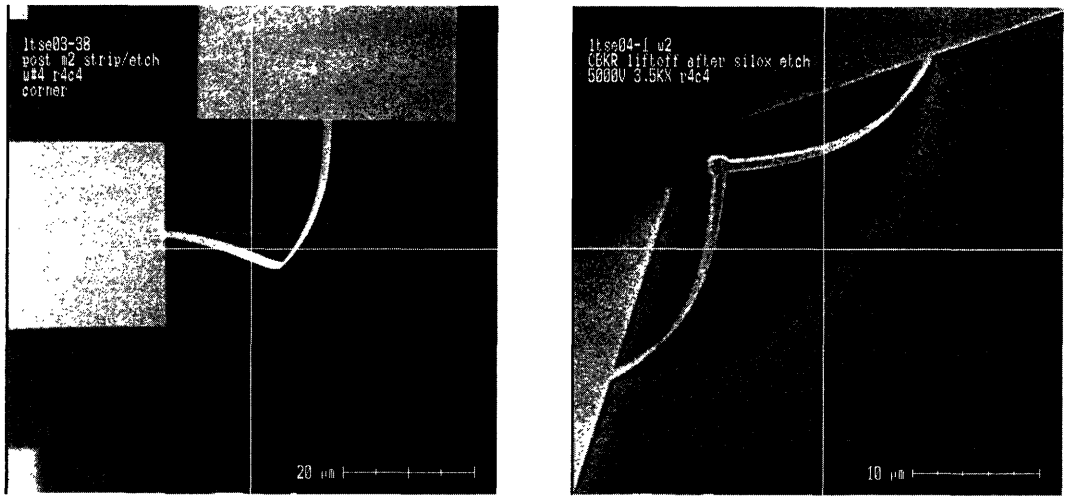

Fig 60: Scanning electron micrographs showing M2 lifting in first (left) and second TD1 lots. Note that the lifting is equally pronounced in the second lot despite the 85% reduction in silox etch time.

When the completed wafers were tested at room temperature (see chapter 4.2), the M2 and M4 features were found to be shorted together on nearly all wafers in both runs. Post-processing analysis of the first run suggested that the likely mechanism of this shorting was residual M4 "stringers" remaining on the M2 sidewalls after the M4 etch.

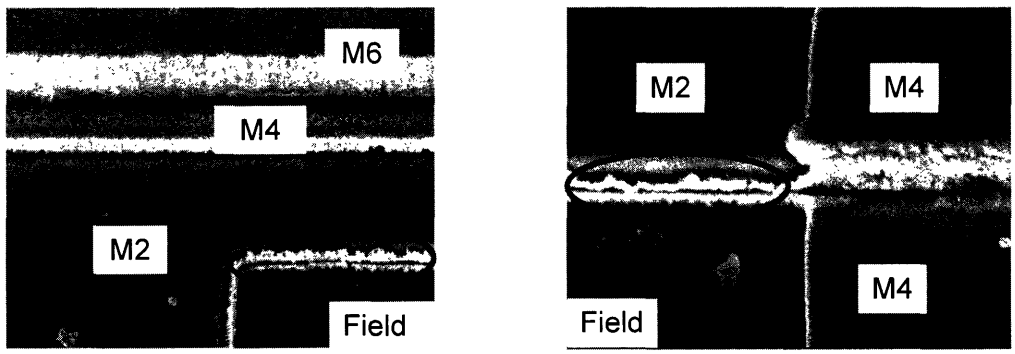

Fig. 61: Scanning electron micrographs taken of primary TD1 lot after processing. The M4 "stringers" along the edges of M2 are clearly visible, both near M4 steps (right) and along the edges of most M2 features.

Initially, it was assumed that this problem could be fixed by simply overetching the M4 layer well past detection of the endpoint. The problem with this "brute force" approach was that the SF_6-based plasma etch recipe used to etch M4 was extremely directional; it did not etch laterally well. This was desirable when vertical sidewall profiles were necessary, such as in the M3 etch, but made removal of residual metal on the sidewalls during the M4 etch problematic, and very long overetch times were needed to make it work. An extensive characterization of the SF_6 etch tool used for the TD1 process is currently in progress, with the goal of finding a more isotropic etch recipe to alleviate this problem.

