Engineering noise resilient super-conducting qubits

Fabrication and control of fluxonium and cos (2phi) systems

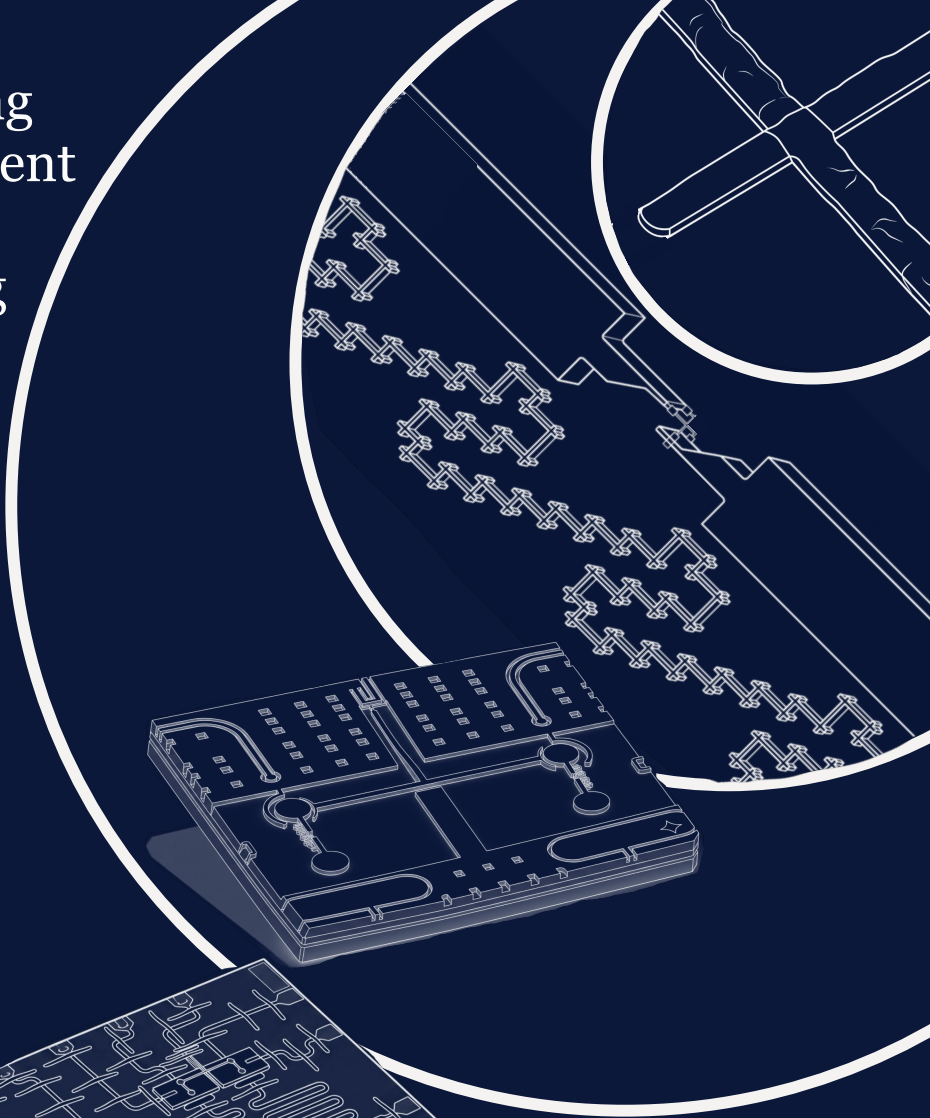

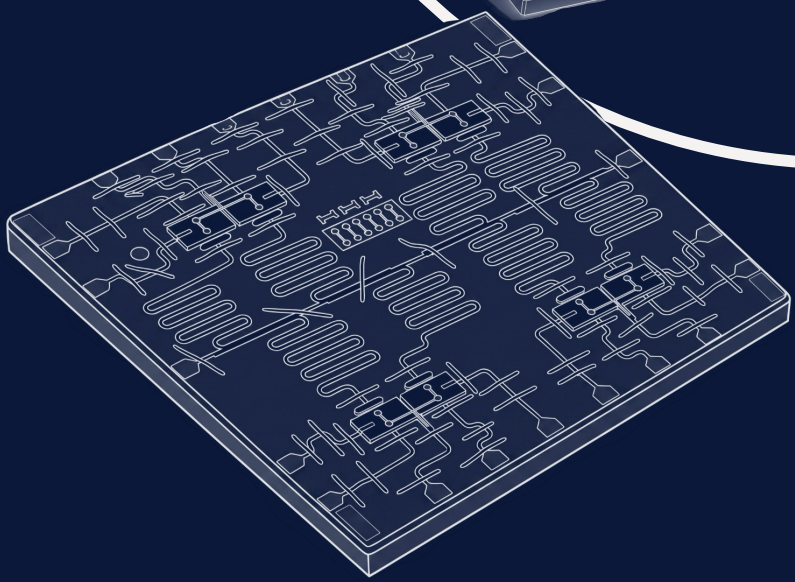

Siddharth Singh

Engineering Noise Resilient Superconducting Qubits: Fabrication and Control of fluxonium and cos(2phi) systems

ENGINEERING NOISE RESILIENT SUPERCONDUCTING QUBITS: FABRICATION AND CONTROL OF FLUXONIUM AND COS(2PHI) SYSTEMS

Dissertation

for the purpose of obtaining the degree of doctor at Delft University of Technology by the authority of the Rector Magnificus, Prof. dr. ir. H. Bijl, chair of the Board for Doctorates to be defended publicly on
Thursday 28 May 2026 at 10:00

by

Siddharth SINGH

This dissertation has been approved by the promotors and the copromotor.

Composition of the doctoral committee:

Rector Magnificus	Chairperson
Prof. dr. L. DiCarlo	Technische Universiteit Delft, promotor
Dr. C. K. Andersen	Technische Universiteit Delft, copromotor

Independent members:

Prof. dr. M. Blaauboer	Technische Universiteit Delft
Prof. dr. ir. M. Veldhorst	Technische Universiteit Delft
Prof. dr. M. Kjaergaard	U. of Copenhagen, Denmark
Dr. Q. Ficheux	CNRS, France

Reserve member:

Prof. dr. M. T. Wimmer	Technische Universiteit Delft

Keywords: Fluxonium, two-qubit gates, noise protected qubits, fabrication

Printed by: Gildeprint, Enschede

Cover by: Saurabh Singh

ISBN: 978-94-6518-327-5

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

Propositions

accompanying the dissertation

Engineering Noise Resilient Superconducting Qubits: Fabrication and control of fluxonium and cos(2phi) systems

by

Siddharth Singh

1. Material science will become a very critical field for industrial-grade fabrication of superconducting processors. *(Chapter 2 of this thesis)*
2. The superconducting qubit community should adopt more surface cleaning practices from the silicon industry. *(Chapter 3 of this thesis)*
3. Analytical pulse-shaping techniques are better than machine-learning-based pulse-shaping techniques. *(Chapter 5 of this thesis)*
4. SIS junctions are simple circuit components, but put in a specific configuration they create unique emergent phenomena, which cannot be directly explained by the physics of individual junctions. *(Chapter 6 of this thesis)*
5. Initial PhD training should include mandatory courses in developing AI driven digital tools concerning the research project.
6. To ensure the long term health and productivity of the scientific community, a healthy work-life balance must be prioritized during the PhD journey.
7. Quality supervision is a critical but undervalued professional skill in the research community.
8. Cleanroom intensive weeks should be followed by equivalent time off to improve fabrication quality.
9. Academic cleanrooms should prioritize hiring more staff with better pay, to ensure continuity of scientific innovation.
10. Establishing permanent, luxury housing for researchers can serve as a powerful catalyst for scientific innovation.

These propositions are regarded as opposable and defendable, and have been approved as such by the promotor Prof. Dr. L. DiCarlo and co-promotor Prof. Dr. C.K. Andersen.

To Mummy, Papa and Bitlu

CONTENTS

SUMMARY

Quantum computing has seen remarkable progress over the past decade, with superconducting qubits emerging as one of the leading hardware platforms for its physical realisation. Among the various superconducting qubit architectures, the transmon has been one of the most widely adopted candidates, owing to its relatively simple design and compatibility. However, the transmon suffers from key limitations, most notably its weak anharmonicity and multi-qubit gates using transmon being difficult to execute faster without errors.

In this context, alternative qubit architectures have attracted interest. The fluxonium qubit, with its large anharmonicity and long coherence times, has emerged as a promising candidate for overcoming shortcomings of the transmon. Similarly, the $\cos(2\phi)$ qubit offers an intriguing design with inherent noise protection, making it a compelling subject of investigation for next-generation superconducting hardware.

Beyond the question of which qubit to build, a central challenge in quantum computing is the realisation of high-fidelity two-qubit gates. Higher gate fidelity directly translates to better quantum operations and is essential for the implementation of error correction protocols. Equally important is the question of how to perform these gates in a fast and practical manner — ideally using simple pulse techniques that minimise the time and effort required for calibration.

This thesis addresses these themes across two broad directions. First, we explore the fabrication of fluxonium and $\cos(2\phi)$ qubits, developing improved processes and cleanroom techniques to achieve more reliable and reproducible devices. Second, we investigate two-qubit gate operations using analytic pulse techniques, with the goal of realising faster and simpler gates that reduce the overhead associated with pulse calibration. Together, these contributions advance the development of superconducting qubit platforms towards more practical and scalable quantum computing.

1

INTRODUCTION

"A great building must begin with the unmeasurable, must go through measurable means when it is being designed, and in the end must be unmeasurable."

Louis Kahn

From age-old mechanical tools like the soroban and abacus to the modern electronic chips, humans have continuously sought better ways to compute. This drive has transformed lives across the world by improving engineering across every domain. Since mid-twentieth century, Moore's law guided progress by shrinking transistors to pack more computing power onto chips. But this approach is reaching its limits — transistors can only get so small. This has pushed researchers toward new hardware architectures that can go beyond what classical computers offer. Quantum computing is one such avenue — rather than working around the limits of binary logic, it adopts quantum mechanics itself as the computational foundation. This shift opens the possibility of solving certain problems exponentially faster than any classical computers, offering a new frontier for computational power. Further, quantum computing also addresses another problem that classical computation does not scale to: the simulation of quantum systems. Modeling even small quantum systems demands exponential classical resources, making larger simulations practically out of reach. Using quantum systems to simulate other quantum systems can sidestep this bottleneck entirely [1], and this acts as a very compelling motivation for the field. The basic unit of quantum information is the qubit, a two-level quantum system at the heart of quantum computation.Qubits can be built from many different physical systems, as long as those systems can be isolated and precisely controlled. Key experiments in Nuclear Magnetic Resonance (NMR) [2, 3] and Cavity Quantum Electrodynamics (cQED) [4, 5] around the turn of the century demonstrated that quantum systems could be manipulated with high precision and exploited for quantum information processing.

These early results helped formalize the field. In the early 2000s, David DiVincenzo proposed five criteria — now called the DiVincenzo Criteria — that any physical platform must meet to work as a viable quantum computer[6]. Since then, researchers have worked to engineer solutions which can fulfill these criteria across several hardware platforms, including superconducting qubits, semiconductor spin qubits, trapped ions, and photonic systems.

This thesis focuses on superconducting qubits, specifically the fluxonium architecture. The work covers three core areas central to building better superconducting quantum processors: the design and properties of fluxonium and $\cos(2\phi)$qubits, their fabrication, and the implementation of two-qubit gates with fluxoniums.

1.1. SUPERCONDUCTING QUBITS:

Circuit Quantum Electrodynamics (cQED) applied the principles of light-matter interactions to electrical circuits. The field took off in 1988 when Martinis, Devoret, and Clarke showed that a Josephson junction — a thin insulating barrier

between two superconductors — could exhibit macroscopic quantum tunneling, meaning quantum mechanical behavior was visible at a scale far larger than individual atoms [7]. This was an important result, because quantum effects had rarely been seen in something as large and "everyday" as a circuit. It opened up the idea that superconducting circuits built from capacitors, inductors, and Josephson junctions could be engineered to produce specific quantum states on demand.

The development of superconducting circuits sits at the heart of cQED. Early groundwork was laid by Clarke, Martinis, and Devoret, who demonstrated both macroscopic quantum tunneling and quantized energy levels in superconducting circuits. The next major step forward came in 1999, when Nakamura and colleagues showed coherent quantum oscillations — Rabi oscillations — in a Cooper-pair box (CPB), making it the first superconducting qubit to be coherently controlled [8]. The CPB, however, was highly sensitive to charge noise, which limited how long the qubit could hold its quantum state. To fix this, Koch and colleagues at Yale introduced the transmon qubit in 2007 [9]. By adding a large shunt capacitor, the transmon greatly reduces sensitivity to charge noise while still being anharmonic enough to work as a qubit. Over the following two decades, research expanded to many different qubit designs, with a consistent focus on reducing decoherence through better materials and cleaner fabrication.

Two practical features have helped superconducting qubits build the necessary tools. First, they are made using thin-film deposition and lithographic techniques — the same basic processes used to make silicon chips. This means the field can draw on decades of semiconductor manufacturing knowledge and use existing fabrication tools to scale up. Second, individual qubits are physically large by quantum standards, typically hundreds of microns across. This makes it much easier to design them together, integrate them on a chip, and control them compared to atomic-scale systems. Quantum states in superconducting qubits are controlled using precision microwave electronics. The field has benefited greatly from decades of progress in radio-frequency engineering, and a growing number of specialized companies now build components specifically for quantum systems — from cryogenic hardware to fast microwave electronics. This commercial ecosystem has accelerated the move from one-off lab experiment to scalable, bigger quantum processors.

1.2. TRANSMON

The transmon qubit, introduced by Koch *et al.* [10], is one of the most widely adopted superconducting qubit architectures. It originates from the Cooper pair

box (CPB), whose Hamiltonian is given by

$$\hat{H}_{\mathrm{CPB}} = 4E_C(\hat{n} - n_g)^2 - E_J \cos\hat{\varphi}, \tag{1.1}$$

where $E_C = e^2/2C_\Sigma$ is the charging energy, E_J is the Josephson energy, $\hat{n}$ is the Cooper-pair number operator, n_g is the offset charge induced by the electrostatic environment, and $\hat{\varphi}$ is the superconducting phase difference across the junction. In the CPB regime, where $E_J/E_C \sim 1$, the qubit transition frequency depends sensitively on n_g, making it highly susceptible to charge noise—a dominant source of decoherence in solid-state environments.

The key insight behind the transmon is that by increasing the ratio E_J/E_C well into the regime $E_J/E_C \gg 1$, the charge dispersion—and hence the sensitivity to charge noise—is exponentially suppressed [10]:

$$\epsilon_m \propto e^{-\sqrt{8E_J/E_C}}, \tag{1.2}$$

while the anharmonicity decreases only algebraically as a weak power law of E_J/E_C. This favourable trade-off makes the transmon robust against charge fluctuations.

The large E_J/E_C ratio is achieved by shunting the Josephson junction with a large external capacitance C_S, as illustrated in Fig. 1.1(a). The equivalent circuit consists of a Josephson junction, characterised by its nonlinear inductance L_J and intrinsic capacitance C_J, connected in parallel with C_S. The total capacitance $C_\Sigma = C_J + C_S$ determines the charging energy $E_C = e^2/2C_\Sigma$, which is deliberately made small by the large shunt capacitance.

Figure 1.1(b) shows the resulting energy spectrum as a function of the superconducting phase φ. The potential energy landscape is a cosine, $U(\varphi) = -E_J \cos\varphi$, within which the lowest quantised levels 0, 1, and 2 are confined. Near the bottom of each cosine well, the potential can be expanded as

$$-E_J \cos\varphi \approx -E_J + \frac{1}{2}E_J\varphi^2 - \frac{1}{24}E_J\varphi^4 + \cdots, \tag{1.3}$$

where the leading quadratic term is identical to that of a quantum harmonic oscillator with frequency $\omega_p = \sqrt{8E_J E_C}/\hbar$, and the lowest energy levels are indeed nearly equally spaced. However, the quartic and higher-order corrections to the cosine potential introduce an anharmonicity

$$\alpha = \omega_{12} - \omega_{01} \approx -E_C/\hbar, \tag{1.4}$$

Despite its widespread success, the transmon architecture has notable limitations. The anharmonicity is typically only 200–300 MHz, which constrains the

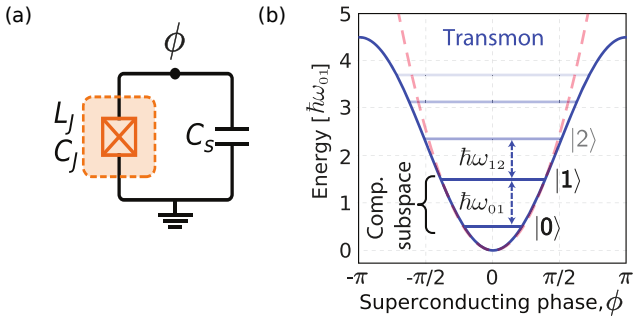

Figure 1.1: Transmon qubit. (a) Equivalent circuit consisting of a Josephson junction (with inductance L_J and capacitance C_J) shunted by a large capacitance C_S. (b) Energy spectrum as a function of the superconducting phase φ, showing the lowest three energy levels 0, 1, and 2 within the cosine potential. Reproduced with permission from [11].

speed of single-qubit gates—faster pulses have broader spectral bandwidth and risk leakage into higher energy levels. Furthermore, the transmon is susceptible to charge-noise-induced relaxation arising from dielectric loss in the circuit materials. This loss channel is characterized by the matrix element $|\langle 0|\hat{n}|1\rangle|^2$, which quantifies the coupling of the qubit to charge fluctuations in the environment. In the transmon, this matrix element remains substantial across the relevant parameter regime, making dielectric loss a dominant contributor to energy relaxation and a fundamental limitation on achievable T_1 times [10, 12].

1.3. FLUXONIUM

The fluxonium qubit, first introduced by Manucharyan [13], offers an alternative approach that overcomes both of the limitations of transmons by adding a high-impedance inductive shunt with the capacitor.

The fluxonium circuit, shown schematically in Fig. 1.2, consists of a single Josephson junction shunted by a large superinductance L and a small capacitance C. The superinductance is typically realised as a chain of large Josephson junctions or a high-kinetic-inductance nanowire, providing inductances of the order of 100–500 nH. The circuit is threaded by an external magnetic flux Φ_{ext}, which provides an additional tuning degree of freedom compared to the transmon.

The Hamiltonian of the fluxonium is given by

$$\hat{H}_{\text{flux}} = 4E_C\,\hat{n}^2 - E_J\cos\hat{\varphi} + \tfrac{1}{2}E_L\big(\hat{\varphi} - 2\pi\Phi_{\text{ext}}/\Phi_0\big)^2, \tag{1.5}$$

where $E_C = e^2/2C$ is the charging energy, E_J is the Josephson energy, $E_L = (\Phi_0/2\pi)^2/L$ is the inductive energy associated with the superinductance, $\Phi_0 = h/2e$ is the

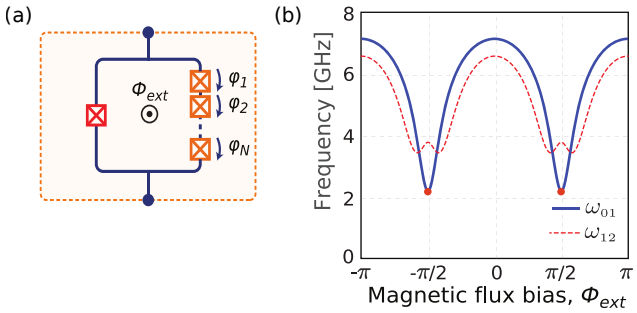

Figure 1.2: Fluxonium qubit. (a) Equivalent circuit consisting of a Josephson junction (with energy E_J and capacitance C) shunted by a superinductance L, threaded by an external flux Φ_{ext}. (b) Energy spectrum as a function of the external flux showing the lowest energy levels. Reproduced with permission from [11].

magnetic flux quantum, and $\hat{\varphi}$ and $\hat{n}$ are the conjugate phase and charge operators satisfying $[\hat{\varphi}, \hat{n}] = i$. Compared to the transmon Hamiltonian in Eq. (1.1), the fluxonium contains an additional quadratic term $\frac{1}{2}E_L(\hat{\varphi} - \varphi_{\text{ext}})^2$ arising from the inductive shunt, where $\varphi_{\text{ext}} = 2\pi\Phi_{\text{ext}}/\Phi_0$ is the reduced external flux. This term confines the phase variable to a single well of the cosine potential at zero flux, fundamentally altering the energy level structure compared to the transmon. At the half-flux-quantum sweet spot, the $0 \rightarrow 1$ transition frequency can be made very small—typically in the range of 100–1000 MHz—while the $1 \rightarrow 2$ transition frequency remains of the order of several gigahertz, The resulting anharmonicity can exceed several gigahertz, representing an order-of-magnitude improvement over the transmon. Furthermore, the qubit transition frequency has a first-order insensitivity to flux noise and the fluxonium has a reduced matrix element $|\langle 0|\hat{n}|1\rangle|^2$ compared to the transmon reducing losses. These properties have enabled fluxonium qubits to achieve T_1 times exceeding 1 ms in recent experiments [14], surpassing the coherence of state-of-the-art transmon devices.