Additionally, the M4 shorting problem was found to be linked to the M2 liftoff problem discussed earlier. When the edges of larger M2 features peeled up during liftoff, they allowed Nb to be deposited underneath the lifted regions during the M4 deposition step. The lifted M2 above this extraneous metal shielded it from etching and made it extremely difficult to remove. In the second TD1 run, severe overetching was successful in removing the sidewall "stringers", but the metal underneath the lifted M2 regions could not be removed and so shorting remained an issue.

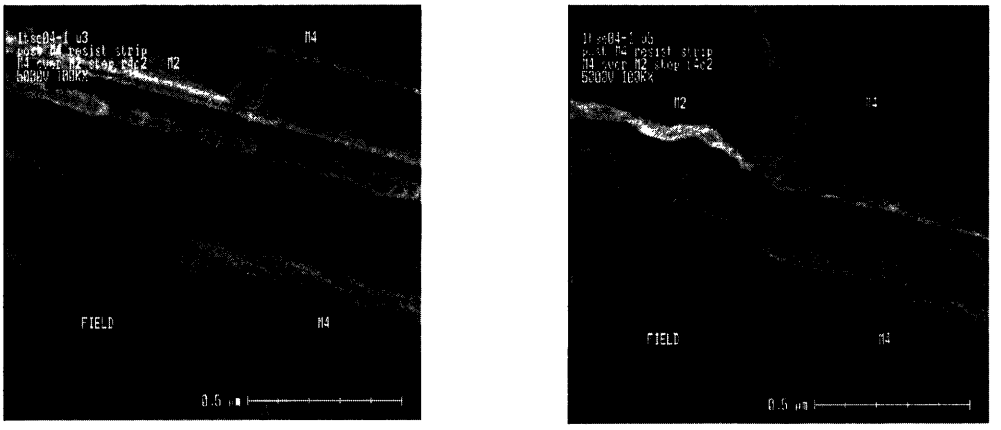

Fig. 62: Scanning electron micrographs of secondary TD1 run showing M4 encroachment under lifted M2 regions.

As a result of these processing issues, no testable devices were found on 19 of the 20 TD1 wafers. A wafer in the first lot, however, was overetched by over 3 min. (nearly twice as long as the etch itself) at the M4 etch step due to operator error. Though this overetch destroyed many of the more complex circuits on the chip, it also removed enough residual metal that several working Josephson junctions were found among the test structures.

4.2 Testing

After processing and dicing into individual chips, the TD1 wafers were first tested at room temperature. A test structure known as a cross-bridge Kelvin resistor (CBKR) was used to quickly determine the critical current of a Josephson junction without requiring measurement at cryogenic temperatures[53, 54]. Essentially a Josephson junction with four leads (two from the top contact and two from the bottom) leading to contact pads, the

room-temperature resistance of a CBKR, measured via four-point probing, can be used to estimate its critical current density at superconducting temperatures.

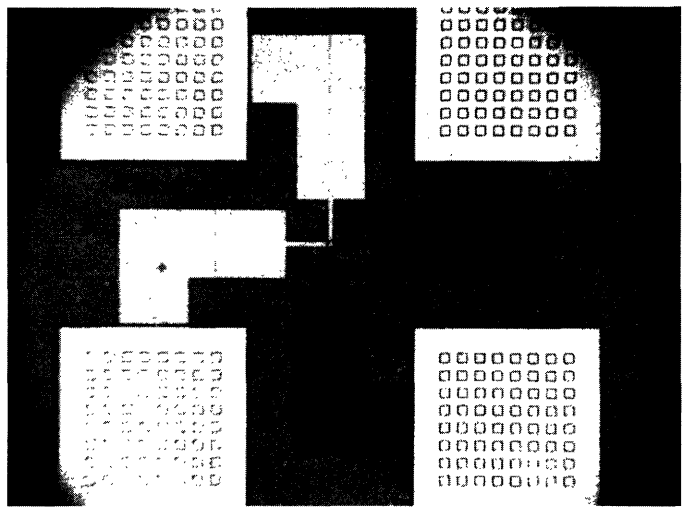

Fig. 63: Optical micrograph of a CBKR four-point test structure.

As mentioned previously, 19 of the 20 TD1 wafers processed were short-circuited to the point of having no testable devices. The remaining wafer (Wafer #1 from the first lot) sporadically yielded some simple test devices, however, and measurements of several CBKRs at different points on the wafer indicated that the critical current density was close to the target and that the junctions were not shorted, discontinuous, or otherwise damaged.

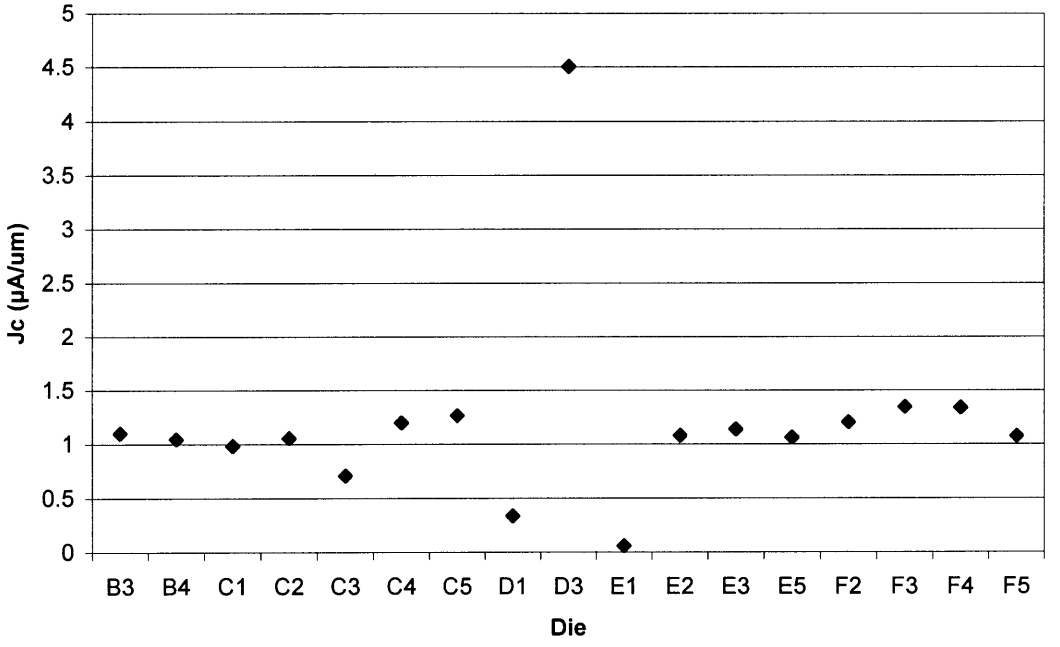

Fig. 64: Estimated critical current density as a function of wafer position, determined via CBKR measurements on wafer #1.

After determining from room-temperature measurements that wafer #1 was likely to contain at least some working junctions, it was cooled down to a temperature of 4.2K using a liquid-helium cryogenic system and several of the CBKRs were re-probed. Several devices showed the characteristic current-voltage behavior of a Josephson junction at superconducting temperatures.