1.4. IMPROVEMENTS IN SUPERCONDUCTING QUBITS

One of the important question facing the field of superconducting qubits today is how to assemble individual qubits into a large-scale quantum processing unit (QPU) capable of solving problems beyond the reach of classical computation. In the current noisy intermediate-scale quantum (NISQ) era [15], quantum processors contain tens to hundreds of physical qubits, each subject to errors from decoherence and imperfect gate operations. The path towards practical, large-scale quantum computation demands fault tolerance, which in turn requires quantum error correction (QEC) [16, 17]. In a QEC scheme, a single logical qubit is encoded redundantly across many physical qubits, enabling the detec-

tion and correction of errors without destroying the encoded quantum information. Prominent examples include the surface code [18], which arranges physical qubits on a two-dimensional lattice with nearest-neighbour interactions—a geometry naturally suited to superconducting circuits. However, QEC imposes stringent requirements on the underlying physical hardware. Each physical qubit must exhibit an error rate below a critical threshold, typically of the order of 1% for the surface code. Current state-of-the-art superconducting processors operate near or only marginally below this threshold [19]. This means that even modest improvements in physical qubit performance will translate directly into a dramatic reduction in the resources required for fault-tolerant computation.

In implementation of surface codes, reducing physical qubit losses has a compounding benefit: it moves the error rates further below threshold but also allows the same logical error rate to be achieved with a smaller distance codes, thus reducing the physical qubit overhead. There is therefore a strong incentive to push physical qubit performance.

The second requirement concerns qubit operations. Implementing error correction requires repeated cycles of single- and two-qubit gate operations alongside qubit measurements, and the fidelity of these operations sets a hard limit on code performance. Lower gate error rates provide more margin below threshold and allow to reduce the logical error rate.

THIS WORK

The two motivations—better physical qubits and better qubit operations—define the scope of the work presented in this thesis. For the qubit platform, we chose fluxonium. The key advantages are its higher anharmonicity and its lower dielectric loss in the qubit frequnecy range. Realising these advantages required building the complete fabrication process. Chapter 2 describes the development of the fabrication processes for devices. Chapter 3 covers the characterisation and optimisation of microwave losses in resonators. Chapter 4 addresses Josephson junction fabrication for fluxoniums.

With the hardware platform established, Chapter 5 addresses two-qubit gate operations, investigating how analytically derived pulse shaping techniques can suppress leakage and control errors.

Finally, Chapter 6 presents experiments on the $\cos(2\phi)$ qubit, a circuit whose Hamiltonian is engineered to suppress noise. This represents a complementary route to improving physical qubit quality, one that operates at the hardware level.

REFERENCES

[1] Richard P. Feynman. “Simulating Physics with Computers”. In: *International Journal of Theoretical Physics* 21.6–7 (1982), pp. 467–488. DOI: 10.1007/BF02650179.

[2] Lieven M. K. Vandersypen et al. “Experimental Realization of Shor’s Quantum Factoring Algorithm Using Nuclear Magnetic Resonance”. In: *Nature* 414 (2001), pp. 883–887. DOI: 10.1038/414883a.

[3] Jonathan A. Jones and Michele Mosca. “Implementation of a Quantum Algorithm on a Nuclear Magnetic Resonance Quantum Computer”. In: *Journal of Chemical Physics* 109.5 (1998), pp. 1648–1653. DOI: 10.1063/1.476739.

[4] J. M. Raimond, M. Brune, and S. Haroche. “Manipulating Quantum Entanglement with Atoms and Photons in a Cavity”. In: *Reviews of Modern Physics* 73.3 (2001), pp. 565–582. DOI: 10.1103/RevModPhys.73.565.

[5] M Brune et al. “Quantum dephasing and the Ramsey interferometer in cavity quantum electrodynamics”. In: *Physical Review Letters* 76.11 (1996), p. 1800.

[6] David P DiVincenzo. “The physical implementation of quantum computation”. In: *Fortschritte der Physik: Progress of Physics* 48.9–11 (2000), pp. 771–783.

[7] John Clarke et al. “Quantum mechanics of a macroscopic variable: the phase difference of a Josephson junction”. In: *Science* 239.4843 (1988), pp. 992–997.

[8] Yasunobu Nakamura, Yuri A Pashkin, and Jaw-Shen Tsai. “Coherent control of macroscopic quantum states in a single-Cooper-pair box”. In: *Nature* 398.6730 (1999), pp. 786–788. DOI: 10.1038/19718.

[9] Jens Koch et al. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Physical Review A* 76.4 (2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319.

[10] Jens Koch et al. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Phys. Rev. A* 76 (4 Oct. 2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319. URL: https://link.aps.org/doi/10.1103/PhysRevA.76.042319.

[11] Philip Krantz et al. “A Quantum Engineer’s Guide to Superconducting Qubits”. In: *Applied Physics Reviews* 6.2 (2019), p. 021318. DOI: 10.1063/1.5089550.

[12] G. Ithier et al. “Decoherence in a Superconducting Quantum Bit Circuit”. In: *Physical Review B* 72 (2005), p. 134519. DOI: 10.1103/PhysRevB.72.134519.

[13] Vladimir E. Manucharyan. “Superinductance”. PhD thesis. Yale University, 2009.

[14] Aaron Somoroff et al. “Millisecond Coherence in a Superconducting Qubit”. In: *Physical Review Letters* 130 (2023), p. 267001. DOI: 10.1103/PhysRevLett.130.267001.

[15] John Preskill. “Quantum Computing in the NISQ Era and Beyond”. In: *Quantum* 2 (2018), p. 79. DOI: 10.22331/q-2018-08-06-79.

[16] Peter W. Shor. “Scheme for Reducing Decoherence in Quantum Computer Memory”. In: *Physical Review A* 52.4 (1995), R2493–R2496. DOI: 10.1103/PhysRevA.52.R2493.

[17] Andrew M. Steane. “Error Correcting Codes in Quantum Theory”. In: *Physical Review Letters* 77.5 (1996), pp. 793–796. DOI: 10.1103/PhysRevLett.77.793.

[18] Austin G. Fowler et al. “Surface codes: Towards practical large-scale quantum computation”. In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324. URL: https://link.aps.org/doi/10.1103/PhysRevA.86.032324.

[19] Google Quantum AI. “Suppressing quantum errors by scaling a surface code logical qubit”. In: *Nature* 614.7949 (2023), pp. 676–681. DOI: 10.1038/s41586-022-05434-1.

2

NANOFABRICATION

"If you want to increase your success rate, double your failure rate."

Thomas J. Watson

In this chapter we discuss the working principles of the fabrication process used in the fabrication of superconducting qubit chips. Further we look at the optimized recipes using the tools inside Kavli Nanolab.

DECLARATION

Here is the list of people I've collaborated with for this chapter. :
Figen Yilmaz (etching, junctions and general optimization), **Sara Bhuktari** (Tantalum fab, 2-inch fab), **Alena S. Kazmina** (junction calibration), **Shikhar Singh** (resonator optimization, junction calibration), **Piran Kumaravadivel** (resonators optimization and air-bridges recipes), **Delphine Brousse** (general fabrication recipes and training), **Nikos Papadopoulos** (junction array structures), **Jana Bauer** (junction fab stack), **Lukas Splithoff** (Gatemon fabrication and junction calibration), **Yen-An** (2-inch fab and flip-chip), **Natasha** (Gatemon fab and SIS fab), **Anand Kamlapure** (junction optimization and tantalum fab).

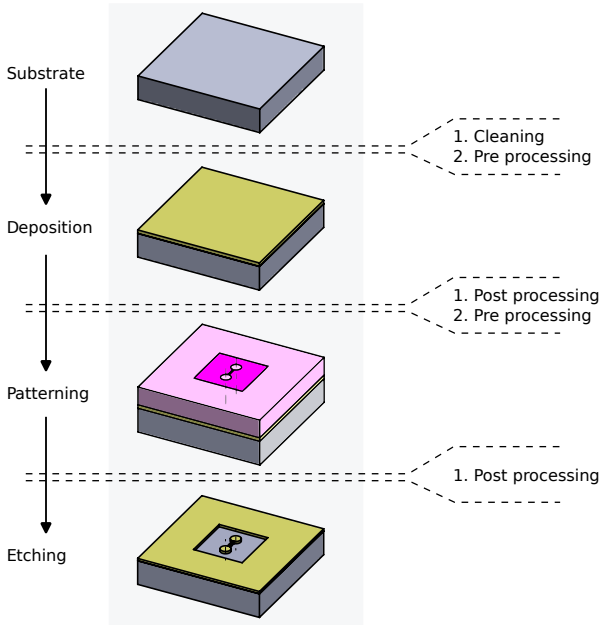

Figure 2.1: General fabrication steps: The fab flow is divided into three categories of fabrication processes. The figure outlines the sequence of the steps for only the baselayer, with intermediate steps of post processing and pre-processing marked between steps.

2.1. INTRODUCTION

Superconducting circuits have typical component dimensions ranging from nm to cm, and materials involving superconducting metals and substrates to make the circuits. In cleanrooms, these circuits can be fabricated using a combination of additive and subtractive fabrication processes. In the schematic shown in fig. 2.1, the overview of the fabrication processes steps for creating a superconducting QPU baselayer is shown. The processes used are divided into three major categories, namely: (1) Substrte preparation (2) Deposition, (3) Lithography, and (4) Etching. Each of these processes are further accompanied by pre and post processing steps.

Given the nature of the QPUs to be extremely sensitive to the presence of noise sources, it is of utmost importance to pay careful attention to the fabrication of the chip. Further, adding to this complexity is the academic nature of the cleanroom design, increasing the complexity of process control over a long time. Thus, this chapter describes all the mechanisms used with the processes, advises on strategies to be used for minimizing uncertainty and lists the recipes developed.

2.2. CHOICE OF SUBSTRATE:

The substrate used in fabricating superconducting qubits serves as the foundational material upon which the quantum circuit is constructed using a thin metallic film. This film, typically less than 250 nm thick, cannot be patterned into the desired circuit geometry in free space for two main reasons (1) The structural integrity of the film is insufficient to support itself, specially in the intricate features required for the circuit design, and (2) There are no practical methods available to form the film into the necessary geometry without a supporting base. Therefore the substrate is chosen based on criteria satisfying both material processing and processing compatibility:

1. **Loss tangents**: The loss tangent characterizes the overall energy dissipation properties of the material, quantifying the amount of power lost inherently. Since substrate selection is typically limited by commercially available materials, Silicon and Sapphire are the selection of substrates. Silicon has a relative dielectric constant of $\epsilon = 11.6$ and sapphire has $\epsilon = 9.6$ in the plane of propagation. This high dielectric constant leads to significant amount of EM energy being stored in the substrate. Consequently substrate's loss tangent limits the bulk dielectric loss for superconducting circuits. The loss tangent of a substrate is temperature-dependent and varies as the circuit is cooled, a behavior thoroughly investigated for silicon [1]. We experimented with both sapphire and silicon substrates, motivated by their intrinsically lower bulk loss tangents. The difference of the bulk loss tangents between silicon and sapphire was explored in depth through dipper experiments [2]. Both silicon and sapphire substrates are typically supplied as wafers, available in various sizes (e.g., 2-inch, 4-inch) and produced using different crystal growth methods. Wafers grown using different methods can differ in impurity levels, resistivity, and consequently in their loss tangents. Even within the same growth method, variations can arise from batch to batch. Furthermore, differences can be observed even within a single batch. These factors introduce variability in material properties, complicating the consistency of reported loss tangent values.

   Pushing this to the extreme, data shown in [3] shows the spatial dependence of the loss tangent within a single wafer. It was observed that the loss tangent tends to increase from the center toward the edge of the wafer. Thus when your quantum circuit is limited by the bulk losses of the substrate, the coupons selected from the center may exhibit better loss characteristics.

2. **Chemical stability:** The chemical behavior of the substrate is an important factor to consider. Over time, substrates may form native oxides or nitrides

upon exposure to the environment, but ideally, their chemical properties should remain stable over time. If the substrate does form oxides or other chemical surface layers, it should still be compatible with surface treatments that enable effective cleaning, without significantly altering the surface morphology. Maintaining a clean, stable surface is essential for high coherence quantum devices.

Additionally, when a metal layer is deposited on the substrate, the substrate should exhibit low chemical reactivity to prevent continuous changes in interface stoichiometry or chemical composition over time. Such changes can degrade the electromagnetic properties of the device, leading to reduced performance or stability. However, the substrate should not be extremely chemically inert, such that it becomes difficult to process. The material should be chemically stable enough to ensure long-term reliability, yet reactive enough to allow necessary fabrication steps like etching and cleaning.

3. **Mechanical stability** When a metal film is deposited onto a substrate, it inherently develops intrinsic stress, which arises from factors such as lattice mismatch, deposition parameters, and the resulting micro-grain structure of the film. These stresses are set during the growth process and can vary based on conditions like deposition temperature, rate, and ambient pressure. The intrinsic stress of thin films can be characterized through a laser based stress-meter in Kavli nanolab. The intrinsic stress is minimised by tracking the sputter target I-V characteristics over the life of the target. Any significant spike or deviation from the expected behaviour can be used to diagnose a mismatch in the stress levels.

   Beyond deposition, additional stress can be introduced during chip processing. Thermal cycling or uneven mechanical clamping can also lead to mechanical stress. This stress is generally reduced by clamping the substrate with even rounded surfaces and using smaller rate of change of tempeartures during processing.

   Further, the thermal stress introduced during device operation becomes critical. As the chip is cooled from room temperature down to milli kelvin temperatures (e.g., 8 mK), materials contract at different rates due to mismatched coefficients of thermal expansion. This significant thermal gradient can result in stress accumulation at interfaces, potentially causing film delamination, cracking, or changes in grain structure. These effects are especially problematic for superconducting devices, where mechanical deformations can impact the phonon environment and contribute to quasi article generation, thereby degrading qubit performance [4]. One of

the methods to avoid thermal change induced stress is to use uniform thermal contact for the chip and avoid using glue based adhesion of the chip, as suggested in [4].

To ensure reliable device behavior, it is essential to minimize all three types of stress - intrinsic, processing-induced, and thermally induced. This places a strong criteria on the substrate to be compatible with all of these kind of stresses.

During the course of this thesis, we used two main substrates - High resistivity Silicon(100) and c-plane Sapphire. The properties are listed in more detail

SILICON

We use silicon wafers with a (100) crystal orientation, grown using the Float Zone (FZ) method. This technique produces a high-purity single-crystal ingot in free space, starting from a seed crystal. Because the molten silicon does not come into contact with a crucible during growth, the risk of contamination from crucible materials is eliminated. Additionally, the FZ method significantly reduces defect-related nucleation events and grain boundaries, leading to superior crystal quality.

The crystal orientation of (100) direction is well-suited for planar device fabrication due to its diamond cubic lattice structure, which results in an atomically flat and stable surface. Compared to (110), the (100) orientation offers a more uniform and controllable etch profile, which is advantageous for precise microfabrication. Its widespread adoption in the semiconductor industry, driven by over five decades of CMOS development means that processing techniques for (100) silicon are highly mature and well-optimized. Furthermore, silicon (100) tends to form a more uniform and higher-quality native oxide, with smoother surface morphology following cleaning or surface treatments, compared to (110) oriented silicon [5]. Losses can still occur due to boron defects, already at very low defect density [6].

SAPPHIRE

For our devices, we use C-plane sapphire wafers of 2-inch and 6-inch diameters, employing different growth techniques to optimize performance. Sapphire's bulk loss tangent(10^{-8} to 10^{-9}) is lower than that of high-resistivity silicon(10^{-7} to 10^{-4}) at cryogenic temperatures. and its lattice constant better matches for niobium and tantalum thin films [7]. Depending on the growth method the substrate's dielectric loss varies, with HEMEX sapphire exhibiting the lowest measured loss tangents [2].

All sapphire substrates have high chemical inertness, which permits more aggressive cleaning procedures but also makes them mechanically and chemically

difficult to etch. This stability translates to high thermal resilience but complicates dicing and patterning steps.

Electrically, sapphire's dielectric constant is anisotropic and depends on the crystal orientation of the wafer cut. As a very poor conductor with resistivity on the order of 10^{14}–$10^{16}\Omega \cdot$cm, sapphire struggles to dissipate surface charge, leading to ESD challenges during fabrication. Optically, sapphire ranges from transparent to translucent based on polish quality, which can introduce handling difficulties in the cleanroom, due to inability to easily distinguish front and backside of the wafer. The choice of SSP wafer over DSP wafer must be made for sapphire to eliminate this handling issue.

Table 2.1: Comparison of Chemical Behavior: Silicon vs. Sapphire Substrates

Property	**Silicon**	**Sapphire (Al_2O_3)**
Crystal orientation	100	c-plane
Native oxide	Forms SiO_2 (amorphous, lossy due to TLS)	Chemically stable surface
Etchability	Easily etched with standard processes (e.g., HF, DRIE)	Chemically resistant; - requires special etching (e.g., hot H_3PO_4 or ion milling)
Surface reactivity	High; prone to contamination and moisture absorption	Low; chemically inert and less hydrophilic
Dielectric loss	Higher, due to amorphous oxide and surface contamination	Lower, due to crystalline structure and inert surface

2.3. DEPOSITION

Critical to the performance of the superconducting qubit is the choice of superconducting material. We deposit a thin film of NbTiN on top of Hi-resistivity silicon for the majority of the qubits in the lab, additionally, we experimented with other combinations of getting a Tantalum layer on both silicon and sapphire 3.3.

Metal deposition on substrate is done in three steps. Firstly, to prepare the substrate wafer for the metal deposition, the surface is cleaned to remove any dirt particles. Followed with a surface treatment to remove potential oxides from the surface. The removal of the oxide is critical to the performance of the qubits, due to the nature of TLS losses springing from this interface 3.1. Thus with two

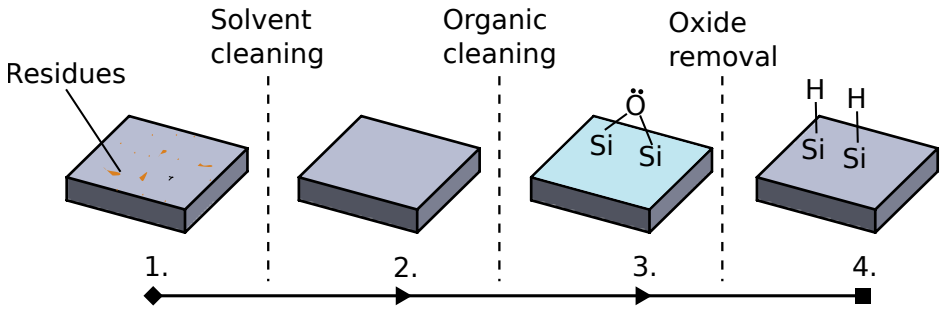

Figure 2.2: Pre-processing : Steps before the deposition are shown to prepare the substrate to deposit the superconducting metal. (1) The substrate piece is shown with some organic residues on the surface. (2) After solvent based clean dissolution of the residues happen on bigger scale. (3) Organic cleaning with acid is performed with highly oxidized surface shown with oxygen termination for Silicon case. (4) Oxide removal leaves the surface with a hydrogen terminated surface.

pre-processing steps, the full process has three steps:

1. Wafer cleaning organics
2. Wafer surface preparation
3. Metal deposition process of sputtering

2.3.1. PRE-PROCESSING

From the controlled environment of manufacturing plant, once the wafer has been sliced and polished, it is cleaned with RCA or SC solutions [8, 9]and then packaged by the manufacturer. Once a wafer has been properly packaged and shipped, it remains free from particulate matter. However, once the wafer box is opened, subsequent contamination takes place during shipping, handling and storage of the wafers.

To minimize and control this exposure, the box must be opened in the wet bench hood and all the transfer must take place at the same spot. However, this will not eliminate the exposure to ambient particles during storage. The wafer must be cleaned to get rid of a broad range of organic particulate matter, which has potentially accumulated on top of the wafer.

The organics on the wafer surface are adsorbed organic materials, including hydrocarbons, airborne molecular contaminants, and residues from packaging or previous handling. The cleaning process is designed to breakdown the particles, dissolve them as much as possible into the used solvent and most importantly it is chemically compatible with the substrate. Two important considera-

tions for selection for the after effect of the cleaning process are - surface morphology and surface chemical composition after cleaning.

The general cleaning processes (RCA-1, SC-1, Piranha clean) focus on breaking down the organic into simpler chains and CO_2 to dissolve them in the solvent. Additionally, the process conditions should assist with detaching the broken particles from the surface. This detachment can be achieved by mechanical vibrations, chemical potential changes and higher mobility of particles and reactions aided by higher temperatures.

Majority of the organic contaminants in the cleanroom consist of polymer varieties. The process of breakdown of polymers can be achieved by using highly oxidizing agents, stripping electrons from the polymer chains, and making them more water soluble and hydrophilic. At high pH[10] the particles and the wafer get negatively charged, leading to repulsion, this weakens the Van-Der Waal forces and assists to lift-off the particles from the surface. Similar dislodgement effect can also be achieved with the application of mechanical agitation in form of sonnication. Megasonnication is used in the industry and desired for smaller contamination, but in Kavli cleanroom we're limited to ultrasonnication.

For the cleaning process of high resistivity silicon, we used 65% concentrated Nitric acid [11], which is a highly oxidizing agent, helping with the breakdown of the organics. The process is assisted with sonnication at highest power. Although the nitric acid is a good reagent to breakdown the carbon based particulate matter, the low pH of the nitric acid, can charge the wafer more positively, which does not assist with the particle lift-off. Further, the nitric acid does not create bubbles at room temperature to the same extent as the peroxide based cleaning agents. There can be bubble formation, from the nitrous oxide formation during reaction but not strongly present. The choice of Nitric acid has been based on the assumption that the RCA clean from manufacturer is good enough to clean the surface and our wafer processing may not require extreme cleaning, unless the particulates are visibly present on the surface.

For the cleaning of the Sapphire wafer we use a 1:3 mix of 49% peroxide and sulfuric acid. This process is good for cleaning due to the presence of very strong oxidizing agent, the peroxide. The continuous creation of oxygen, creates micro bubbles and leads to a gentle scrubbing effect. Additionally the mixing of the sulfuric acid with the peroxide is a highly exothermic reaction. The heat generated during the mixing necessitates slow mixing and careful handling, making it a deterrent for the selection of the process. But at the same time local self heating improves the efficiency of the process in breaking down of the organics. One drawback of the process is the relatively low pH still, that does not assist with the charging effects of particles.

The discussed cleaning processes are compatible with both Silicon and sap-

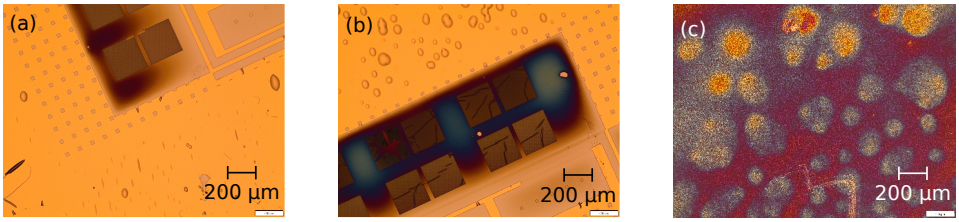

Figure 2.3: Improper water Rinsing (a),(b) Chip surface in the dark Coloured area is the damaged surface due to improper DIW rinsing. (c) Zoomed in zone of the chip surface with damage stains.

phire substrates. After metallization is done, we use the organic cleaning step one more time to clean the metal layer after RIE etching(sec 2.5.2). At this step RCA-1 is not compatible with the NbTiN, due to it's metal etching property.

WATER RINSE AND DRYING:

After the extensive organic cleaning, it's important to remove any trace chemicals on the wafer. This is accomplished by dipping the wafer in deionized water (DIW) beakers. It's dipped into two DIW beakers back to back, first a shorter dip dissolves the gross chemicals and subsequent long dip dissolute any trace chemicals. Further, the wafer can be rinsed further by running the wafer in a a continuous flow with the water gun. After the water rinse, the wafer is dried with nitrogen blow -drying. Careful attention must be given in this step, as trace carryover from oxidizing agent to reducing agent can have severe chemical reaction. In Fig. 2.3, reaction was observed on the wafer surface due to mixing of the oxidizing agent with the reducing agent. Even though the wafer surface is rinsed with water, trace chemical can still be stuck in forms of droplets near the holder or tweezers. The same step is repeated after oxide etching step, discussed in next section.

WAFER SURFACE TREATMENT:

Cleaning the wafer with the solvent for organics chemically alters the underlying substrate surface. The objective of surface treatment process is to prepare the surface chemistry, before the next step of metal deposition. In this process we do three sub-processes, depending on the substrate. First is cleaning of potential oxides, second is covering the surface with HMDS and additionally annealing in case of the sapphire.

In case of the silicon, cleaning with a strong oxidizing agent leads to a formation of silicon oxide. This oxide is deterrent towards high coherence due to the nature of silicon oxide to host significant amount of TLSs 3.1. To select the correct chemical for this process, considerations of chemical compatibility with

materials on the chip and effect on surface morphology and chemical composition of the outcome of the process must be considered.

A Hydrofluoric acid (HF) 40% solution is used for etching the oxide. The wafer is dipped in the HF beaker for 6 min long cleaning. The HF removes the native silicon oxide with a high etch rate [12], although the pH of the solution is relatively low. After the oxide etch, the wafer is again rinsed with DI water and dried. Rinsing is done using the hand water gun

HMDS

The oxide etching on silicon does not lead to a long lasting effect, infact re-oxidation of silicon starts within minutes after oxide etch, leading to silicon oxide reformation. To delay this oxide formation, HMDS is used immediately after the BOE step, extending the time allowed to load before sputtering [13]. The priming of the surface with HMDS works by the reaction with HMDS leading to a monolayer formation of tri-methylsilyl group on the wafer surface. This blocks immediate oxygen reaction with the surface silicon, leading to prolonged protection form native oxide. After HMDS priming, the wafer is immediately loaded into the sputtering chamber.

Process	**Chemical**	**Time**
Acid clean	HNO_3	7 mins; ultrasonnication power 9
Water rinse	Water gun	2 DIW beakers
Nitrogen gun drying	Nitrogen gun	
Oxide clean	HF 40%	6 mins
Water rinse	Water gun	2 DIW beakers
Nitrogen gun drying	Nitrogen gun	
Monolayer passivation	HMDS	recipe 3

Table 2.2: Current pre-processing recipe for silicon wafer cleaning

ANNEALING

In case of Sapphire substrates, higher quality factors have been observed for devices which have been annealed [14]. This process involves heating the bare wafer to high temperatures in an environment of oxygen or nitrogen gases. This heating leads to rearrangement of the grain boundaries, removal of defects and stress relief inside the substrate. In case of c-plane sapphire, the annealing creates atomically flat terrace like structures, allowing to create a more unifrom substrate metal interface. The following recipe was used to anneal the sapphire, based on [15]. This was mainly used during few deposition for the tantalum deposition flows.

Table 2.3: Annealing steps for substrate preparation in Tempress furnace.

Step	Condition	Duration
Initial heating	400°C in N_2 atmosphere	Ramp phase
Oxygen purge	O_2 flow introduced at 400°C	5 minutes
High-temperature anneal	1000°C in O_2 atmosphere	2 hours
Cooling phase	Substrates cooled in furnace	Overnight

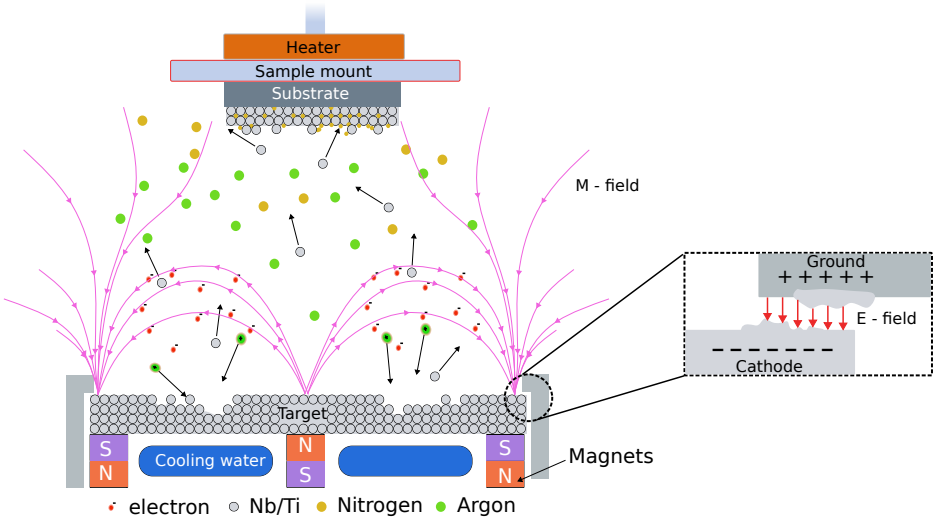

Figure 2.4: Sputtering: The sputtering chamber is shown with the plasma formation in process. Inset: shows the growth of metal flake features.

2.3.2. SPUTTERING:

In the duration of my PhD project, chips were mainly fabricated using three different kinds of superconducting films- Niobium (Nb), Niobium titanium Nitride (NbTiN) and Tantalum (Ta). NbTiN is an alloy with Type 2 superconductivity, while Tantalum and niobium are elemental superconductors. Both of these metals were deposited using magnetron sputtering technique. With the parameters of the material deposition techniques influencing the microstructure and the superconducting properties. In this section we further discuss the material deposition.

MAGENTRON SPUTTERING:

Sputtering is a physical vapor deposition (PVD) process in which atoms are layer by layer onto a substrate.

To understand the working of PVD, the material to be deposited is shaped

into a high-purity disc called target. A negative potential is applied to the target cathode, while the shield surrounding it acts as the anode. Argon is introduced into the chamber to create a plasma. When voltage is applied, the electrons ionize the gas, and the resulting electric field accelerates the ions toward the negatively charged cathode target. As the ions strike the surface, they eject electrons with secondary electron emission, helping sustain the plasma. Atoms are displaced from the target surface if the incoming ion energy exceeds the interatomic bonding energy.In this process, energy transfer is most efficient when the mass of the accelerated ions is similar to that of the atoms in the target material—this efficiency is defined as the sputter yield. Argon is therefore a suitable choice due to both its inert nature and its favorable mass.

The pressure of the gas inside the chamber, required to create the plasma, is an important factor during the deposition process. If the gas pressure is high, the mean free path of the ions decreases, reducing the energy an ion can impart to the target. Therefore, a balance between pressure and sputtering efficiency is desired. The sputtering process is made more efficient by applying a magnetic field near the target cathode, which traps more electrons in a locally created magnetic field. This increases the plasma density near the cathode and allows for more efficient sputtering, as the plasma remains concentrated near the target rather than spreading throughout the chamber. This configuration enables the creation of plasma at lower pressures and is known as magnetron sputtering.

Magnetron sputtering leads to better target utilization, denser films, and faster deposition rates. To further improve the quality of the metal film the atoms should have more mobility during crystal formation, for this purpose magnetron sputtering can be used in an unbalanced configuration. In this setup, the magnetic field lines do not close in on the target but extend beyond it, allowing electrons and ions to escape beyond the target area. This increases the ion current density on the substrate. Although the ions arriving at the substrate are not accelerated due to the substrate's neutral nature they still impart energy, modifying the surface energy. This change in surface energy enhances the mobility of the incoming target atoms after adhesion, resulting in an overall improvement in film quality.

Compressive stress: When we deposit the metal film, during the deposition the atoms have significant kinetic energy, thus heating the substrate and move on the surface of the substrate. However, once the deposition is finished, as the wafer cools down, the atoms in the film loses mobility and the film tries to contract to reduce interatomic distance. This mismatch in the substrate's tendency to retain it's shape and the metal film's tendency to contract leads to development of compressive stress in the film. This stress can be measured using the stressometer setup in the Kavli cleanroom. The stressometer measurements

must be done individually for all the wafers before and after deposition. Significant bending has been observed(>20μm) for the native wafer after deposition, in such cases the target I-V characteristics must be confirmed to check health. Additionally, more stress in the film can lead to delamination of the film and ruin the deposition of the wafer, see fig. 2.5(c),(d) and (e).

Target property changes: As the target material is depleted, the target thickness in the middle decreases in a horse track pattern. Careful observation should be kept to make sure the target does not deplete to too thin, otherwise it can damage the underlying magnets and water cooling structures, leading to more severe damage to the system. The stress measurements and resistance measurements can aid to probe the changes from the film, simultaneously the I-V characteristics indicate the amount of the depletion of the target compared to the previous target.

Further, the I-V characteristics also indicate potentially any presence of flake development. The cathode target is placed at fixed distance from the ground produces a fixed current value when the voltage is applied. The maximum current allowed on the current source is 999 mA, exceeding it means there is a short in the current path. As the target is sputtered, the material deposition decreases the distance between the cathode and the ground, changing the I-V characteristic. The value of stress changes from the start to a stable value in the middle and rapid decrease towards the end of the life of the target.

During the depositions, if there is observation of flickering plasma, that is an indication of arcing behavior, made possible by shorts happening from material flakes or rapid changes in the current leading to arcing. Arcing is very deteriorating for the system due to the sudden thermal shock experienced by the system, leading to developments of crack in the target cathode. At the same time it can accelerate the development of flakes in the intermediate space between the cathode and the ground.

NbTiN DEPOSITION ON SILICON SUBSTRATE

NbTiN is a disordered crystalline alloy with type -2 superconductivity, in which the crystal structure and film properties depends on the nitrogen content [16, 17]. We deposit a 200 nm thick NbTiN film at the highest rated power of 250W for the SuperAJA system. The target material for the NbTiN is a Nb:Ti 70:30 stoichiometric mix in their alloy form. When the deposition is done using the magnetron sputtering, the atoms for both Nb and Ti species are dislodged individually but reach an equilibrium to maintain the stoichiometry in the target film. For the creation of NbTiN, we introduce nitrogen gas into the chamber, which reacts with the NbTi on the surface of the substrate and forms a NbTiN film. The ratio of Argon to nitrogen is (50:3.5) for majority of the depositions done during the project,

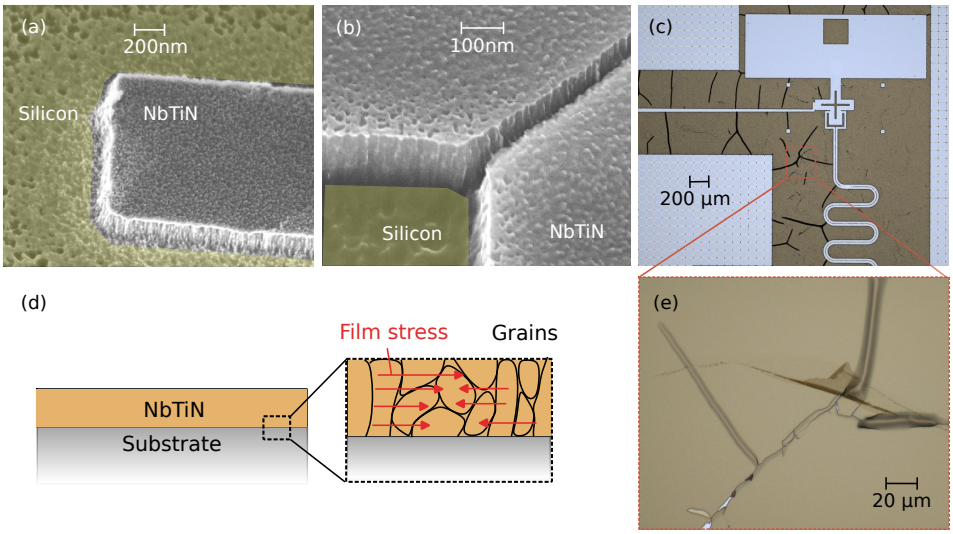

Figure 2.5: NbTiN film characteristics: (a),(b) SEM image of an etched NbTiN film island with silicon as the stop layer (c) Micrograph of a cracked NbTiN layer on the chip (d) Illustration of the stress buildup int he grains of the NbTiN film (e) Zoomed in micrograph of the (c)

but it can be changed to adjust for the desired properties. In Fig. 2.5(a) and (b), the grains of the NbTiN film are shown. Due to room temperature deposition, the grains growth potentially following the Thornton structure zone model [18], where the substrate temperature is less then 30% of the melting temperature of the deposited material (T/T_m<0.3).

Four probe measurements on the fixed thickness of 200 nm NbTiN films show a resistance between 1 to 4 Ω, depending on when the target was replaced and utilised. Kavli has two different sputtering systems, the SuperAJA with a circular 3 inch target of (70:30) stoichiometry and the Nordiko system with a capsule target of (80:20) stoichiometry. We used the SuperAJA system throughout the project, but Nordiko can be used in future to explore different material growth properties.

Table 2.4: Sputtering steps for NbTiN

Step	**Condition**	**Duration**
Pre-sputtering	250 W; 50 sccm Ar + 3.5 sccm N_2 (2.3 mTorr)	2 mins
Sputtering	250 W; 50 sccm Ar + 3.5 sccm N_2 (2.3 mTorr)	200 nm

TANTALUM DEPOSITION

Tantalum is an elemental type-1 superconductor which can grow in two different phases - α phase and β phase, with α having a higher T_c. When Tantalum is deposited in thin films using sputtering, it grows in the β phase but the losses in the β phase are detrimental for low photon number regime of a superconducting device. Thus the aim of optimizing the tantalum growth for superconducting devices is to make a reproducible α phase.

In our experimentation, tantalum is deposited on Sapphire using the same magnetron sputtering process in SuperAJA. A 2 inch target is used, unlike the NbTiN. The deposition rate is < 10nm/s making the deposition of a 200 nm film a relatively longer deposition process. The stability of the plasma is important due to the long deposition, furthermore the target depletion is higher due to the slow deposition rate but longer deposition time.

For the creation of the correct alpha phase of the tantalum on top of sapphire we did experiments in the growth conditions, mainly depositing at a higher temperature and introducing a seed layer of Nb. Further we also developed a deposition process for depositing Niobium, capped with Tantalum to see the effect of a thin capping on the performance 3.3. The different approached with their deposition parameters are listed in this table.

Approach	**Substrate**	**Metal**	t (s)	P (W)	p (mTorr)	Ar (sccm)	T (°C)
Heated	Sapphire	Ta	1200	150	2.1	50	475
	Sapphire	Ta	1500	150	2.4	50	700
Seedlayer	Sapphire	Nb	90	180	2.3	50	RT
	Sapphire	Ta	1500	150	2.3	50	RT
	Silicon	Nb	90	200	2.1	25	RT
	Silicon	Ta	1350	150	2.1	25	RT
Encapsulation	Silicon	Nb	1500	180	3.0	25	RT
	Silicon	Ta	85	150	2.3	25	RT

Table 2.5: Deposition parameters for tantalum depositions using different approaches

TANTALUM DEPOSITION AT ELEVATED TEMPERATURES:

The body centred cubic (bcc) α Ta is a stable phase at higher temperatures, while the β Ta is a metastable phase easily formed at lower temperatures. The substrate heated to higher temperatures allows for the migrations of the adatoms adsorbed on the surface of the substrate. This leads to rearrangement into more energetically favorable crystal orientation.

The heating mechanism in the SuperAJA works by placing a halogen lamp besides the substarte holder. The lamp is a halogen lamp working by radiative

heating, see fig. 2.4.Since the heat flow travels from the heater to the sample mount first and then to the substrate through contact, in the heating phase, a slow heating rate of 2°C/min is adopted to remove improper thermal gradients and heat the substrate uniformly. Further to improve the heat flow between the sample mount and the substrate, a flow of argon at 20 sccm is flowed inside the chamber.

In fig. 2.6 (a) and (b), SEM pictures of the Ta grains are shown for growth at 475 and 700 °C. At 475, the structure is more grainy and disordered, resembling a more β Ta. While the growth at 700, shows two regions with different structure, on top there is more ordered α Ta resembling structures and transitions to more granular structure. This points to the spatial variation in the nature of growth induced control over the Ta phase.

However, we bought one wafer with tantalum grown on sapphire substrate from Starcryo [19], which was grown with high temperature growth conditions. In fig. 2.6(e), large domains with needle like grains can be seen uniformly thus to achieve similar results for the in-house grown wafers, further optimization should be done to find good growth conditions.

TANTALUM DEPOSITION WITH NIOBIUM SEED-LAYER

An alternative to the elevated temperature induced appropriate growth conditions, is to use a niobium seed layer. This layer closely resembles the α Ta growth and thus helps with inducing the α Ta without the need for heating. This is particularly useful in the case when silicon is used as a substrate, as heating silicon with Tantalum deposition leads to formation of Tantalum silicides [20, 21].

Figen Y. optimized the tantalum deposition recipe with the Nb seed layer. 10 nm of niobium is deposited first, followed by 200 nm Ta. In fig. 2.6(c) and (d), seed layer deposition on sapphire and silicon are shown respectively. The grains resemble needle like features and look more ordered, and match α Ta domains [22]. More thorough material characterization still needs to be performed to definitively distinguish phases in the growth.

Finally one experimental deposition of niobium with 200nm deposition and a thin capping of 10nm was done. The Tantalum growth for the thin layer looks like the niobium seed-layer growth on sapphire, but there are open voids visible in deposition as seen in fig. 2.6(f).

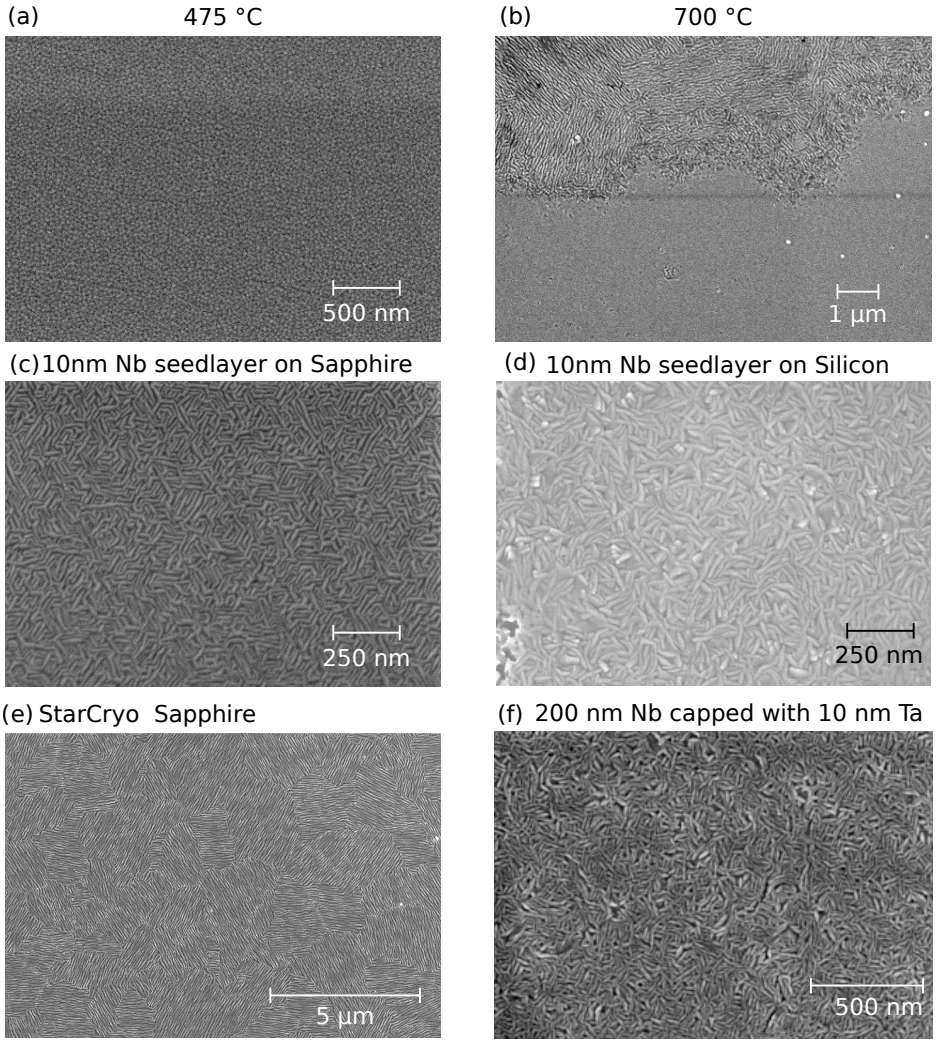

Figure 2.6: Scanning electron microscope images of Tantalum films: (a) Ta deposited on sapphire substrate heated to 475° (b) Ta deposited on sapphire substrate heated to 700° (c) Ta deposited on sapphire substrate with a niobium seed layer of 10 nm (d) Ta deposited on silicon substrate with a niobium seed layer of 10 nm (e) Ta deposited by StarCryo (f) Nb deposited on silicon substrate capped with a 10 nm tantalum layer

2.3.3. ALUMINUM DEPOSITION FOR JOSEPHSON JUNCTIONS:

Josephson junctions (Ch. 4) are fabricated using aluminum deposition in our qubit fabrication. We use electron beam evaporation (EBE) based PVD technique to deposit aluminum, performed using the Plassys system in the Kavli cleanroom.