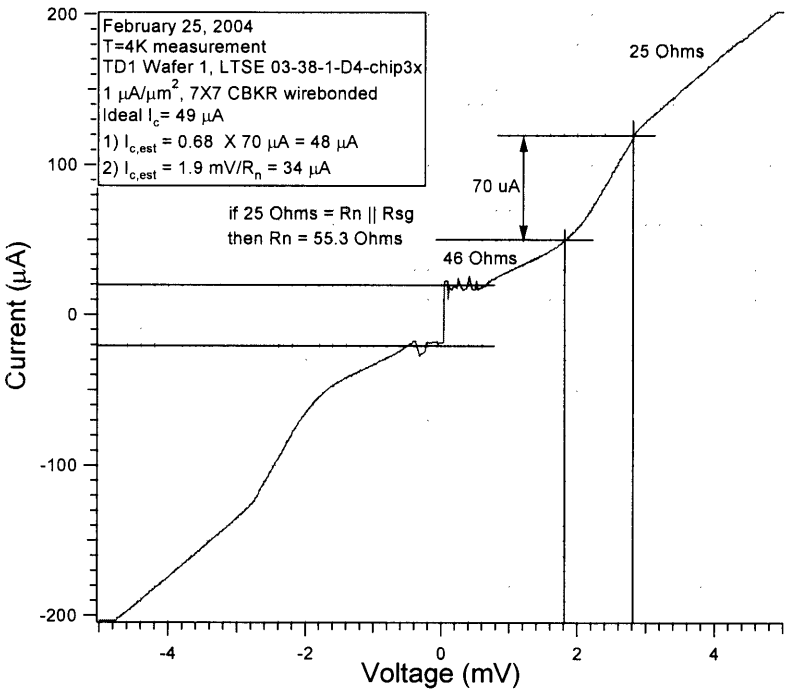

Fig. 65: Current-voltage characteristic of a 7 μm diameter Josephson junction measured at liquid-helium temperature (4K). Note that, while the critical current is nearly on target (see numbers in upper-left), the subgap resistance is extremely low, indicating that the junction is "leaky". This is most likely due to very small M4 stringers shorting the junction. (Graphic courtesy of T. Weir, MIT-LL)

While the surviving junctions were generally leaky and low-quality, the fact that they survived the various processing problems alone is significant. Moreover, the existence of working junctions proves that the TD1 process is not inherently flawed and can be used to fabricate working Josephson junctions if the processing problems it currently faces are solved. The majority of the processing problems that plagued the devices fabricated so far have been addressed, and future runs should be able to produce small junctions without serious leakage.

5 Further Investigation

While it has been demonstrated that working Josephson junctions can be fabricated using the TD1 process, substantial improvement and troubleshooting work remains before it can be used to generate deep-submicron devices, as well as qubits and other complex circuits.

The main issue plaguing the two fabricated TD1 lots was the liftoff of M2 features during the silox immersion to remove the adhesion layer. This problem has recently been solved by replacing the silox strip with a CF_4 plasma (see chapter 3.4), but has not yet been incorporated into a full TD1 processing run. Eliminating the liftoff should both preserve the small features that were lost in the first two TD1 runs and eliminate the most severe cause of the shorting (M4 under lifted M2, see chapter 4.1) that destroyed most of the remaining devices.

The M4 "stringers" along the M2 sidewalls is a more subtle problem, and one still being investigated. A partial solution is to develop an M2 etch recipe that creates a sloped (rather than vertical) sidewall. In addition to improving M4 step continuity (see chapter 3.2), this would allow the stringers to be more easily etched in a directional plasma recipe. Similarly, an M4 etch recipe with less directionality (etching laterally as well as vertically) than the current one would improve removal of residual metal from the M2 sidewalls. As mentioned at the end of chapter 3.2, an extensive design-of-experiment (DOE) characterization of the plasma-etching tool used for the M3, M2, and M4 etches is

underway, with the goal of determining the impact of pressure, temperature, and RF power on etch directionality and sidewall profile angles.

In addition to perfecting the TD1 process, a parallel project is currently underway to incorporate the deep-UV lithography techniques developed for TD1 into the full Lincoln Laboratory DPARTS-based process. The previous lithographic resolution of this process was roughly 0.75 μm; migration to from i-line to deep-UV lithography could improve it by nearly a factor of two. This modification would allow the fabrication of deep-submicron junctions in complex circuits requiring additional resistor and wiring layers.