The EBE works by accelerating electrons through an applied voltage and redirecting them onto the metal bucket using a magnet, see fig. 2.7(a), with the electrons emanating from the slit after the voltage source. Compared to thermal evaporation, where the whole material is melted and heat transferred throughout, in EBE the electrons impart with high kinetic energy onto the metal, leading to a local melt zone. The electrons collide and travel in the lattice structure for a few hundred nanometers before losing all their energy. The electron temperature of 1500° C material is roughly on the order of meV, so an electron accelerated to keV will be able to travel further and lead to significant local melt. The electrons are accelerated with a applied voltage of 100 kV.

The geometry of the sample mounted in the deposition chamber is described in the fig. 2.7 (a), with the arrows indicating the degree of freedom to enable a double angled deposition. The figure shows the top view, at the bottom the mounted sample looks as fig. 2.7 (b).

However evaporating Al with EBE is a tricky task, especially in vacuum. Although the melting temperature of Al is only 600° C, The vapor pressure required to attain the material deposition rate of 0.5nm/s leads to a temperature requirement above 1000° C. This can easily lead to more material being melted in the pocket. The melted pocket accumulates a lot of thermal energy and sloshes around in the bucket, referred to as spitting. As the material gets depleted inside the pocket, this high temp requirement further complicates the spitting problem. The hot spot of the vapor formation gets deeper into the pocket and the overall amount of material gets less, see fig. 2.7(c). If the beam remains centered, the thermal conduction path becomes smaller, thus requiring ever more current, necessitating the calibration of the beam to off-center, see fig. 2.7(d). The highest current with the deposition rate being fixed to 0.5nm/s changes over the pocket age, indicating the transition at some unstable equilibrium points.

In the current Al pocket configuration, no liner is used, instead aluminum is directly melted in the copper grooves. This reduces sloshing, reducing thermal buildup in the melted Al. Further liner can deteriorate the spitting, leading to uncontrollable deposition rate and crack due to thermal fatigue.

During the deposition, with the vaporization of aluminum progresses, the incoming electrons hit the vapor cloud and ionize the vapor, see fig. 2.8. This leads to charging of the vaporized material, and the amount of charge changes with change in the current of the electron beam. Further, radiation is also emit-

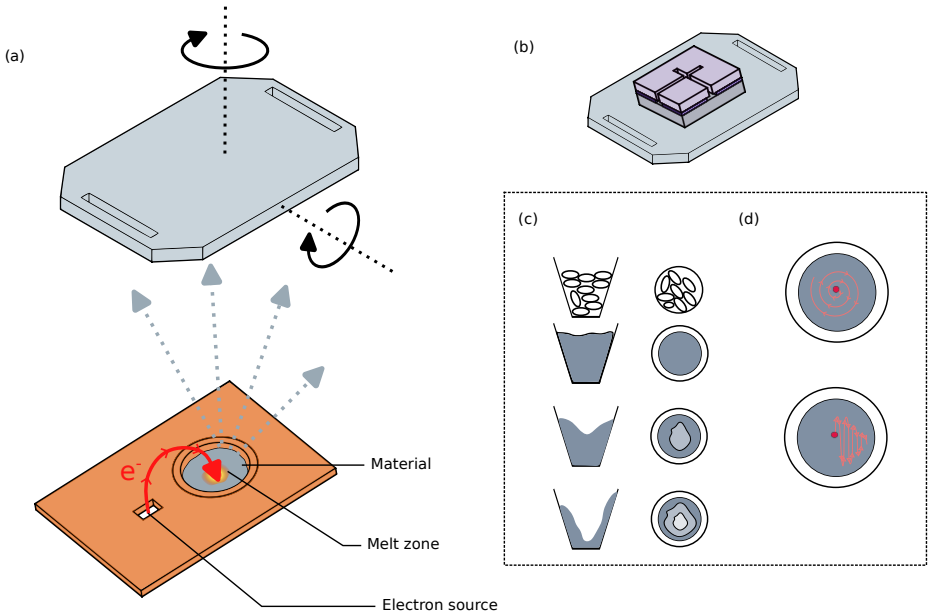

Figure 2.7: EVG: (a) Sample holder orientation shown on top of the e-beam material deposition source. The alignment is not to scale (b) inverted image of the sample holder with sample on top (c) Material depletion shown with progressive removal of material from the bucket, during different stages of deposition (d) Movement of the e-beam on top of the bucket

ted in the form of X-rays and thermal radiation As the material left in the bucket depletes, it has a combined effect of change in amount of ionized vapor, IR radiation and secondary electrons. In the deposition chamber, these factors influence the grains properties of the aluminum and it's properties.

2.4. LITHOGRAPHY

Lithography is the process of pattern generation used as the precursor for the creation of lateral structures on the chip. In tandem with additive and subtractive material processes, lithography enables creation of well-defined patterns at scales from nm to mm, which are generated using a combination of (1) a writing medium and (2) a writing tool. The dimensions of the structures defined on the chip, dictate the requirements on the writing parameters.

In the Circuit QED chip fabrication, patterns are defined with dimensions ranging from 100 nm to a few mm. For the writing tool, usage of photolithography and electron beam lithography (EBL) is generally used in various labs. We specifically use the electron beam lithography as the writing tool, as it enables a sub 5 nm writing resolution. For a single electron accelerated using voltage V, the

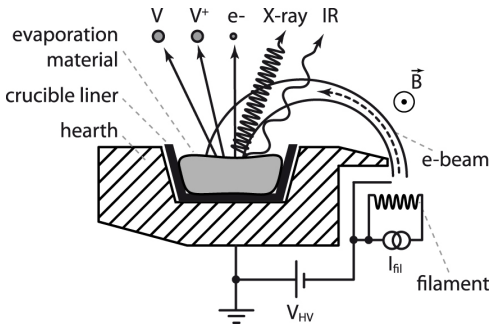

Figure 2.8: Molten material: The e-beam is shown on the material source, with different particle emissions and radiations coming off from the bucket. Figure reproduced from [23]

wavelength is-

$$\lambda = \frac{h}{\sqrt{2m_e K}} = \frac{12.27\text{A}^\circ}{\sqrt{K[eV]}} \tag{2.1}$$

where the enegy of elctron is in eV. With a 100 kV acceleration, a resolution of less then a nm can be achieved. In practice the combination of focusing mechanism of the electrons and writing medium combined give a resolution of 5 nm.

Along with the writing tool, the writing medium used is called as the resist, which is a radiation sensitive organic polymer chains mixed with a suitable solvent. For pattern generation this polymer is spun in a thin film, subsequently baked to evaporate solvent and exposed to electrons. Following pattern exposure, the resist-coated substrate undergoes a development process, wherein it is immersed in a developer chemical. This step selectively dissolves the exposed regions of the resist, thereby revealing the intended pattern on the substrate surface.

A brief understanding of the resists and the EBL machine is described first, followed with the process developed inside the cleanroom.

WORKING PRINCIPLE:

The EBL works by focusing a narrow beam of electrons onto a organic polymer resist. For our process we use a Raith Electron beam pattern generator (EBPG) 5200 system. In the 5200 machine, electrons are generated using a thermal field emission (TFE) gun , which are focused using a column of electromagnetic lenses and directed onto the writing sample. The electron source is the most important part of the system, TFE sources can generate electrons in a narrow tip of tungsten, enabling better focusing capability. The spread in the electron energies within the beam dictates the extent of chromatic aberration suffered by the beam.

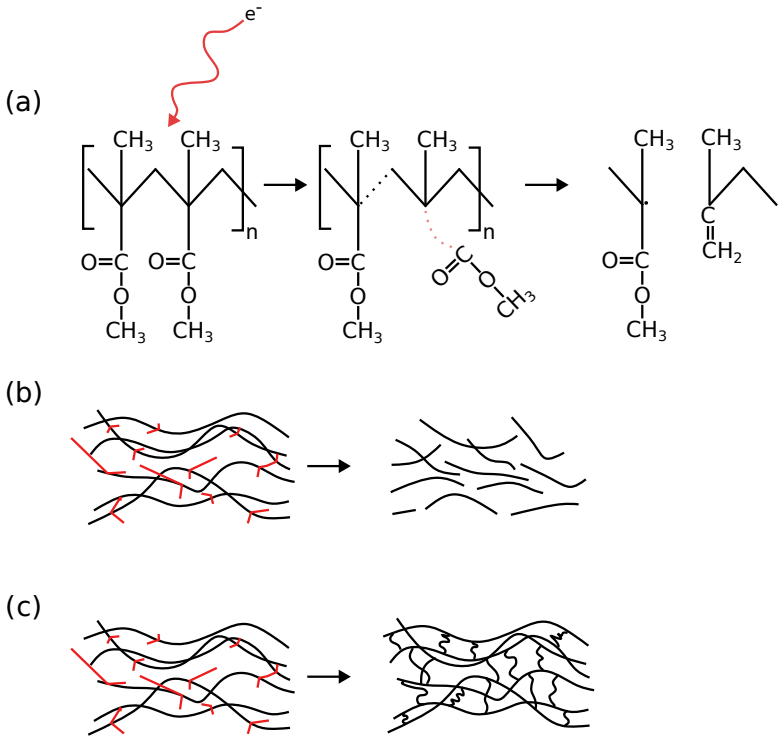

Figure 2.9: Polymer resists: (a)PMMA polymer chain unit is shown with an electron incident on it. After the electron exposure the chain breaks into smaller units. (b) Positive resist impacted by the electrons in red leads to chain scission events. (c) Negative resist cross linking process is shown after impact with electrons.

Stability of the source/filament overtime is necessary to maintain a good reproducibility of the system. extensive use of the source will eventually lead to degradation in focusing parameters, including the width of the beam emanating forming on the tip and the spread in the electron energies. The filament must be reset on a regular basis to achieve stable source characteristics. There are still peculiar effects which can occur even when the filament is new. Bigger exposures with a requirement of higher current emission from the source, can lead to more aberrations and faster degradation due to thermal instabilities. It can lead to drift in the beam spot sizes. This issue was encountered in the new Raith system bought in 2024 Q3. One must take into account if the patterns drift overtime.

When electrons strike the resist, a combination of physical and chemical processes occurs to form the desired pattern. The resist is composed of carbon based polymeric chains, when the electrons at high energy hit the atoms in the chain, some of the electrons suffer inelastic collision with the atom and transfer energy to the atom, these electrons are called the forward scattering electrons. This energy will lead to ionization and break the bond of the polymer and lead to broken fragments, see fig. 2.9(a). Simultaneously, higher energy electrons will travel further through the resist and can hit the nucleus of the atoms in the substrate. This can lead to the electron elastically colliding with the atom and taking back almost most of the energy. These are called as the back-scattered electrons and this takes place at the substrate resist interface, see Fig. 2.10(b). When the primary polymer chains break into smaller fragments as a result of the bond breaking ionization, the smaller fragments have a higher solubility in the developer solution, compared to the longer chain. This difference in solubility enables a exposure enabled pattern generation. This typ eof resist is called a positive resist, see fig. 2.10(b).

Alternatively, the energy of electrons can also be used to chemically alter the polymer is revrse manner, to increase the molecular weight of the polmyer upon leectron exposure. This process forms bonds in the polymer, cross-linking it and decreases resist solubility in the expodsed region. This type of resist is called as negative resist, see fig. 2.10(c).

PATTERNING PROCESS:

The EBL pattering is done in four steps: (1) Substrate prep (2) Resist spinning (3) Patterning (4) Resist development. The process starts with the Here we discuss them further and understand common failure modes.

(1) SUBSTRATE PREPARATION

The objective of this step is to eliminate any particulate matter acquired from different processing steps performed after metal deposition. To eliminate the dirt

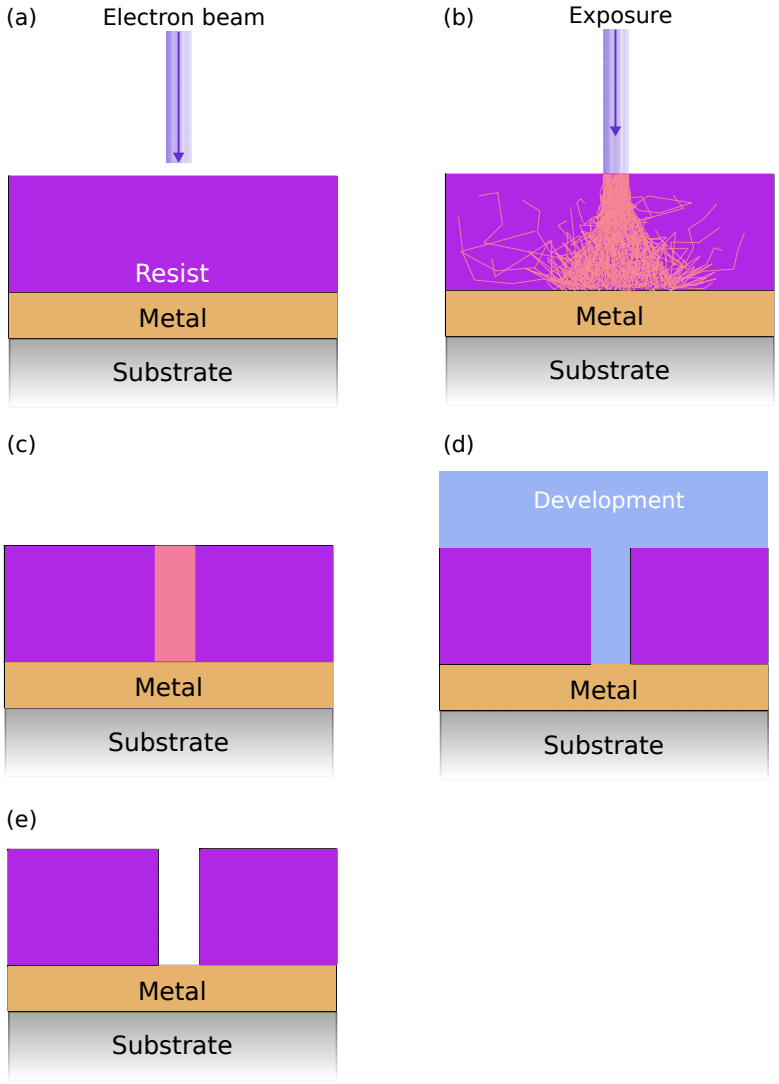

Figure 2.10: Process of pattern generation: Steps of E-beam lithography is shown from (a) to (e). The first step is to expose the resist with electrons, after the exposure the resist breaks into smaller fragments, which can be dissolved using a developer solution, leading to a exposed pattern.

and chemicals on the surface cleaning must be done. Additionally, the surface properties must be suitable for the resist to stick completely to the surface.

After metal deposition and dicing, the substrate is left covered in dicing resist in sample box. This resist layer acts as a barrier to decrease any dirt exposure to metal surface and reduces oxidation on the metal surface [24]. This resist is removed by dissolving in a resist stripper, either NMP or PRS-3000. In the same step initially acetone was used but should be avoided as a stripper, due to it's low boiling point and high vapor pressure. Dicing resist dissolves rather quickly, within 15 mins, in the NMP at 80 degrees. We rinse the NMP away by subsequently dipping the coupon into Acetone for 1 min, followed by IPA for 1 min and blow dry until no droplets are visible. At this point, the user must check the coupon optically with naked eye and subsequently with the optical microscope. Checking in both bright and dark field allows the user to discern any apparent particulate matter on the surface.

Once cleaned from leftover organic residues, the surface is physically and chemically altered to improve the adhesion of the resist on the surface. First a quick 2 min oxygen plasma descum is done to change the surface energy and further cleaning it from organic residues. When the process of plasma cleaning is used, it heavily oxidizes the surface of the coupon. This change in chemical composition enabling highly wettable/hydrophilic surfaces, which adsorbs a certain degree of water molecules on the surface. Adsorption of water molecules on the surface of the coupon will lead to the -OH bonding to the surface. We heat the coupon for 10 minutes at 150 ° to vaporize any residue amounts of the water on surface and follow it up with priming the surface with a HMDS monolayer.

The initial objective of introducing the process of plasma cleaning and HMDS priming was to dehydrate the surface and improve the resist adhesion, it has not found to be significantly different when these processes were not used. In favor of reducing complexity and steps, I would recommend to remove these steps and just do 5 min baking before spinning.

(2) RESIST SPINNING

The resist is the writing medium and thus needs to be evenly spread out on top of the substrate. This is done by spinning the resist in the solvent, followed by baking to evaporate the solvent. The resists are stored in glass beakers by the cleanroom staff, as the resist ages, the solvent evaporates and forms solid residues inside the bottle. When taken out of the bottle, the resist is filled in a syringe with a filter to eliminate the solid residues. Careful storage and filtering must be done to avoid the solid residues from the resist bottle. This can be done by replacing resist in-time and accessing the resist from the middle of the beaker to eliminate denser stratified resist mixing.

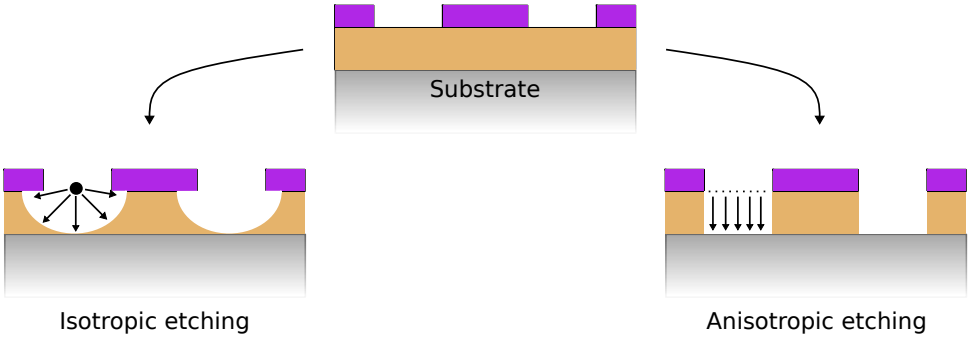

Figure 2.11: **Etching profile**: On top a substrate is shown with the patterned resist, the arrows follow two different types of etch profiles. The Isotropic etch profile is seen when the ecthing happens with the same rate in all directions. While the Anisotropic etchins ahppens when the etching rate is different in different directions, speciofically in this figure it is faster in the vertical direction.

The process of filling the syringe can also lead to bubble formation in the resist and eventually end up on the substrate. These bubbles are undesired as they leave an uneven resist surface on the substrate. Glass pipettes can be used to eliminate the bubbles. Once the resists has been dispersed on the substrate, selection of the spinning speed is done based on the desired thickness. Every resist has a spin-speed curve, which can be accessed from the datasheet and should be validated in the cleanroom for more accurate thickness. The initial spin speed and the exact solvent can influence the spin-speed curve.

2.5. ETCHING

Etching is the process of removing material from the metal and substrate layers, transferring the resist pattern into the lateral features on the circuit. The material can be removed by choosing either methods or a combination: (1) using a chemical process generally referred to as wet etching, or (2) using a physical removal process referred to as dry etching. The choice of etching and it's parameters influence the shape of the etched material, characterized with the etch profile. Two types of etch profile can be seen- an isotropic profile or an anisotropic profile, made possible in both wet and dry etching, see Fig. 2.11.

In the following sections, we look at a detailed description of the mechanism of the etch processes and their underlying parameters.

2.5.1. WET ETCHING:

Wet etching is done by using a liquid chemical solution, that selectively removes the metal film, while not attacking the resist and the underlying silicon substrate.

In our fabrication process we use the wet etching process to define the Silicon nitride based dielectric patch underneath the nanowire and experimented with wet etching of NbTiN metal layer.