Assuming the success of the etch-characterization experiment, the TD1 process should soon be able to quickly fabricate simple circuits with Josephson junctions as small as 0.3 μm. Furthermore, the problem of resolution limitations imposed by the incompatibility of ARC with our process may be solved in the future, as recent preliminary experiments have shown good results by using a low-temperature plasma to strip the ARC. If ARC is able to be incorporated into the process, achievable resolution may fall to 0.2 μm or smaller. The secondary goal of the project (rapid turnaround) was successful as well, as the process took approximately three weeks to run and could be easily streamlined to fewer than two as the process became more robust and less "hands on". This is a marked improvement over the three-month turnaround of the QC process and should prove useful for rapidly fabricating simple test devices in the future.

A Appendix: Design of Experiment[55]

Overview

Robust engineering and manufacturing process design commonly requirse detailed characterization of complex systems. Design of experiment (DOE) is a general methodology for this type of characterization, designed to replace inefficient and costly trial-and-error methods. In DOE, a complex system (such as a piece of fabrication equipment) is modeled as a combination of parameters (inputs) and responses (outputs). A matrix of experiments with different variations of input parameters is then designed, the outputs of which are interpolated to give a complete picture of the relationships between system parameters and responses.

Response Surface Design

Several methods exist for analyzing the results of a DOE experiment, but one of the most useful and popular is known as response-surface design (RSD). RSD fits quadratic polynomials to experimental data (via least-squares approximation) to create an analytic, multidimensional "surface" representation of the simultaneous effect of an arbitrary number of parameters on a given response. The analytical nature of the results produced with RSD is particularly useful, as it makes pinpointing minima and maxima (typically the goal of a DOE experiment) straightforward and simplifies extrapolation beyond the model limits.

Central Composite Design

Similarly, a variety of methods can be used to design the actual matrix of experimental runs in a DOE model. The most widely used is the Central Composite Design (CCD), although it can occasionally be advantageous to use other methods, such as the Box-Behnken Design (BBD) or full-factorial design. The CCD method uses three types of points in generating a run matrix:

- *Fractional factorial points*: runs in which all parameters (factors) are set at some point between their midrange and minimum/maximum values. In a three-parameter space, these points form the vertices of a cube.
- *Center point*: The point at which all parameters are set at their midrange values. This run is typically repeated at least once to verify the repeatability of the responses.
- *Axial Points*: For axial-point runs, all parameters but one are set to their midrange value, while the remaining parameter is set to its extreme maximum or minimum. In a three-parameter space, these points form a "star" shape and are sometimes called "star points".

For a system with k inputs, the fractional-factorial points will require 2^k runs, the axial points will require 2k runs, and the center point will require n_c runs, where n_c is usually at

least two to ensure repeatability. Thus, the total number of experimental runs in a k-dimensional CCD matrix is:

$$n = 2^k + 2k + n_c$$

For a three-dimensional system with one repetition at the center point, a CCD analysis requires 16 runs. By contrast, a 4-level full-factorial matrix (every possible combination of 4 different levels of parameters) requires 4^k runs, or 64 in this case, not counting the center point. RSD curve-fitting is nearly as effective with a CCD model as with a full-factorial model, so as the number of input parameters increases it becomes exponentially more efficient to use a CCD-based matrix (by contrast, a 5-dimensional experiment requires 44 runs in a CCD matrix and 1024 runs in a 4-level full-factorial matrix!).

Implementation

In order to obtain a detailed understanding of the SF_6- based plasma etching tool used for Nb etches in the TD1 process, DOE methodology was employed. The etch tool was modeled using the following parameters and responses:

- Parameters
    - RF power (watts)
    - Substrate temperature (°C)
    - Chamber pressure (mTorr)

- Responses
    - Etch rate (nm/sec)
    - Lateral etch/bias (nm)
    - Lateral etch uniformity (%)
    - Sidewall profile angle (degrees)

The lateral etch extent and uniformity was characterized by measuring the diameters of posts of several different sizes at various sites on each wafer, both before and after etch. By making pre-etch measurements, any extraneous bias or nonuniformity produced by the lithography process could be accounted for and ignored.

JMP, a computer program designed for DOE and other statistical-analysis applications, was used to generate a CCD matrix based on the three parameters above (Table 1). The minimum and maximum values of each parameter were based on the standard recipe values (50W, 50°C, 25mTorr); the range used for each parameter was between 0.5X and 2X the standard value.