WET ETCHING FOR DIELECTRIC:

The SiN dielectric is deposited on the patterned base layer with underlying parts partially on NbTiN metal and silicon. The etch solution chemical must selectively remove the dielectric, simultaneously not etching the underlying metal and the silicon. Additionally the resist must maintain it's shape during the process. In the process, the choice of etch solution is 20:1 BOE which acts by providing a constant pH buffered solution with ammonium fluoride and HF ions:

$$Si_3N_4 + 12\ HF \rightarrow 3\ SiF_4\ (gas) + 4\ NH_3\ (gas) \tag{2.2}$$

followed by

$$SiF_4\ +\ 2\ H_2O \rightarrow H_2SiF_6\ +\ 2\ HF \tag{2.3}$$

Fluorine ions (F^-) act as the complexing agent in this reaction, bonding with the silicon forming the gaseous by-product (SF_4). The reaction proceeds slowly due to a strong Si-N bond in the nitride, leading to a surface reaction limited rate. The chemical solution is used for etching, the gaseous nature of the products allow for faster dissolution and transport in the solution[[25, 26]].

When etching the dielectric on the chip, the area and height of the actual patch after etching is smaller compared to the lithographically defined pattern. This is due to two effects: (1) The etch solution leads to an isotropic etch profile, see Fig. 2.12 (a), with the effective area of the patch being undercut from all sides. (2) The Chemical solution seeps underneath the resist due to the capillary action, reducing the effective height of the patch, shown in fig. 2.12 (b). To mitigate the effect of the undercuts, adhesion promoter is used before spinning the negative resist and subsequently a bigger area pattern is defined. Additionally, special care must be taken to maintain the pH of the BOE solution. Lower pH can lead to a seeping of the HF through the resist and lift-off of the resist pattern, thus a 20:1 solution was chosen. Further, inconsistent etching is observed, as seen in fig. 2.12, across larger areas, potentially due to the inconsistent flow characteristics in the solution while performing the etching.

2.5.2. DRY ETCHING:

To understand the working principle of the dry etching process, we draw a comparison to the PECVD process, focusing on the chemical reaction to do etching rather then deposition. In dry etching, reactive gases interact with the exposed regions of the substrate surface forming volatile by-products that are subsequently removed by the gas flow within the chamber. While these chemical

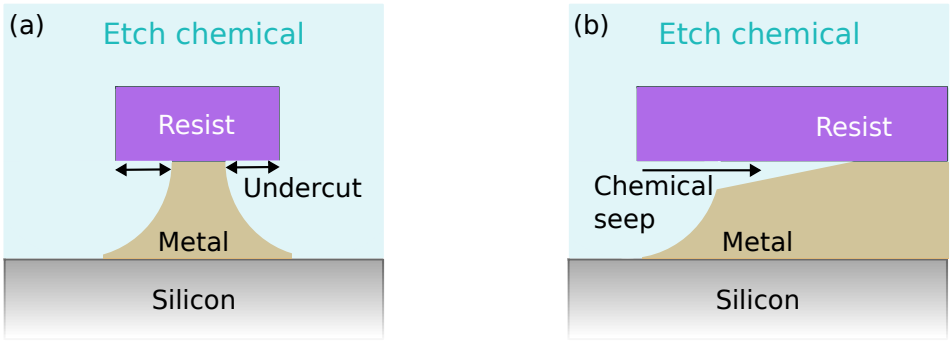

Figure 2.12: Mechanism for inconsistent wet etch (a) Undercut profile development due to isotropic etching at the edge of the resist is shown. (b) Peeling of resist form one side with chemical seep

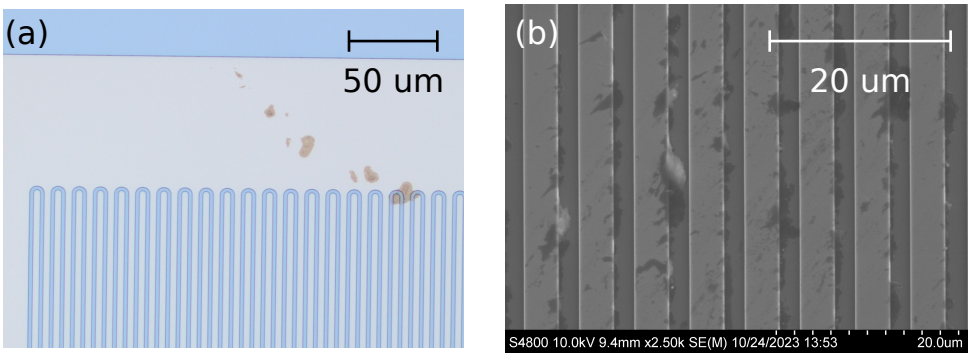

Figure 2.13: Wet etch residues: (a),(b) residues from the wet etch of SiN

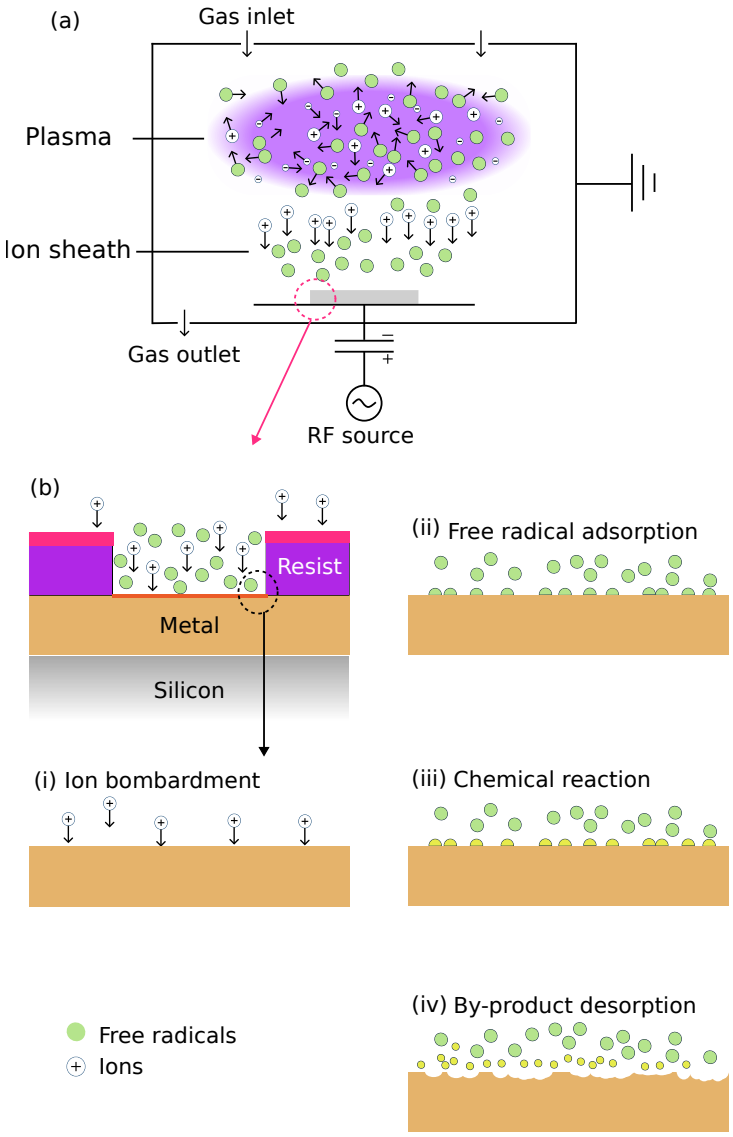

Figure 2.14: Principles of Reactive ion etching: (a) Geometry and working of the direct etch plasma system is shown with distribution of gaseseous species inside the chamber. (b) Exploded view of the sample being etched. (i)-(iv) zoomed in view of the RIE process in the exposed area of the resist in (b).

reactions require high activation energy, generally achievable at high temperatures, the same is achieved by using a plasma. The combination of plasma and gaseous species are responsible for the rate and nature of the chemical reaction deciding the etching effects.

WORKING OF THE PLASMA RIE ETCHING:

In Fig. 2.14(a) we explain the detailed working of the Reactive ion etching. We use a RIE plasma etcher system, the Sentech Ethclab 200, locally refered to as F1 system. This system generates capacitively coupled plasma with the samples directly loaded in the reactor. The system uses a laser in-point detection system for in-situ measurements.

The process starts with loading the sample inside the chamber and pumping it to low pressures of 10^{-5} mBar pressures. Subsequently, a gas mixture of selected chemistry is introduced from the holes in the top plate of the chamber. With a constant pump and gas inlet, sufficient pressure is allowed to buildup for the plasma ignition.

Next, the system RF source is switched on to apply oscillating electric fields with a frequency of 13.56 MHz. This frequency provides the energy to start a glow discharge in the chamber, generating a mix of positive ions (White) and neutral radicals (Green), see Fig. 2.14(a), with the radicals being higher in number compared to the positive ions.

As the charges build-up, the electrons go farther from the cathode due to coulomb repulsion, this effect leads to a buildup of positive ions closer to the cathode, termed as ion sheath. Within this region, all the positive ions accelerate towards the cathode and hit the cathode and the sample. When the ions hit the cathode, it also leads to emission of secondary electrons from the cathode surface, leading to continuous maintenance of plasma.

At the same time, free radicals are generated that are highly reactive neutral species which chemically react with the underlying film and the substrate to remove material as a volatile by-product. The role of ions in the RIE process, is to impart energy to the material being etched and break bonds on the surface, this lowers the energy required for chemical reaction between the radicals and the thin film. The etching reaction goes in this manner: (1) In Fig. 2.14 (i) the ions impart the energy (2) Free radicals are adsorbed on the surface of the sample (3) Chemical reaction takes place, forming the by-product (4) The by products are de-adsorbed from the surface. Each step has a rate and the whole reaction is limited by the slowest step among these four. The ions play an additional role of catalyst for each of the four steps and will influence the rates. Further, the ions travel towards the cathode and the sample surface impacting only the horizontal surface due to the electric field direction, generating the anisotropic etching

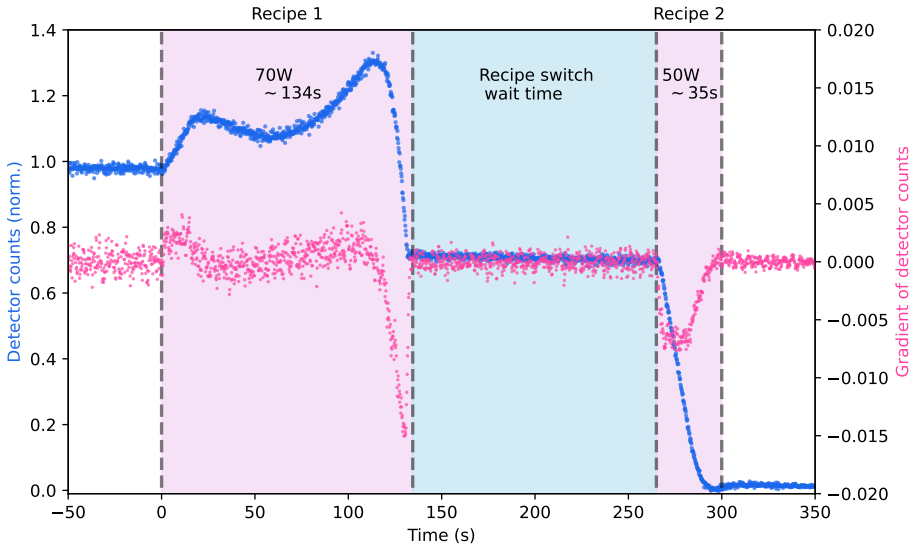

Figure 2.15: End point detection: plot of end point detector values vs the time. The time progresses form before start of etching and goes on till after etching is finished.

profile of the film.

In the etching process, the role of chamber pressure and DC bias decides the mean free path of the ions which bombard the substrate. Sdditionally, the density of the free radicals is also influenced by the pressure. For a fixed RF power, as the pressure is increased, the free radicals increase in number, increasing the chemical etch rate, but the energy of ions incident on the substrate decreases due to more shortened mean free path. Similarly, for a fixed pressure, increasing the RF power leads to more DC bias and more ion flux at the substrate, making the etching more physical. The optimization of power and pressure must be done to achieve the desired etching.

Alongside the etching of the film, the user must keep notice of the etch rate of the resist film, see fig. 2.14(b). CSAR resist has a particularly high etch resistance to a wide variety of chemistries, but it will still get removed if a oxygen rich mix is used for a long duration. This is portrayed by the etch on top of the resist.

IN-SITU END POINT DETECTION:

To characterize the thickness of the film that has been etched in the process, the F1 system uses a laser interference based end point detection scheme. The scheme works by pointing a laser towards the chip, specifically on the developed area that is open to plasma. Ideally this area is chosen to be big enough to place

the laser spot, carefully with help of an onboard camera.

When the laser hits the surface, it gets reflected back to the laser detector, creating an interference in the detector. As the metal layer gets etched slowly, the interference changes as a result of change in the length of the path of the light. Further as the etching process reaches the transition from metal to substrate, the laser reflection changes significantly, leading to a significant drop in the detector count.

In fig. 2.15, the data from a etching run is plotted, where the blue curve is the collection of reflection values in the detector and the pink points are the first derivative of this curve. The process is divided into two sub-processes with different etching power, labeled as recipe 1 with 70 watt RF power and recipe 2 with 50 watt RF power. As the process starts with recipe 1 at 0 sec, we see a shift in the detector values, this happens due to change in the reflection happening from etching of the metal. The etching transitions from the metal to the substrate, observed by the global value drop in the derivative at around 130 s. As the derivative starts dropping, we stop the recipe and start the process to recipe 2. At the transition, the values of detector remain stable, shown by the blue shaded region in the plot. next, as the system starts etching with the recipe 2, a drop in detector is measured. The gradient of the detector counts as a measure to see when the recipe has reached a stable point in silicon, as the gradient stabilizes after recipe 2, we stop the recipe and start pumping any residue gases in the chamber.

CHEMISTRY OF THE ETCHING COMBINATION:

The gas mixture introduced in the chamber is decided based on the chemistry of the material that is being etched. The combination of gases has two components, one is halogen based gas like SF_6, CF_4, Cl_2 CHF_3, BCl_3 and other is oxygen (O_2). Halogen based gases are highly reactive, producing reactive radicals upon exposure to plasma and the by products formed are generally volatile. Further, these compounds have different selectivity for the metals and silicon. Simultaneously, introduction of O_2 in the reaction chamber, acts as a carrier to remove the volatile by-products and fluoro-polymers, which are constantly being formed in the etching reactions. The fluorpolymers stick to different surfaces and deteriorate the etching of the surface, thus removal of fluorpolymers is an important part to carry forward the chemical reaction for etching. A balance between halogen and oxygen needs to be optimized to create the correct etch profile, etching rate and surface morphology.

For the etching optimization of the baselayer of our chip, we used different combinations of gases for the metal and substrate etching, listed as follows:

NBTIN

NbTiN is a hard refractory metal alloy, making it harder to etch. We used a SF_6 and O_2 mixture to etch the thin film on the silicon substrate.

In the presence of plasma, SF_6 disassociates down into a variety of positive and negative ions and forms free radicals. Additionally, the SF_x species also break down further to produce more fluoride ions as-

$$e + \mathrm{SF}_x \rightarrow \mathrm{SF}_{x-1} + \mathrm{F} + e \tag{2.4}$$

This cascade of reactions, reacts with the metals forming fluroide compounds -

$$\mathrm{Nb} + 5\mathrm{F}\cdot \rightarrow \mathrm{NbF}_5 \quad \text{(low volatility)}$$
$$\mathrm{Ti} + 4\mathrm{F}\cdot \rightarrow \mathrm{TiF}_4 \quad \text{(higher volatility)}$$

In these reactions, the products formed have different volatility and the removal of the low volatile compounds are accelerated by the help of ion impact. This leads to a complex balance of physical and chemical etching.

The role of oxygen in this etching reaction is to improve the volatility and carry the by-products away from the substrate. Oxygen does this by reacting with the fluorides formed as -

$$\mathrm{NbF}_5 + \mathrm{O} \rightarrow \mathrm{NbOF}_3 + \mathrm{F}_2$$

The resultant oxy-fluroide is more volatile and again makes the etching more chemical.

Optimum surface morphology: Two - step etching recipes:

To produce a smooth silicon surface at the end of the etching procedure, we optimized the process into two steps. The first step achieves a higher etching rate into NbTiN with higher selectivity towards NbTiN, but produces a rougher surface. The RF power used is higher and the ratio of SF6 is also higher. We terminate the first step as soon as the etching starts transitioning to the silicon, characterized using the in-situ laser interferometer.

In second step we use a lower power higher oxygen content recipe, producing a more chemical etching and leading to smoother surface. This step is stopped as the etching rate reaches a plateau of the gradient of the detector count. A heavily physical etching leads to more roughness from the ion bombardment, as seen in 2.16(a), this will happen if the recipe was continued all the way till it etches silicon. When the etch time for recipe 1 and recipe 2 are calibrated correctly for the NbTiN thin film thickness, it corresponds to a surface morphology like 2.16(b). Further, the dose of the features shall also be optimized along with the RIE recipe, to get the correct dimensions of the structures on the chip, as compared between 2.16(c) and (d).

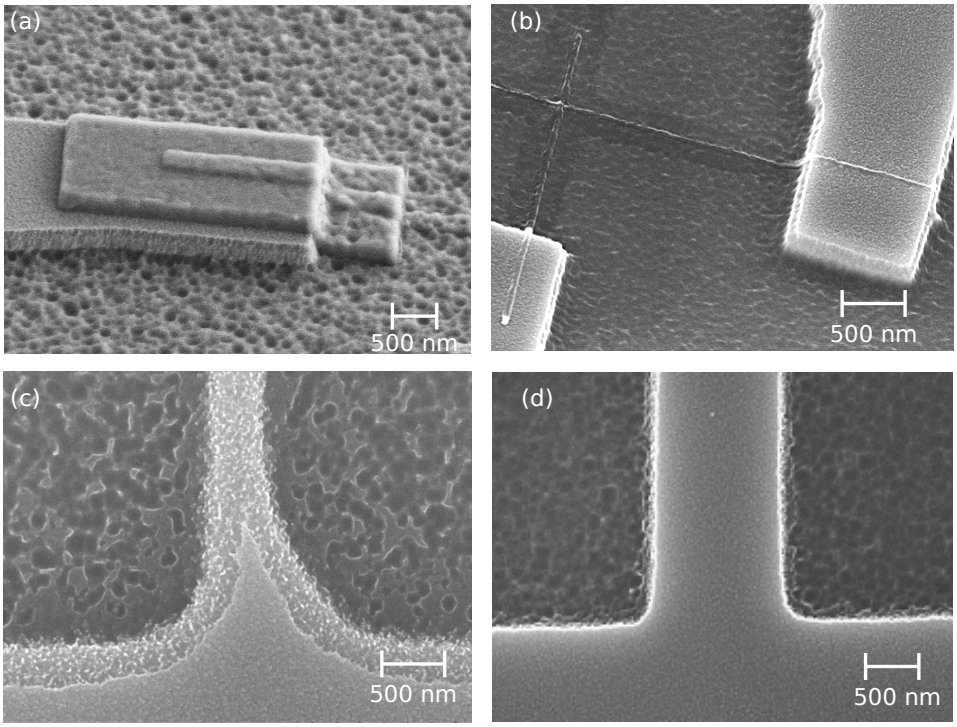

Figure 2.16: Etching optimization: (a) SEM of a chip surface with RIE etching performed with more physical etching recipe (b)SEM of a chip surface with RIE etching performed with more chemical etching recipe (c),(d) comparison of metal feature on the chip with non-optimized vs. optimized etch flow respectively.

Name	parameters	Gas (sccm)
NbTiN		
Recipe 1		
	0.01 mBar ; 70 W	SF_6:O_2 (13.5 : 4)
Recipe 2		
	0.08mBar ; 50 W	SF_6:O_2 (4 : 15)
Tantalum		
CF_4 recipe		
	0.01mBar ; 50 W	CF_4 (20)

TANTALUM:

Tantalum is a refractory metal, which is harder to etch due to chemical stability. We experimented with the SF_6 and O_2 mixture for etching, but formation of grass like spikes on the substrate were observed, as seen in 2.17(a) and (b). Exact mechanism of formation and the nature of these spikes was not investigated. Potentially these can be a result of low volatile compounds being redeposited on the substrate or residues of polymers formed on the surface.

In 2.17 (c), we switched to using just CF_4 and found very good etching results at a RF power of 50 W. The reaction of etching of tantalum proceeds by formation of Tantalum pentafluoride as

$$Ta + 5F\cdot \rightarrow TaF_5 \tag{2.5}$$

The better etching of Ta with the CF_4 can be potentially explained by more physical nature of the etching, leading to more effective removal of the byproducts. Additionally, the substrate used in this case is sapphire, which is very hard to etch, thus it can sustain ionic impact to produce a smooth surface.

Similar etching with CF_4 is harder to optimize for the silicon substrate and for the samples where Ta was etched on top of silicon, we switched to usage of the SF_6 and O_2 mixture, but albeit with different parameters for recipe 1 compared to NbTiN.

2.6. ETCH RELATED ARTIFACTS AND OPTIMIZATION OF BASE LAYER FEATURES ON THE CHIP:

Post etching inspection in our fabrication process has routinely shown artifacts, which are suboptimal for feature designs on the chip. In fig. 2.18(a), one of the common artifacts is seen as the presence of small islands of metal where they should have been completely etched away. This has been observed although also with varying severity across different metals and processes. In fig. 2.18(b), SEM of one small island left over on top of silicon is shown. The metal islands are

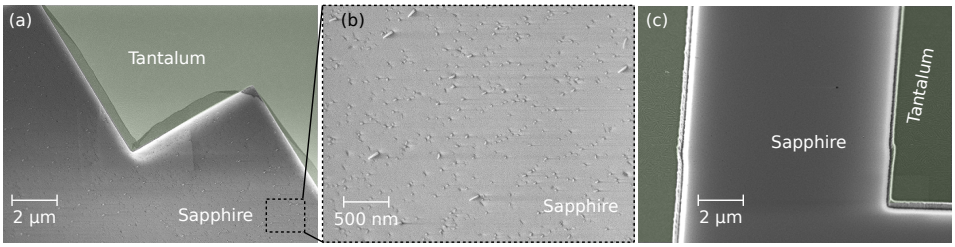

Figure 2.17: Tantalum Etching: (a) Tantalum surface SEM after etching with two step recipe as NbTiN. Tantalum surface is false colored in green. (b)Zoomed in surface SEM of the sapphire surface (c) Surface SEM after CF_4 based etch recipe of the Tantalum.

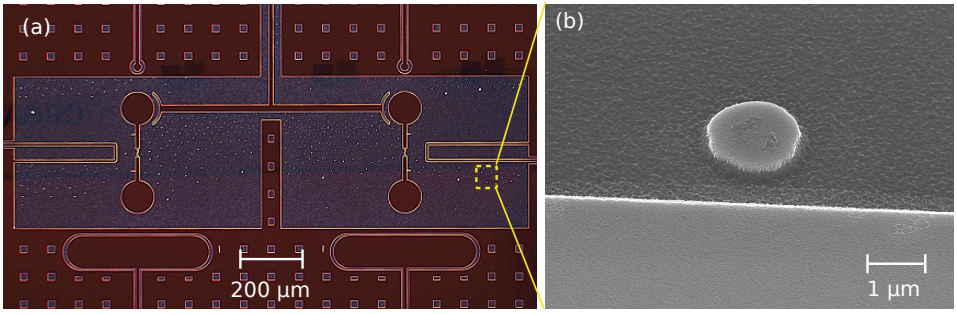

Figure 2.18: Metal etch residues (a) Dark filed micrograph of the features on an etched chip. The bright islands are metal residues in the silicon. (b) Zoomed in SEM image of the metal island

generally visible in larger etching pockets and smaller features are relatively more cleaner.

From our current observations and optimizations, the presence of these artifacts is attributed to two main reasons - (1) Resist residues during etching, (2) Loading effect from the plasma and re-deposition of residues during the RIE process. These two factors influence locally how the ions are removing the material from the substrate. At the same time, the volatility and the carbon-fluoropolymer formation in the reaction chamber can be leading to micromasking. Further reposition of the etched material onto the substrate can be one potential explanation as well. To mitigate these effects, the development process was optimized and steps of ashing before and after the RIE etching have been included in the fabrication flow of the base layer. Further changes in the resist cleaning of the post-etch residues were added. Resist stripper PRS-3000 has replaced the usage of NMP, due to the improved surfactant component of PRS. The optimized recipe from metal deposition to just before the surface treatment of the baselayer is -

Table 2.6: Base layer latest fabrication steps

Step	Parameters	Duration
Development	Amyl Acetate (60 s); MIBK (30 s); IPA (60 s)	
Oxygen ashing	20 sccm O_2 flow at 200 W	2 minutes
RIE	recipe 1 and 2	
Oxygen ashing	20 sccm O_2 flow at 200 W	3 minutes
Resist stripping	PRS-3000 (80° C)	2 hours

2.7. CONCLUSION & OUTLOOK:

2.7.1. DEPOSITION:

1. **NbTiN** The process of deposition for the superconducting thin film has been used NbTiN as the base layer metal. The recipe 2.3.2 was optimized previously in the lab and used to deposit 200 nm thick films. Resonators and qubit chips have been extensively fabricated using these thin film dispositions.

   During the different wafer NbTiN depositions, the compressive stress was not monitored on a throughout the three years, this can be improved and frequent monitoring can lead to avoidance of highly stressed thin films.

2. **Tantalum** Deposition recipes were experimented with in-house sputter system to test the possibility of generating α phase of the tantalum. For Ta deposition on sapphire using heated deposition conditions and a Nb seed layer were successfully executed. In these recipes, features resembling the α-Ta were observed using SEM based imaging and resonators and qubits were fabricated using these wafers. Additionally, Annealing recipe of the sapphire was tested and executed before the deposition of the Tantalum.

   Further, extensive characterization with XRD based measurements will validate the presence of the correct α-Ta phase and help narrow down the recipe parameters. The annealing recipe further needs to be performed in a timely manner before the deposition for reproducing the exact substrate conditions before each deposition.

3. **Pre-processing** Add solvent based organic cleaning as the first step to remove dissolvable organic residues, this can reduce the amount of bigger easily removable particles to carry over in the organic step cleaning.

4. **Wet organic clean:** RCA -1 composes of a mixture of NH4OH, H_2O_2 and H_2O in a ratio of 1:1:5. The net pH of this prepared solution is between 10-11, leading to deprotonation of surface hydroxyl groups, creating the per-

fect conditions for particles lift-off. H_2O_2 is a very strong oxidizing agent to breakdown organics. Further, the NH4OH creates a surfactant like effect. This improves the solubility of the broken chunks into the solvent better and assists with the lifting up of particles. Additionally, breaking down of the peroxide into oxygen , generates micro bubbles which assist with agitation of the solution and dislodging the particles from the surface.

5. **Dry organic clean:** Alternatively, since the wafers are cleaned with the RCA clean by manufacturers before dispatch, we can use plasma based organic cleaning . This reduces chemical preparation for already clean substrates to clean the surface of the organics and follow it up with oxide removal step directly.

6. **Oxide etching:** The HF solution should be replaced with a BOE (7:1) solution for three reasons: (1) To remove any potential surface roughness introduced as a result of the reaction condition have lower pH with 40%HF (2) A buffered solution leads to a more consistent etching with the pH of the solution being stable in a buffered solution, and (3) Oxide etch rate depends on the pH of the BOE solution [27], thus choosing 7:1 solution for a period of 5 mins is recommended to etch the native silicon oxide, alternatively QDR bath can be explored to maintain rinse consistency. The drying can be done using the spin dry, eliminating handling based residue droplets drying on the surface. Atomically flat surfaces with Silicon 100 can be achieved by using higher pH BOE solutions allowing to build a uniform crystalline interface for the next step of metal deposition. Alternatively, suggestions to use ammonium fluoride solutions after the HF oxide clean have been demonstrated to reduce the TLS populations [28]. Additionally, sonication should be used to remove any bubbles formed during the oxide etch to create more smooth surfaces. [29]

7. **HMDS** The desorption and decomposition of the tri-methylsilyl group attached to the wafer surface leaves a monolayer between the metal deposition and silicon. This can act as a source of losses. Since the time after the water rinse followed by drying is done on wafer and transferred over to the HMDS station, it would take similar time to load the wafer into the sputtering chamber. Considering oxide formation is slower then the loading time, it is recommended to quickly and directly load the wafer into the sputtering chmaber, reducing one process step.

2.7.2. PATTERNING

Recipes with both the CSAR and PMMA, MMA based resist were developed for the whole fabrication flow of the qubits. Dose tests gave the optimum feature size and the desired outcomes in multiple cases. Problems with base layer features were observed during different phases of the project, potentially due to aging related resist hardening. Further, issues with metal islands were observed due to resist residues and plasma characteristics during the RIE process. This was troubleshoot with optimization of better development recipe, including sonnication instead of manual stirring. Changes were made to improve the adhesion of the resist by priming with HMDS were done. Additionally towards the same goal, plasms ashing step was added to change the surface energy of the surface.

For future changes, the resists can be filled in dedicated bottles to keep track of any resist hardening based residues. Further pipette mechanism with a micro pipette should be incorporated. Better mechanisms to mitigate presence of residues and bubbles will go a long way in making the fabrication easier and less tedious for the fab user.

Optimizing the development flow of the resists can be a good avenue to invest more time into understanding how it affects the optimum feature size. Further, eliminating the additional steps of priming and plasma ashing can be done. For this first the effect of such adhesion promotion steps needs to be seen, do we actually need these steps or whether the metal layer is not as detrimental to the adhesion of resist. Although, such steps are needed for negative tone resists, because of the resists chemistry but not equally applicable to all resists and metals. Thus elimination of these steps will save time and simplify the process of writing patterns for the users.

2.7.3. ETCHING:

- **Wet etching**: Etching of SiN was performed with BOE solution, this was done to fabricate gate dielectric structure to make gatemons. The residues of the SiN were observed and made the etching process complex.

  Wet etching of NbTiN was explore with RCA solution, but due to time limitations, the recipe could not be optimized to actually use it for any chips. In thee future, NBtIn wet etching recipe can be explored, the need for eliminating the two step recipe arises out of the system down time considerations.

- **Dry etching**: Recipes for both the Tantalum and NbTiN base metal layers have been optimized. For NbTiN, a two step etching recipe with SF_6 and O_2 mix was developed, which produces smooth etched surfaces. For the tantalum on sapphire a CF_4 based etch recipe was performed to etch the

metal layer.

For future study, the fluctuations in the surface morphology along with RIE led residues can be optimized more in depth. The fluctuations arise from the significant change in the system parameters due to continuous use of the system. This degrades the chamber walls by deposition of metal based residues and polymer based residues. The same residues also degrade the vacuum performance of the system.

Further, the calibration of the recipe parameters with each wafer run can improve the tracking of system performance over extended periods of time.

REFERENCES

[1] M. Checchin et al. "Measurement of the Low-Temperature Loss Tangent of High-Resistivity Silicon Using a High-Q Superconducting Resonator". In: *Phys. Rev. Appl.* 18 (3 Sept. 2022), p. 034013. DOI: 10.1103/PhysRevApplied.18.034013. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.18.034013.

[2] Alexander P. Read et al. "Precision Measurement of the Microwave Dielectric Loss of Sapphire in the Quantum Regime with Parts-per-Billion Sensitivity". In: *Phys. Rev. Appl.* 19 (3 Mar. 2023), p. 034064. DOI: 10.1103/PhysRevApplied.19.034064. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.19.034064.

[3] Alexander P. Read. "The dielectric dipper: experimental comparisons of common dielectric substrates in isolation". In: *APS March Meeting 2024*. Session K47: Superconducting Loss Characterization, March 5, 2024, Room 200CD. American Physical Society. 2024. URL: https://meetings.aps.org/Meeting/MAR24/Session/K47.1.

[4] W. A. Page et al. "A stress-induced source of phonon bursts and quasiparticle poisoning". In: *Proceedings of the National Academy of Sciences (PNAS)* 121.15 (2024), e2316588121. DOI: 10.1073/pnas.2316588121. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11292027/.

[5] Takeo Hattori. "Interface structure of Si/SiO2 systems". In: *Solid State Phenomena* 37 (1994), pp. 317–332.

[6] Zi-Huai Zhang et al. "Acceptor-Induced Bulk Dielectric Loss in Superconducting Circuits on Silicon". In: *Physical Review X* 14.4 (2024), p. 041022. DOI: 10.1103/PhysRevX.14.041022. URL: https://doi.org/10.1103/PhysRevX.14.041022.

[7] Loren D. Alegria et al. “Two-level systems in nucleated and non-nucleated epitaxial -tantalum films”. In: *Applied Physics Letters* 123.6 (Aug. 2023), p. 062601. DOI: 10.1063/5.0157654.

[8] Werner Kern and D. A. Puotinen. “Cleaning Solutions Based on Hydrogen Peroxide for Use in Silicon Semiconductor Technology”. In: *RCA Review* 31 (1970). Original publication introducing the RCA Standard Clean (SC-1 and SC-2), pp. 187–206.

[9] Werner Kern. *Handbook of Semiconductor Wafer Cleaning Technology*. Chapter 1 covers RCA Standard Clean SC-1 and SC-2 in detail. Park Ridge, NJ: Noyes Publications, 1993.

[10] Masaru Itano et al. “Particle removal from silicon wafer surface in wet cleaning process”. In: *IEEE Transactions on Semiconductor Manufacturing* 11.4 (1998), pp. 478–484.

[11] Sigma-Aldrich. *Ammonium Fluoride, NHF, ReagentPlus®, 98%*. 2024. URL: https://www.sigmaaldrich.com/NL/en/product/sial/30709m.

[12] P. T. J. Gennissen and P. J. French. “Sacrificial Oxide Etching Compatible with Aluminum Metallization”. In: *Proceedings of the 1997 International Conference on Solid-State Sensors and Actuators (Transducers ’97)*. Chicago, IL, USA, June 1997, pp. 225–228.

[13] A. Bruno et al. “Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates”. In: *Applied Physics Letters* 106.18 (2015), p. 182601. DOI: 10.1063/1.4919761. URL: https://doi.org/10.1063/1.4919761.

[14] Archana Kamal et al. “Improved superconducting qubit coherence with high-temperature substrate annealing”. In: *arXiv preprint* arXiv:1606.09262 (2016). https://arxiv.org/abs/1606.09262.

[15] Suhas Ganjam et al. “Surpassing millisecond coherence in on-chip superconducting quantum memories by optimizing materials and circuit design”. In: *Nature Communications* 15 (2024), p. 3687. DOI: 10.1038/s41467-024-47857-6. URL: https://www.nature.com/articles/s41467-024-47857-6.

[16] Julien Zichi. “NbTiN for Improved Superconducting Detectors”. MA thesis. KTH Royal Institute of Technology, 2019. URL: https://kth.diva-portal.org/smash/get/diva2:1316907/FULLTEXT01.pdf.

[17] Siddhesh Sanjay Yeram et al. “Structural transitions in superconducting NbTiN thin films”. In: *arXiv preprint arXiv:2311.14154* (2023). arXiv:2311.14154 [cond-mat.supr-con]. URL: https://arxiv.org/abs/2311.14154.

[18] J. A. Thornton. “High Rate Thick Film Growth”. In: *Annual Review of Materials Science* 7.1 (1977), pp. 239–260. DOI: 10.1146/annurev.ms.07.080177.001323. URL: https://doi.org/10.1146/annurev.ms.07.080177.001323.

[19] STAR Cryoelectronics. *Foundry Services*. Accessed: 2025-07-02. 2025. URL: https://starcryo.com/foundry-services/.

[20] R. Hübner et al. “Comparison of the annealing behavior of thin Ta films deposited onto Si and SiO substrates”. In: *Analytical and Bioanalytical Chemistry* 379.4 (2004), pp. 568–575. DOI: 10.1007/s00216-004-2602-5. URL: https://link.springer.com/article/10.1007/s00216-004-2602-5.

[21] Veronika Rinnerbauer et al. “Low emissivity high-temperature tantalum thin film coatings for silicon devices”. In: *Journal of Vacuum Science Technology A* 31.1 (2013), p. 011501. DOI: 10.1116/1.4766295. URL: https://pubs.aip.org/avs/jva/article/31/1/011501/246127/Low-emissivity-high-temperature-tantalum-thin-film.

[22] Loren D. Alegria et al. “Two-level systems in nucleated and non-nucleated epitaxial alpha-tantalum films”. In: *Applied Physics Letters* 123.6 (2023), p. 062601. DOI: 10.1063/5.0157654. URL: https://pubs.aip.org/aip/apl/article/123/6/062601/2905681/Two-level-systems-in-nucleated-and-non-nucleated.

[23] Frank Volmer et al. “How to solve problems in micro- and nanofabrication caused by the emission of electrons and charged metal atoms during e-beam evaporation”. In: *Journal of Physics D: Applied Physics* 54.22 (Mar. 2021), p. 225301. DOI: 10.1088/1361-6463/abe89b. URL: https://doi.org/10.1088/1361-6463/abe89b.

[24] Jeroen Van Damme et al. “Advanced CMOS manufacturing of superconducting qubits on 300 mm wafers”. In: *Nature* 634 (2024), pp. 74–79. DOI: 10.1038/s41586-024-07941-9. URL: https://www.nature.com/articles/s41586-024-07941-9.

[25] Stanley Wolf and Richard N. Tauber. *Silicon Processing for the VLSI Era, Volume 1: Process Technology*. Sunset Beach, California: Lattice Press, 1986.

[26] M. J. Loboda, A. R. Duggal, and S. W. Pang. “Wet Etching of Silicon Nitride in HF Solutions”. In: *Journal of The Electrochemical Society* 136.9 (1989), pp. 2782–2786.

[27] M. Virginia P. Altoé et al. “Localization and Mitigation of Loss in Niobium Superconducting Circuits”. In: *PRX Quantum* 3.2 (2022), p. 020312. DOI: 10.1103/PRXQuantum.3.020312. URL: https://doi.org/10.1103/PRXQuantum.3.020312.

[28] Cameron J. Kopas et al. “Enhanced Superconducting Qubit Performance Through Ammonium Fluoride Etch”. In: *Materials for Quantum Technology* 4.4 (2024), p. 045101. DOI: 10.1088/2633-4356/ad88cc. URL: https://doi.org/10.1088/2633-4356/ad88cc.

[29] Brandon S. Aldinger et al. “The same etchant produces both near-atomically flat and microfaceted Si(100) surfaces: The effects of gas evolution on etch morphology”. In: *Journal of Applied Physics* 107.10 (2010), p. 103520. DOI: 10.1063/1.3402580. URL: https://doi.org/10.1063/1.3402580.

3

RESONATORS

"If you want to find the secrets of the universe, think in terms of energy, frequency, and vibration."

Nikola Tesla

In this chapter we discuss the working principles, theoretical framework, and experimental characterization of microwave resonators used in superconducting qubit architectures. A comparative study of resonators fabricated from niobium titanium nitride, tantalum, and niobium is presented, alongside optimized fabrication recipes and measured quality factors to guide material selection for high-performance superconducting qubit platforms.

In a circuit QED setup, controlling and measuring the qubit state requires a carefully engineered interaction between the qubit and its environment. This is achieved through dedicated control lines — the drive line and flux line — which manipulate the qubit state, while resonators serve as readout channels, all operating via the exchange of microwave photons. These same interactions, however, inevitably open leakage channels through which environmental noise can degrade qubit coherence, necessitating careful engineering of the microwave mode structure to suppress unwanted coupling. In two-dimensional circuit QED architectures, this mode engineering can be realized using coplanar waveguide (CPW) transmission line geometry, which provides tight transverse confinement of electromagnetic fields, significantly reducing both radiation and dielectric losses compared to alternative geometries. Further, implementation of CPW features benefit from the micro-fabrication techniques and their recent advancements[1].

The microwave energy in the CPW geometry propagates using the quasi-TEM mode[2], and the line is designed to match a 50 Ω impedance environment through careful selection of geometric parameters. As shown in Fig. 3.1(a), the cross-section of the CPW line displays the electromagnetic field distribution, where **s** denotes the signal line width and **d** the gap between the signal line and the ground plane. The characteristic impedance is determined by the ratio **s/d** and the substrate height **h**, calculated using conformal mapping formulas [3].

For qubit readout, a resonator is formed by defining standing-wave electro magnetic (EM) modes along a finite-length CPW segment. The resonance frequency is inversely proportional to the resonator length, and the resonator must be coupled to a feedline transmission line to allow microwave photons to enter and exit. This coupling must be carefully engineered: sufficiently strong coupling is required to achieve an adequate signal-to-noise ratio (SNR) during readout, yet excessive coupling introduces a significant photon leakage channel that degrades the resonator quality factor and, in turn, qubit coherence.

Fig. 3.1(b) shows a $\lambda/2$ resonator with the electric field intensity (voltage) profile indicated in pink. The voltage is maximum at both ends of the resonator and vanishes at the center, where the current reaches its maximum — a field distribution characteristic of a half-wavelength standing EM wave. The characteristic impedance and resonance frequency of such a resonator can be expressed as [4, 5]:

$$Z_0 = \sqrt{\frac{L}{C}} \tag{3.1}$$

$$\omega_r = \frac{\pi}{l\sqrt{LC}} \tag{3.2}$$

where L and C are the inductance and capacitance per unit length of the CPW,

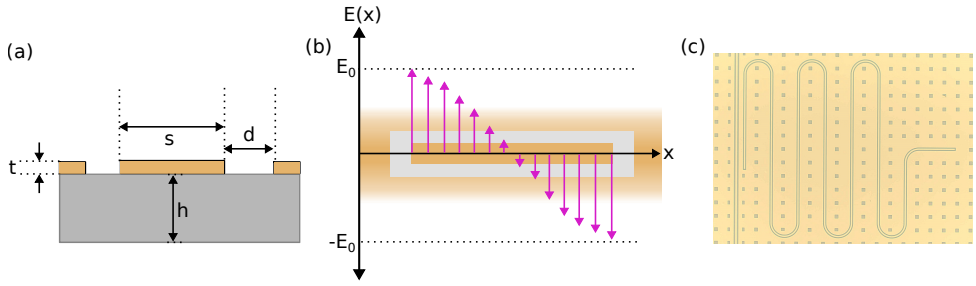

Figure 3.1: CPW geometry:(a) The cross -section view of a CPW based microwave line is shown, with different dimensions of structures. (b) A CPW $\lambda/2$ resonator is shown with electric field distribution along the length of the resonator. (c) Optical microscope image of a single resonator coupled to a feed line is shown.

and l is the physical length of the resonator. In the planar CPW geometry, the qubit is capacitively coupled to one end of the resonator, where the voltage antinode ensures the electric field amplitude is maximum, thereby maximizing the coupling.