Wafer	Temperature (°C)	Power (W)	Pressure (mT)
1	20	25	12
2	20	100	12
3	20	25	50
4	50	100	31
5	80	100	12
6	80	100	50
7	80	25	50
8	50	63	31
9	50	63	50
10	50	25	31
11	80	63	31
12	80	25	12
13	20	100	50
14	50	63	12
15	50	63	31
16	20	63	31
17	50	50	25
18	80	50	25

Table 4: CCD matrix of experimental runs for SF_6 etch characterization.

The first 16 runs are the computer-generated CCD matrix; the last two are the two versions of the "standard" etch recipe already used in TD1 and were added manually, for a total of 18 runs. Note that the order of the runs is randomized by the computer to avoid systematic errors as much as possible. Unfortunately, practicalities of the etch tool (changing the electrode temperature can take several hours) forced the runs to be re-grouped by temperature, though the randomness of the other parameters was preserved.

The 18 wafers to be used were coated with 200 nm of thermal SiO_2 and 150 nm of Nb, patterned with the M3 mask layer from TD1, and etched to the SiO_2 endpoint with each recipe. SEM analysis was used to determine the feature dimensions and uniformity both

before and after the etch, and cross-sectioning and additional SEM analysis was used after the etch to determine sidewall angle. The etch rate was determined by simple division.

When fully completed and analyzed, the DOE run should provide a detailed picture of the effect of power, temperature, and pressure on the TD1 Nb etch. The data produced will be useful both for improving the current process (identifying the parameters that produce minimal bias and nonuniformity) and for implementing new processes with different requirements in the future.

References

1. S. Lloyd, *Quantum-Mechanical Computers.* Scientific American. **273**(4), 44 (1995).

2. S. Lloyd, *A Potentially Realizable Quantum Computer.* Science. **261**(5128), 1589 (1993).

3. C.H. Bennett, *Quantum Information and Computation.* Physics Today. **48**(10), 24 (1995).

4. D.P. DiVincenzo, *Quantum Computation.* Science. **269**(5234), 255 (1995).

5. T.P. Spiller, *Quantum Information Processing: Cryptography, Computation, and Teleportation.* Proceedings of the IEEE. **84**(12), 1719 (1996).

6. B. Schumacher, *Programmable Cesium Rydberg Wave Packets.* Physical Review A. **52**(6), 4719 (1995).

7. S. Lloyd, *Almost any Quantum Logic Gate is Universal.* Physical Review Letters. **75**(2), 346 (1995).

8. D. Deutsch, A. Barenco, and A. Ekert, *Universality in Quantum Computation.* Proceedings of the Royal Society of London Series A. **449**(1937), 669 (1995).

9. C. Monroe, D.M. Meekhof, B.E. King, W.M. Itano, and D.J. Wineland, Physics Review Letters. **75**, 4714 (1995).

10. N.A. Gershenfeld and I.L. Chuang, *Bulk Spin-Resonance Quantum Computation.* Science. **275**(5298), 350 (1997).

11. Q.A. Turchette, C.J. Hood, W. Lange, H. Mabuchi, and H.J. Kimble, *Measurement of Conditional Phase Shifts for Quantum Logic.* Physical Review Letters. **75**(25), 4710 (1995).

12. 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. **60**(22), 398 (1999).

13. V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. Devoret, *Quantum Coherence with a Single Cooper Pair.* Physica Scripta. **T76**, 165 (1998).

14. P. Joyez, A. Pilipe, D. Esteve, and M. Devoret, *Observation of Parity-Induced Suppression of Josephson Tunneling in the Superconducting Single Electron Transistor.* Physical Review Letters. **72**(15), 2458 (1994).

15. M. Matters, W.J. Elion, and J.E. Mooij, Physics Review Letters. **75**, 721 (1995).

16. M. Kohler, Etching in Microsystems Technology. New York: Wiley-VCH, 1999.

17. T.P. Orlando and K.A. Delin, Foundations of Applied Superconductivity. New York: Addison-Wesley, 1991.

18. P. Carruthers and M.M. Nieto, Revolutions in Modern Physics. **40**, 411 (1968).

19. M. Tinkham, Introduction to Superconductivity. 2 ed. New York: McGraw-Hill, 1996.

20. H. Grabert and M. Devoret, eds. Single Charge Tunneling. New York: Plenum, 1992,

21. P. Lafarge, M. Matters, and J.E. Mooij, *Charge Representation of a Small Two-Dimensional Josephson-Junction Array in the Quantum Regime.* Physical Review B. **54**(10), 7380 (1996).

22. D.V. Averin and K.K. Likharev, in *Mesoscopic Phenomena in Solids*, Edited by B.L. Altshuler, P.A. Lee, and R.A. Webbs. 1991, North-Holland: Amsterdam.

23. L.J. Geerligs, V.F. Anderegg, and J.E. Mooij, *Tunneling Time and Offset Charging in Small Tunnel Junctions.* Physica B. **165-166**, 973 (1990).

24. G. Zimmerli, T.M. Eiles, R.L. Kautz, and J.M. Martinis, *Noise in the Coulomb Blockade Electrometer.* Applied Physics Letters. **61**(2), 237 (1992).

25. M.G. Castellano, R. Leoni, and G. Torrioili, *Switching Dynamics of $Nb/AlO_x/Nb$ Josephson Junctions: Measurements for an Experiment of Macroscopic Quantum Coherence.* Journal of Applied Physics. **80**(5), 2922 (1996).

26. S. Han, J. Lapointe, and J.E. Lukens, in *Activated Barrier Crossing*, Edited by G.R. Fleming and P. Hanggis. 1993, World Scientific: Singapore. p. 241.

27. J. Diggins, T.D. Clark, H. Prance, R.J. Prance, T.P. Prance, J.R. Spiller, and F. Brouers, *Limits to the Observation of Coherent Oscillations in a SQUID Ring.* Physica B. **215**(4), 367 (1995).

28. R. Rouse, S. Han, and J.E. Lukens, *Observation of Resonant Tunneling Between Macroscopically Distinct Quantum Levels.* Physics Review Letters. **75**(8), 1614 (1995).

29. S. Han and J.E. Lukens, *Observation of Coherent Charge-State Mixing in Asymmetric Bloch Transistors.* Journal of Superconductivity. **12**(6), 3247 (1999).

30. M. Hidaka, H. Tsuge, and Y. Wada, *Thermal Stability of $Nb/AlO_x/Nb$ Josephson Junctions.* Advances in Cryogenic Engineering. **34**, 765 (1988).

31. T. Shiota, T. Imamura, and S. Hasuo, *Fabrication of High-Quality Nb/Al-AlO$_x$/Nb Josephson Junctions: III - Annealing Stability of AlO$_x$ Tunnel Barriers.* IEEE Transactions on Applied Superconductivity. **2**(4), 222 (1992).

32. A. Oliva and R. Monaco, *Annealing Properties of High Quality Nb/Al-AlO$_x$/Nb Tunnel Junctions.* IEEE Transactions on Applied Superconductivity. **4**(1), 25 (1994).

33. K.K. Berggren, E. Macedo, D.A. Feld, and J.P. Sage, *Low-T_c Superconductive Circuits Fabricated on 150-mm-Diameter Wafers Using a Doubly Planarized Nb/AlO$_x$/Nb Process.* IEEE Transactions on Applied Superconductivity. **9**(2), 3271 (1999).

34. T.P. Orlando, Private Communication. (2004)

35. K. Segall, D.S. Crankshaw, D. Nakada, B. Singh, J. Lee, T.P. Orlando, K.K. Berggren, N. Markovic, S.O. Valenzuela, and M. Tinkham, *Experimental Characterization of the Two Current States in a Nb Persistent-Current Qubit.* IEEE Transactions on Applied Superconductivity. **13**(2), 1009 (2003).