While the $\lambda/2$ standing wave mode is an option, other modes can also be chosen to make a resonator. When selecting the standing mode of the resonator, both coupling efficiency and the physical footprint of the resonator on the chip must be considered. For example, a $\lambda/2$ resonator occupies more space compared to a $\lambda/4$ resonator, making the latter preferable when minimizing the footprint of the resonator on the chip. The research has also explored the development of lumped-element resonators [6] to further reduce the required footprint. However, increasing component density on the chip introduces other challenges. Compact layouts can lead to unwanted parasitic self-coupling within resonators. Additionally, the standard solution to Maxwell's equations for CPW geometry-using conformal mapping assumes infinite ground planes on either side of the signal line. In practice, if the ground planes are asymmetric, the altered boundary conditions can distort the electromagnetic field profile, introducing additional asymmetric slot line modes which can travel along the CPW [5, 7]. To address this, techniques such as wire bonds and air-bridges are commonly employed to equalize the ground planes and maintain the intended mode structure.

3.0.1. Evaluating the losses

When quantum information is encoded in the states of superconducting circuits on a chip, it is susceptible to losses. In an ideal scenario, complete isolation of the qubit from their environment would suppress most dissipation channels; however, even in vacuum, spontaneous emission remains unavoidable. In practice, this manifests as relaxation of the qubit state or loss of the photons carry-

ing quantum information. The problem is further compounded in circuit-QED, where scaling up the system introduces an increasing number of loss channels. Unlike the textbook case of a single atom coupled to a single photon mode, a superconducting qubit constitutes a macroscopic quantum state in which a multitude of coupled degrees of freedom are present. These degrees of freedom also contribute additional loss channels. Fortunately, the resonator and the qubit share the similar underlying material stack and fabrication processes albeit the Josephson junction, making the resonator a convenient experimental proxy for assessing qubit quality. Improvements that reduce losses in the resonator translate directly to improvements in qubit coherence. We therefore use the resonator as the primary diagnostic tool to understand loss mechanisms and guide fabrication optimization, before examining what ultimately determines the losses experienced by the superconducting qubit[8].

Throughout this PhD project, the experimental work was focused primarily on the fabrication of superconducting resonators and qubits, and on understanding how fabrication choices determine device loss. To provide a practical framework for this investigation, losses are categorized into two groups based on their origin relative to the chip. This categorization is not a universally adopted convention in the community, but serves as a useful organizing principle for the work presented here:

(1) Extrinsic losses: Losses whose sources lie outside the chip — in the surrounding circuitry or enclosure, are classified as extrinsic. The principal contributors can be the microwave lines running through the dilution refrigerator, which connect the chip to room-temperature electronics and can conduct thermal noise or stray magnetic fields to the device. Imperfect enclosure engineering introduces additional extrinsic loss channels, including photon leakage into the package, spurious low-frequency electromagnetic resonances of the enclosure, and inadequate thermalization of the chip or its immediate environment.

(2) Intrinsic losses: Losses that originate from within the chip itself — from its constituent materials, interfaces, and fabrication-induced defects — are classified as intrinsic. These are considerably more difficult to suppress through system-level design, as they are fundamentally tied to the materials and processes used to fabricate the device. Identifying the precise microscopic pathways responsible for intrinsic loss, and developing improved fabrication processes to mitigate them, formed the central focus of this experimental work. In the next section we focus exclusively on intrinsic losses.

3.1. INTRINSIC LOSSES

The CPW resonators share the vast majority of their loss channels with qubits, as the relevant material interfaces, surface conditions, and fabrication-induced de-

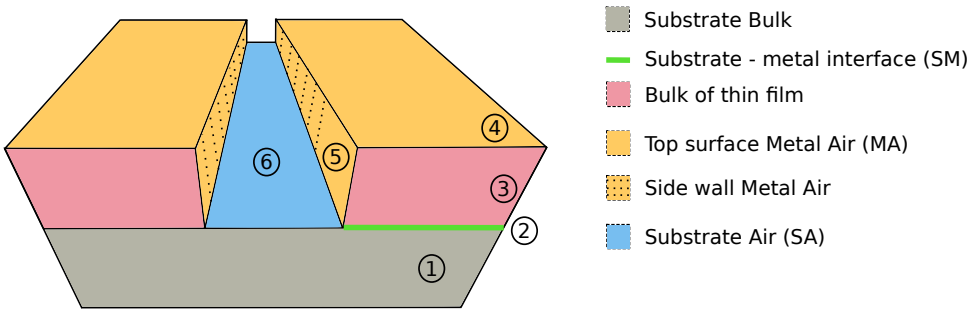

Figure 3.2: Regions of a resonator: Cross-section view of a CPW resonator with different regions is shown

fects are common to both. Specifically, both the resonator and the qubit are patterned with the same lithographic and etching processes. As a result, the critical loss-bearing interfaces are present and geometrically similar in both structures. Although the qubit has additional fabrication step of the Josephson junction, but for a good base layer fabrication process development, resonators are a convenient probe.

A cross-section of a CPW resonator is shown in Fig. 3.2, highlighting five distinct regions representative of those encountered across a circuit QED chip. The loss mechanisms associated with each region are described

1. **Substrate bulk:** The quasi-TEM field distribution of the CPW geometry results in a significant fraction of the electric field energy residing within the bulk of the substrate, meaning that any dielectric loss in the substrate bulk directly contributes to the overall resonator loss. In superconducting qubit chips, silicon and sapphire are the most widely used substrate materials (see Section 2.2), both offering exceptionally low bulk loss tangents relative to other loss channels in the system. However, as losses from other channels are progressively engineered away, substrate loss eventually emerges as a fundamental limit to coherence.

   The primary loss channels within the substrate are associated with two-level system (TLS) defects, present both in the bulk and at the substrate surface. In the bulk, TLS defects can arise from several distinct mechanisms- (1) Crystalline defects within the substrate lattice can host TLS-active impurity species — for example, acceptor impurities such as boron have been shown to contribute measurable dielectric loss [9], (2) fabrication processes that involve physical bombardment of the substrate surface, such as ion milling, can damage the near-surface crystal structure, introducing a disordered layer that hosts TLS defects [10], (3) bulk impurities incorporated

during wafer growth — particularly when the growth method has not been optimized for quantum device applications — can introduce additional TLS-active sites throughout the substrate volume. All of these mechanisms manifest collectively as an enhanced dielectric loss tangent of the substrate material.

Mitigating substrate loss therefore requires: (1) careful selection of the high-purity, low-loss tangent substrates at the outset, and (2) rigorous control of handling and processing conditions throughout fabrication to avoid introducing processing-induced TLS defects.

2. **Substrate metal (SM) interface:** The SM interface is one of the most critical regions governing resonator and qubit performance. In a CPW geometry, the superconducting metal film terminates at this boundary, and it is precisely at this interface that the electric field originates — driving the electric field directly into the substrate from this plane. As a consequence, the electric field intensity is at its maximum at this interface, and any dielectric loss channel present within the interfacial region, couples with maximum strength to the electromagnetic mode. Further, this interface is influenced by the deposition process to a great degree, making process control critical during preparation of wafer before deposition. Few mechanisms through which the losses can occur have been discussed here -

   (i) Surface chemistry: The presence of native oxides or hydrides on the substrate surface prior to metal deposition introduces TLS-active species that can become buried beneath the metal film, where they interact strongly with the electromagnetic field[11, 12]. Proper chemical surface preparation of the substrate immediately before deposition is therefore critical to suppressing this loss channel. In the work presented here, significant effort was devoted to optimizing this surface preparation procedure for our specific substrate materials. 2.3.1

   (ii) Organic residues: Fabrication processes — particularly those involving polymer-based photo resists — can leave organic residues on the substrate surface that are highly detrimental if incorporated into the SM interface[10]. Oxygen plasma ashing and wet chemical cleaning are standard approaches to removing such residues. In our processing flow, a nitric acid cleaning step is included as part of the substrate pre-treatment prior to metal deposition to ensure a contamination-free surface. 2.3.1

   (iii) Ion milling damage: Ion milling performed prior to metal deposition can amorphize the near-surface crystal layers of the substrate. For the incoming metal deposition this increases the effective lattice parameter mis-

match at the interface and disrupts the grain growth of the subsequently deposited metal film, leading to degraded micro-structural quality and increase losses [10]. For this reason, ion milling is generally avoided as a surface preparation step when depositing the superconducting base layer.

(iv) Lattice matching and surface quality: The lattice constants of the chosen metal film should be as close as possible to those of the substrate, since significant mismatch introduces strain into the deposited film, disrupts grain growth, promotes the formation of crystalline defects, and can generate additional TLS loss sites 2.3.2. In the case of niobium on silicon, for instance, the large lattice mismatch results in suboptimal film quality, which was observed in our own deposition experiments. The substrate surface can be further conditioned to promote better film growth through techniques such as high-temperature annealing, which in the case of sapphire improves crystal surface orientation and enhances the quality of the subsequently deposited metal film 2.3.1.

(v) Interfacial reactions: The combination of deposition temperature and chemical affinity between the substrate and metal can drive interfacial reactions that introduce lossy secondary phases. For example, tantalum deposited at elevated temperatures on sapphire does not react adversely, whereas the same process on silicon leads to the formation of tantalum silicide species at the interface, which act as loss channels[13]. A similar silicidation reaction can occur with niobium on silicon at high deposition temperatures, and must be carefully avoided through appropriate control of process conditions[14].

Thus the SM interface loss channels should be collectively addressed through a well-considered choice of material combinations and a rigorously optimized deposition process. In particular, achieving the correct polycrystalline microstructure of the superconducting film requires careful tuning of the growth conditions within the deposition chamber, as discussed in detail in Section 2.3.2.

3. **Thin film:** An ideal superconducting thin film is perfectly conductive with zero microwave loss. In practice, however, Cooper pairs within the superconductor can break, generating quasi-particles that drive local regions of the film into the normal state and introduce resistive loss. The quasiparticle losses is an insignificant contributor to resonator losses compared to other losses, but a major one for the qubit[15]. The primary sources of quasi-particle generation are thermal excitations and stray radiation; the thermal contribution is negligible for materials such as NbTiN and Ta

cooled to operating temperatures of around 10 mK, where the thermal energy is far below the superconducting gap. The second intrinsic loss mechanism arises from vortices, creating a normal-state core surrounded by circulating supercurrents [16]. These normal-state cores introduce dissipation as vortices move in response to the microwave currents in the resonator. Vortex motion can be suppressed when a vortex becomes trapped at a defect or inhomogeneity in the film, known as a pinning site. Pinning sites arise naturally from material defects, but can also be introduced deliberately through microfabrication — patterned apertures in the ground plane, commonly referred to as flux holes, provide controlled pinning sites that immobilize vortices and thereby reduce vortex-induced loss[17].

4. **Metal-air (MA) interface:** The metal-air interface, labeled as region 4 in Fig. 3.2, represents the largest surface area of the superconducting film exposed to ambient conditions during fabrication. This exposure inevitably leads to the formation of a native oxide layer on the metal surface, which reforms after the surface treatment steps performed at the end of the fabrication flow. The oxide is a well-established host for TLS defects and contributes significantly to dielectric loss. The magnitude of this contribution is governed by two key parameters: (1) the chemical composition of the oxide, which determines its intrinsic loss tangent, and (2) its thickness, with thicker oxides hosting a proportionally larger density of lossy TLS defects. Minimizing the MA interface loss therefore requires both minimizing oxide thickness and selecting materials whose native oxides are chemically less lossy[18]. Another important region is the sidewall of the metal film at the etched edges of the CPW, which constitutes a subset of the MA interface lying directly in the region of highest electric field concentration, labeled region 5 in Fig. 3.2. This makes the sidewall oxide disproportionately influential on the total loss, and careful control of the etch process and sidewall morphology is therefore critical to achieving less losses[19, 20].

5. **Substrate-air (SA) interface:** The reactive ion etching step that defines the CPW geometry exposes the substrate surface in the gaps between the signal line and ground planes to ambient conditions throughout the remainder of the fabrication process, labeled region 6 in Fig. 3.2 . In the case of silicon substrates, this exposure leads to the formation of a native silicon oxide layer, whereas sapphire substrates do not oxidize under ambient conditions. The SA interface presents a large exposed surface area, and the timescale for silicon oxidation is of the order of minutes, making this interface a significant and practically unavoidable source of TLS loss. Compounding this, the plasma environment during RIE can produce

a chemically distinct oxide at the silicon surface compared to that formed by ambient oxidation — the plasma-driven oxidation yields a different oxidation state and chemical character, which can result in a more lossy interfacial oxide [21]. Post-fabrication surface treatment of buffered oxide etch is employed to remove the process oxide and reduce the SA interface loss contribution 2.3.1.

3.2. QUANTIFYING THE LOSSES:

The resonators are probed by measuring the complex transmission amplitude (S_{21}) of the microwave signal through the transmission line on the resonator chip. The measurement is performed with a vector network analyzer (VNA) using a two-port configuration. The resonators are capacitively coupled to the transmission line in a hanger configuration [22], absorbing power at their respective resonance frequencies. This absorption manifests as a dip in the transmission amplitude at the resonance frequency.

To quantify the losses in the resonator, the quality factor Q is the primary figure of merit. It is defined as the ratio of the energy stored in the resonator to the energy dissipated per cycle of oscillation, and can be written as

$$Q = \frac{\omega}{\kappa}, \tag{3.3}$$

where ω is the resonance frequency and κ is the total loss rate. The loss rate κ receives contributions from two channels: the coupling to the input/output transmission line, κ_{coupling}, and all other internal dissipation mechanisms, κ_{internal}. The total loaded quality factor is therefore given by

$$\frac{1}{Q} = \frac{1}{Q_{\text{internal}}} + \frac{1}{Q_{\text{coupling}}}, \tag{3.4}$$

where Q_{internal} describes the intrinsic loss channels of the resonator, and Q_{coupling} accounts for energy lost to the transmission line through the coupling capacitor. The measurement of (S_{21}) can then be fitted using the relation in [23], to extract the internal and the coupling quality factors individually.

The extraction of the internal quality factor tells us about the extent of the losses inside the resonators. However, since the losses can originate from multiple sources as described in previous section, it is hard to disentangle the source and understand exact pathways in one experiment. Thus, the observed Q provides an estimate to the degree of losses[24]and will not be able to distinguish in understanding which region the loss is coming from, as the different loss mechanism are simultaneously present in multiple regions. To understand the individual contribution, experiments are designed with one specific loss mode in de-

sign and keeping all the other processes and parameters constant. Further, this requires that the process control is good enough to deliver the same resonator fabrication parameters ib each run. In the next section we discuss the measurement in the experimental setup and look at the results.

3.2.1. EXPERIMENTAL SETUP

The chip is measured in a BlueFors LD-400 dilution refrigerator. To connect the chip to external electronics, it is glued into a PCB assembly as shown in Fig. 3.3(a) and (b), and wire-bonded using aluminium wires to a microwave PCB. The PCB is designed with a stripline geometry and a correctly matched impedance of 50 Ω. The chip is further shielded from the environment by an aluminium enclosure, which is designed to eliminate any resonant modes near the qubit frequency range, as shown in Fig. 3.3(c) and (e).

The absolute transmission amplitude is shown in Fig. 3.4(a), with 12 distinct dips visible in the spectrum corresponding to the 12 resonators on the chip. The designed resonator frequencies are typically specified in the range of 4–8 GHz and can be adjusted according to the requirements of each chip. Beyond the resonator features, the baseline of the spectrum does not remain at a consistent 0 dB level. This variation arises from impedance mismatches throughout the full microwave signal path, starting from the VNA port, through the fridge cabling, and back to the VNA. Careful engineering of the in-house microwave packaging has been carried out to account for these impedance mismatches across the wiring and chip housing.

After an initial characterization of the resonator frequencies as shown in Fig. 3.4(a), measurement with sweeping the power of signal is done for each individual resonator, as shown in Fig. 3.4(c). The loaded quality factor is extracted from the full-width at half-maximum (FWHM) of the resonance line-shape in the frequency domain, Δf_r, and the resonance frequency f_r, according to

$$Q_l = \frac{f_r}{\Delta f_r}. \tag{3.5}$$

The power of the signal input on the transmission line (at the RT stage the power is higher and attenuates as it propagates in the microwave lines) can be related to the extracted quality factors by the formulae[25]-

$$\langle n_r \rangle = \frac{Q_l^2 P_{\text{in}}}{\hbar \omega_r^2 Q_c},$$

The photon number dependence of quality factor for one such resonator is shown in 3.4(b) with a increase in quality factor with the increase in input power. The quantity of interest is the quality factor of the resonators at the lowest

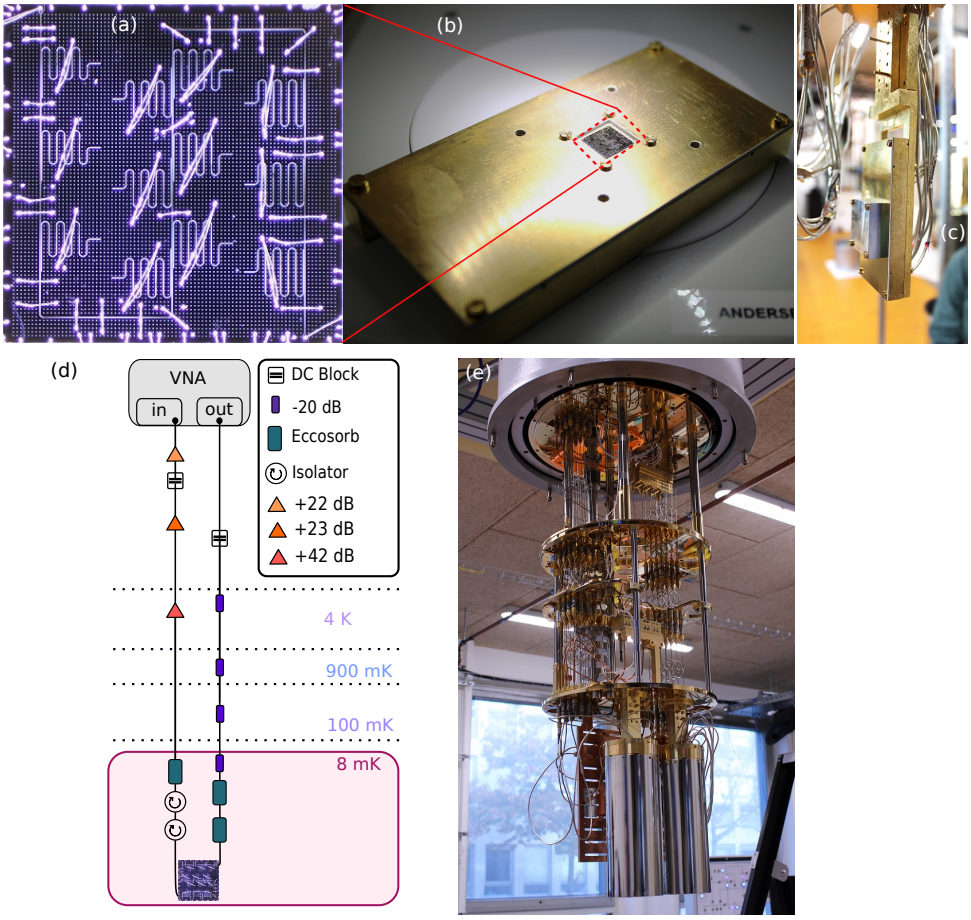

Figure 3.3: Measurement setup : (a) A resonator chip with 12 resonators of different frequencies is imaged under the microscope. Aluminum wire-bonds can be sene in silver color crossing over the resonators.(b) The chip is wire bonded into a microwave PCB packaging, with the bottom side of the package shown. (c) The PCB is mounted into the fridge cold-finder with an aluminum shield on top of the chip as shown. The wires have been connected on the front side. (d) Wiring layout inside the fridge, connected to the VNA (e) The PCB assembly enclosed in the shields is shown

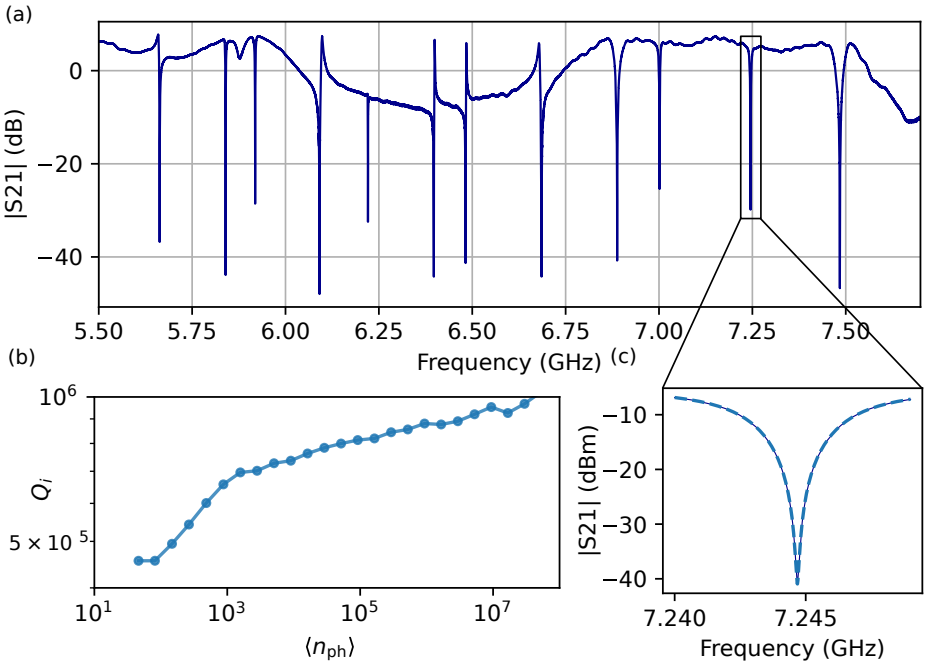

Figure 3.4: Resonator measurements: (a) VNA measurements with the scattering parameters is shown for the probed frequency range around resonators (b) Internal quality factor dependence on photon numbers is plotted (c) Zoomed in view of the resonator is shown for the calculation of FWHM

powers, corresponding to lower photon regime, which is closer to the operating regime for a qubit during operations.