36. Y. Yang, D. Nakada, J. Lee, B. Singh, D.S. Crankshaw, T.P. Orlando, K.K. Berggren, and W.D. Oliver, *Energy Relaxation Time Between Macroscopic Quantum Levels in a Superconducting Persistent-Current Qubit.* Physical Review Letters. **92**(11), 117904 (2004).

37. R. Dolata, T. Weimann, H.-J. Scherer, and J. Niemeyer, *Submicron Nb/AlO$_x$/Nb Josephson Junction Fabricated by Anodization Techniques.* IEEE Transactions on Applied Superconductivity. **9**(2), 3255 (1999).

38. W.S. DeForest, Photoresist. New York: McGraw-Hill, 1975.

39. D.J. Elliott, Microlithography: Process Technology for IC Fabrication. New York: McGraw-Hill, 1986.

40. D.W. Hess and K.F. Jensen, eds. Microelectronics Processing: Chemical Engineering Aspects. Washington, DC: American Chemical Society, 1989,

41. S. Nonogaki, T. Ueno, and T. Ito, Microlithography Fundamentals in Semiconductor Devices and Fabrication Technology. New York: Marcel Dekker, Inc., 1998.

42. L.P.S. Lee, A.R. Arambula, G. Hanaya, C. Dang, R. Sandell, and H. Chan, *RHEA Process for Fine-Geometry Josephson Junction Fabrication.* IEEE Transactions on Magnetics. **27**(2), 3133 (1991).

43. H.I. Smith, Class Notes for 6.781: Submicrometer and Nanometer Technology. (2004)

44. M. Chen and Y. Lee, *Heating Effects in Reactive Etching of Nb and Nb_2O_5.* IBM Research Bulletin. **131**(9), 2123 (1984).

45. J. Sasserath, *Dry Etching of Niobium Using CCl_2F_2 and CF_4: A Comparison.* Journal of Applied Physics. **68**(10), 5324 (1990).

46. A.C. Seabra, P. Verdonck, W.L. Xavier, and V. Baranauskas, *Dry Etching of Niobium Oxide Thin Films.* SPIE Proceedings. **1185**, 80 (1989).

47. T. Imamura and S. Hasuo, *A Submicrometer $Nb/AlO_x/Nb$ Josephson Junction.* Journal of Applied Physics. **64**(3), 1586 (1988).

48. E.A. Garcia, *Fabrication of All-Nb SIS Tunnel Junctions for Submillimeter-Wave Detectors*, in *Electrical Engineering and Computer Science*. 1992, Massachusetts Institute of Technology: Cambridge, MA, USA.

49. D. Nakada, K.K. Berggren, E. Macedo, V. Liberman, and T.P. Orlando, *Improved Critical-Current-Density Uniformity by Using Anodization.* IEEE Transactions on Applied Superconductivity. **13**(2), 111 (2003).

50. N. Magnussen, L. Quinones, D.C. Dufner, D.L. Cocke, E.A. Schweikert, B.K. Patnaik, C.V. Barros Leite, and G.B. Baptista, *Analysis of Anodic Oxide Films on Niobium.* American Chemical Society. **1**, 220-225 (1989).

51. H. Kroger, L.N. Smith, and D.W. Jillie, *Selective Niobium Anodization Process for Fabricating Josephson Tunnel Junctions.* Applied Physics Letters. **39**(3), 280 (1981).

52. V. Bolkhovsky, *Private Communication.* (2004).

53. K.K. Berggren, M. O'Hara, and J.P. Sage, *Evaluation of Critical Current Density of $Nb/AlO_x/Nb$ Josephson Junctions Using Test Structures at 300K.* IEEE Transactions on Applied Superconductivity. **9**(2), 3236 (1999).

54. M. O'Hara and K.K. Berggren, *Room-Temperature Testing for High Critical-Current-Density Josephson Junctions.* IEEE Transactions on Applied Superconductivity. **10**(4), 1669 (2000).

55. G. Goodwin and R. Payne, Dynamic System Identification: Experiment Design and Data Analysis. New York: Academic Press, 1977.