3.3. RESULTS:

3.3.1. PATHFINDER WITH OPTIMIZED SURFACE MORPHOLOGY:

The first step towards fabricating CPW resonators was to establish a set of base processes for the full fabrication flow. Towards this goal, we began by fabricating basic resonator chips, which were characterized in a dip-stick experiment using a VNA to observe resonance features. The quality factors obtained from this initial process were in the range of 10^4. One important outcome of this work was the process validation of the complete resonator chip fabrication flow and its integration with a microwave measurement environment. This was further validated through cryogenic microwave characterization and by confirming that the electromagnetic simulations were in agreement with measured resonator frequencies and coupling strengths. While these steps are well established in the field, having set up fabrication processes in a new laboratory, they represented an im-

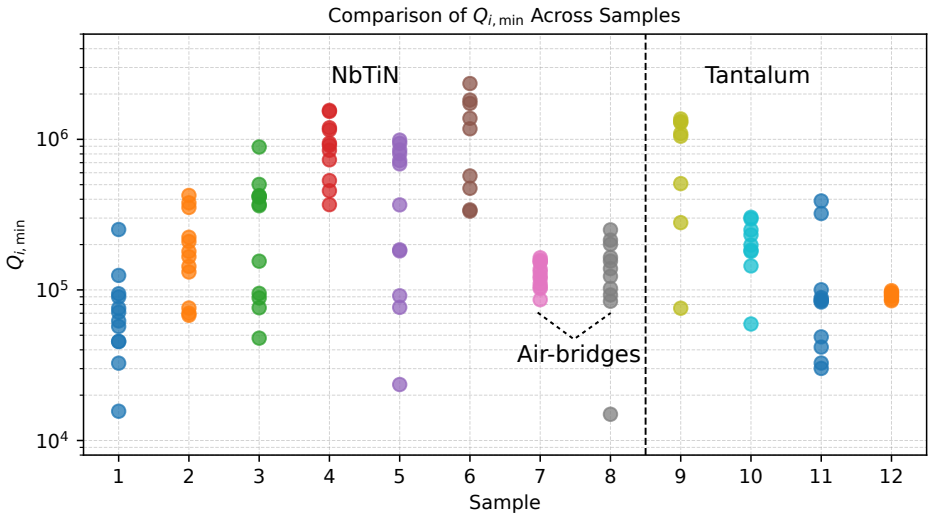

Figure 3.5: Internal Quality factors of NbTiN based resonators for different samples. The quality factor numbers are for the respective lowest measured power corresponding to each sample.

portant milestone in confirming that our processes were functioning correctly and that we were proceeding in the right direction.

A further important finding from this initial process was that our etching recipe was not optimized. The silicon surface morphology after etching was highly irregular, exhibiting crater-like features across the substrate. This was attributed to a non-optimal RF power and gas composition in the reactive ion etching (RIE) recipe 2.5.2. With the assistance of Figen Y., we were able to optimise the RF power and the SF_6/O_2 gas ratio to achieve a significantly smoother surface morphology. Resonators fabricated using the resulting optimized two-step etching recipe are discussed in the following section.

3.3.2. OPTIMIZATION OF SURFACE CLEANING RECIPES:

The majority of losses in CPW resonators originate from surface dielectric loss, particularly from metal oxides formed on the superconducting film. The goal of this optimization was to first understand the dominant loss mechanisms and then improve the surface treatment step performed at the end of the fabrication flow. Although qubit fabrication requires additional processing steps beyond the resonator level, achieving higher internal quality factors at the resonator stage validates the improvements made to the fabrication process.

In the case of NbTiN, niobium forms multiple oxide species — NbO, NbO_2, and Nb_2O_5 — depending on the local coordination of the Nb atoms within the

Sample #	Metal film/Substrate	Summary of surface treatment
1	NbTiN/ Si	IPA
2	NbTiN/ Si	BOE(30 s)
3	NbTiN/ Si	HNO3(1min) -> BOE(15 min)
4	NbTiN/ Si	BOE(10 min)
5	NbTiN/ Si	BOE(10 min)
6	NbTiN/ Si	HNO3(1min) -> BOE(30 min)
7	NbTiN/ Si	Airbridge processing
8	NbTiN/ Si	Airbridge processing
9	Ta/ Sapphire (StarCryo)	HNO3(1min) -> BOE(30 min)
10	Ta/ Sapphire (StarCryo)	HNO3(1min) -> BOE(30 min)
11	Ta/ Sapphire (475°C)	HNO3(1min) -> BOE(30 min)
12	Ta on Nb/Sapphire	HNO3(1min) -> BOE(30 min)

Table 3.1: Summary of the resonator samples fabricated and measured in this study. Each sample is identified by number, with the corresponding superconducting metal film and substrate listed alongside the surface treatment applied at the end of the fabrication flow.

material. Each of these oxides has a different loss tangent, with losses increasing in the order NbO < NbO_2 < Nb_2O_5. The formation kinetics of these oxides are time-dependent and differ between species: while NbO begins to form relatively early, Nb_2O_5 requires significantly longer timescales to develop, on the order of 100 hours. This implies that the cumulative loss tangent of the oxide layer at the metal-air (MA) interface evolves over time. Furthermore, the growth of niobium-based oxides does not self-terminate; rather, the oxide thickness continues to increase, reaching values on the order of 5–10 nm at the MA interface [21].

This behavior contrasts with aluminum, which grows a comparatively thin (2–3 nm) and self-terminating native oxide. The loss tangent of AlO_x is also relatively lower than that of the niobium oxides, making aluminum more favorable when assessed purely on the basis of its native oxide properties.

In addition to the MA surface, the substrate-air (SA) interface in the etched gaps of the CPW also contributes to the total dielectric loss. Since silicon has been used as the substrate for the majority of our devices, silicon oxide — primarily SiO_2, is the dominant loss source at the SA interface.

Both the niobium-based and silicon-based surface oxides can be removed by wet etching using buffered oxide etch (BOE). A 1:7 BOE solution has been used in our process for this purpose 2.3.1.

SAMPLE 1-2:

An initial experiment was carried out to establish the final cleaning step using IPA. Following the resist stripping step, which removes the majority of organic

residues, a complete dip in an IPA beaker cleans the surface of any remaining NMP residues, leaving a hygroscopic surface. The quality factors for sample 1, based on this IPA-only cleaning recipe, average around 7×10^4. The low internal quality factor can be attributed to the presence of native oxides at both the SA and MA interfaces.

For sample 2, the final surface cleaning step was extended to include a short BOE dip of 30 s. This resulted in a modest improvement in Q_i, which can be partially attributed to the removal of SA interface oxides. In [21], the etch rate of silicon oxide was characterized, showing that the SA oxides are removed on a timescale of approximately 30 s. However, the regrowth of silicon oxide is rapid, occurring on a timescale of a few minutes, such that the SA interface is again covered with oxide before the sample can be put int he fridge. This fast regrowth makes it difficult to directly attribute the differences in Q_i between sample 1 and sample 2 solely to the removal of SA oxides. An alternative explanation, supported by XPS measurements discussed in [21], is that the process oxide grown during fabrication may have a different bonding environment at the silicon surface compared to the oxide that regrows under ambient conditions following BOE etching. This difference in bonding environment can influence the density and character of TLS defects present at the interface, and may therefore account for the observed difference in losses between sample 1 and sample 2.

SAMPLE 3-6:

In the follow-up experiments with samples 3-6, the BOE etching time was increased based on the assumption that a longer etch would also remove oxides at the MA interface and thereby improve Q_i. Additionally, a nitric acid cleaning step was introduced prior to BOE etching to remove any trace organic contamination remaining from the fabrication process. An important consideration here is that nitric acid is compatible with NbTiN and effectively removes organic residues without etching the metal, making it a suitable choice at this stage of the process. RCA cleaning was not considered as an option due to its aggressive etching behavior with metals. A 1 min nitric acid clean was adopted for organic removal.

Samples 3 and 4 were prepared in the same fabrication run and measured together in a single cooldown. For sample 3, the surface treatment consisted of 1 min in nitric acid followed by 15 min in BOE. This yielded an improvement in the highest Q_i achieved, but was accompanied by a larger spread across the resonators on the chip. Sample 4 was prepared without the nitric acid step, in order to isolate the effect of nitric acid on Q_i. From the perspective of MA interface chemistry, nitric acid is a strong oxidising agent and may promote the formation of a thicker oxide layer, which could explain the slightly lower average Q_i observed in sample 3 compared to sample 4. However, no strong conclusion can be

drawn from this comparison alone, as measurements from the similar samples 5 and 6 show variation that complicates the interpretation.

Samples 5 and 6 were likewise fabricated in the same run to further investigate the role of BOE etching duration on MA interface oxide removal. Sample 5 was processed using the same BOE etching time as sample 4 but exhibited a larger spread in Q_i across resonators. No definitive explanation for this variation has been identified, though it may arise from spatial non-uniformity of the oxide across the chip, introduced during handling in the fabrication process.

For sample 6, the BOE etching time was extended further to 30 min. The rationale for this was that the BOE etch rate for niobium oxide is considerably slower than that for silicon oxide [21], and a longer etch is therefore required to appreciably reduce the MA interface oxide thickness. The nitric acid step was retained for organic removal, with the expectation that the extended BOE etch would more than compensate for any additional oxide growth induced by the oxidising acid. The combined surface treatment of 1 min nitric acid followed by 30 min BOE yielded the highest Q_i values with the lowest spread across resonators. This recipe was subsequently adopted as the standard surface treatment in our NbTiN fabrication flow and has been used for all NbTiN-based chips since.

3.3.3. INTEGRATION OF AIR-BRIDGES: SAMPLE 7-8

As chips scale to integrate more components — such as additional resonators and qubits — the inter-component spacing shrinks, making it increasingly difficult to maintain the correct electromagnetic mode structure. The integration of air-bridges into the fabrication flow therefore became the next logical step, and this work was primarily carried out by Kumaravadivel P.

The air-bridges were fabricated using the recipe described in Appendix 3.5. The fabrication of air-bridges introduces additional lithography and deposition steps on top of already-defined resonators, involving multiple resist coating and baking cycles, as well as wet processing steps. This wet processing degrades the surface quality achieved by the preceding nitric acid and BOE cleaning steps.

Samples 7 and 8 were fabricated in the same run for direct comparison. Sample 7 had air-bridges integrated, while sample 8 did not have air-bridges but was subjected to the same wet processing steps used in the air-bridge fabrication, in order to isolate the effect of the wet processing on Q_i. The internal quality factors of both samples were comparable, but limited to approximately 10^5, which is one order of magnitude lower than the best values achieved with the acid and BOE surface treatment alone. This confirms that the wet processing associated with air-bridge fabrication is a significant source of contamination, and that organic residues introduced during these steps severely degrade the internal quality fac-

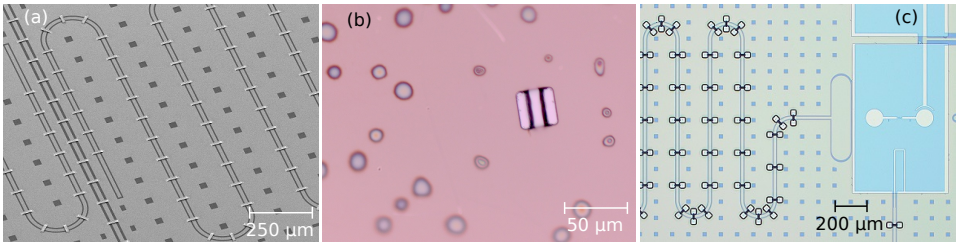

Figure 3.6: Air bridges (a)SEM of the resonators integrated with the airbridges (b) Droplets after the airbridge processing with vapor HF (c) Air-bridges with the PMGI not stripped, with the qubit island area developed to expose the silicon

tor.

While air-bridge integration is necessary for the chip architecture, the associated processing clearly deteriorates Q_i. Since the air-bridges are made of aluminium, it is not possible to repeat the BOE cleaning step to remove surface oxides without etching the bridges. Two alternative approaches were therefore explored to recover higher quality factors after air-bridge integration.

(1) Vapour HF cleaning: Vapour HF was considered as a means to remove surface oxides similarly to BOE, while being compatible with aluminium. A test experiment was carried out on a chip with air-bridges, but liquid droplets were observed across the chip surface following the treatment, likely arising from a reaction between the HF vapour and the surface, as shown in Fig. 3.6(b). Although further chemical analysis would be needed to identify the exact nature of these droplets, the approach was abandoned due to the concern that the droplets may leave fluoride-based residues on the chip surface.

(2) Selective resist masking during BOE cleaning: Since the qubit area and in particular the region where Josephson junctions are to be fabricated — is a critical interface, a chip was prepared in which the air-bridges were present but the qubit island regions were left free of resist, as shown in Fig. 3.6(c). The resist was developed such that the silicon surface in the qubit island area was exposed. This approach allows the Josephson junction resist stack to be subsequently spun and patterned, and the junction fabrication flow to proceed normally. Critically, the BOE step used for silicon oxide removal prior to junction deposition would clean the exposed qubit island area without affecting the air-bridges, which remain protected by the overlying resist. Although a chip was fabricated to test this concept, the subsequent experiments did not ultimately require the full junction integration to be demonstrated. This approach can be further assesed for future usage.

3.3.4. CHANGING TO TANTALUM FOR IMPROVING Q: SAMPLE 9-12

The metal-air interface in NbTiN is one of the limiting factors in achieving very high internal quality factors, due to the formation of lossy surface oxides. To mitigate the influence of these oxides and improve Q_i, we explored replacing the superconducting thin film with tantalum. Recent experiments have demonstrated that Ta-based resonators and qubits can achieve significantly higher quality factors and coherence times [26].

Samples 9 and 10 were fabricated using the same recipe but in separate fabrication runs. Both were made from 200 nm thick tantalum deposited on sapphire substrates, commercially supplied by Starcryo. The internal quality factors of sample 9 were already approaching 10^6 in the first set of devices, which was a strong indication that the fabrication flow for Ta was well suited to yielding high-performance resonators without extensive optimisation. The fabrication recipe followed closely that of [26]. One of the critical factors governing the performance of Ta-based resonators is the crystalline phase of the tantalum film: the α-phase is associated with low loss, while the β-phase is lossy. The role of the α- and β-phases and the conditions required to grow each are discussed in Section 2.3.2. In the commercial wafers used for samples 9 and 10, α-phase tantalum was confirmed (see Fig. 2.6), consistent with the high quality factors observed.

To reduce dependence on a commercial supplier and to develop an understanding of the growth conditions required for high-performing tantalum films, we proceeded to grow tantalum in-house. Sample 11 was fabricated by sputtering Ta on sapphire using the deposition recipe described in 2.3.2, with deposition carried out at an elevated substrate temperature of 475 °C, which is expected to promote α-phase growth. The crystalline phase of the resulting film could not be fully characterized by XRD, but preliminary SEM data suggests a mixed α/β-phase microstructure (see Fig. 2.6). This phase mixing is a likely contributor to the lower Q_i observed for sample 11.

Sample 12 was fabricated using a 200 nm thick niobium film capped with a 10 nm tantalum layer. This approach is motivated by the concept of metal capping [27], whereby the lossy niobium surface oxides at the MA interface are replaced by the comparatively more benign tantalum oxide. Despite this, the measured Q_i of all resonators in sample 12 was limited to approximately 10^5. The mechanism responsible for this limitation has not been conclusively determined, but may be related to a non-optimized niobium deposition recipe.

3.4. CONCLUSION & OUTLOOK:

Surface treatment: The resonator experiments presented in this chapter have allowed us to identify the dominant loss mechanisms and systematically refine the fabrication process. Beginning with initial pathway-finding experiments to es-

tablish a working fabrication recipe for NbTiN resonators, we progressively optimised the surface treatment steps through successive cooldown measurements. Over time, surface treatment experiments were run in parallel with qubit development, leading to iterative refinements. The final adopted recipe consists of an organic cleaning step with nitric acid, followed by a long BOE dip for oxide removal. This optimised surface treatment resulted in an improvement of the internal quality factor by approximately two orders of magnitude compared to the initial devices.

The improvement in Q_i has been attributed primarily to the removal of niobium-based oxides at the MA interface and silicon-based oxides at the SA interface. A deeper understanding of the remaining TLS loss could be gained by investigating the role of hydrogen incorporation into the niobium film. Hydrogen trapped at interstitial sites in niobium is known to increase TLS density, and niobium hydrides are themselves a source of microwave loss. Since the BOE etching time has been significantly extended in the optimised recipe, the exposure of the niobium surface to hydrogen in the BOE solution is also increased. Further characterization is therefore warranted to assess whether switching to an alternative oxide removal chemistry — one with reduced hydrogen exposure — could yield further improvements in Q_i.

Air-bridges: As the number of components on the chip increased, air-bridge integration was tested to mitigate grounding issues associated with higher component density. The results showed a reduction in Q_i due to the additional wet processing steps required for air-bridge fabrication, as well as increased surface constraints introduced by the bridges themselves. Techniques for integrating Josephson junctions in the presence of air-bridges, as well as vapour HF cleaning as an alternative oxide removal method, were also explored, but did not yield conclusive results.

A surface cleaning procedure that is compatible with air-bridge integration remains a desirable objective for future work. One potential route is to fabricate air-bridges from tantalum or NbTiN rather than aluminium, which would make the bridges compatible with BOE-based oxide cleaning without risk of etching the bridge material. If the air-bridges are to be deposited by sputtering, lift-off-related issues would need to be addressed, and the deposition of a sufficiently thick metal layer onto resist presents an additional process challenge. One advantage of using the Super-AJA sputter system is the ability to deposit tantalum, followed by a titanium adhesion layer, and then tantalum again in a single pump-down without breaking vacuum, which could facilitate the development of such a process.

Tantalum: Tantalum on sapphire and niobium on silicon with a tantalum capping layer were explored as alternative material systems. Commercial tantalum-

on-sapphire wafers were tested first, yielding significantly improved Q_i without requiring extensive process optimisation. Attempts to replicate this performance using in-house-grown tantalum films did not achieve the same results, with the quality factors from in-house wafers falling below the best values obtained from NbTiN resonators. A potentially promising direction identified in these experiments is the tantalum-on-silicon material combination, which warrants further investigation to validate whether high quality factors can be achieved on this substrate.

Additionally, the long BOE exposure optimised for NbTiN surface treatment may not represent the most appropriate recipe for tantalum-based devices. Further studies are needed to understand the interplay between BOE etching of tantalum oxides, the role of hydrogen exposure in degrading tantalum film performance, and the trade-off between oxide removal and etch-induced surface damage, in order to develop an optimised surface cleaning procedure for tantalum resonators.

3.5. APPENDIX A

Table 3.2: Details of the individual fabrication steps in the airbridge fabrication flow.

Process	Details
Spin coat PMGI resist (Step 1)	For the support pads (anchorages) of the bridge. PMGI SF15 @ 2500 rpm (recipe 6), bake @ 190°C for 5 min, used plastic pipette
Exposure (EBPG)	Airbridge dose test – Step 1 (the anchorages of the bridge). Dose range: 300–550 μC/cm^2 (tested with 50 μC/cm^2 increments). Beam: 250nA_400um.beam_100, BSS: 0.1
Development of PMGI (Step 1)	Post EBPG development (the anchorages of the bridge). AZ 400k : H_2O (1:5) 60 s; then AZ 400k : H_2O (1:10) for 30 s and another 30 s H_2O
Reflow PMGI	Bulging of resist for the bridge. @ 225°C, 5 min on a hot plate, closed. To create the bulge of PMGI between the two EBPG written pads
Spin coat double layer PMMA resist (Step 2)	For the bridge. 495 K A8 PMMA @ 1000 rpm baked at 185°C for 3 min; then 950 K A8 PMMA @ 2000 rpm baked at 185°C for 3 min
Exposure (EBPG, Step 2)	For the airbridge (the bridge). Dose range: 700–1200 μC/cm^2 (tested with 100 μC/cm^2 increments). 40nA_200um.beam_100, BSS = 0.02 nm.
Post EBPG develop double layer	MIBK : IPA (1:2) for 1 min followed by 30 s IPA rinse
RIE O_2 cleaning	Before Al deposition. 20 W power, 45 sec recipe
Metal deposition	Plassys 100 nm Al + 200 nm Ti + 100 nm Al
Lift-off (Step 1)	In Anisol, overnight and then transferred to warm IPA. Took out from warm IPA (70°C), no blow drying
Lift-off (Step 2)	In NMP, 80°C for approx. 3 hrs and then transferred to warm IPA (70°C). Took out the substrate from warm IPA (70°C)

DATA AVAILABILITY

The data for the experiments along with the scripts to anlayse the data can be found at [28] with an open CC by 4.0 license.

REFERENCES

[1] A. Wallraff et al. “Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics”. In: *Nature* 431 (2004), pp. 162–167. DOI: 10.1038/nature02851.

[2] David M. Pozar. *Microwave Engineering*. 4th. Hoboken, NJ: Wiley, 2012. ISBN: 978-0-470-63155-3.

[3] Philippe Veyres and Victor Fouad Hanna. "Extension of the application of conformal mapping techniques to coplanar lines with finite dimensions". In: *International Journal of Electronics* 48.1 (1980), pp. 47–56. DOI: 10.1080/00207218008901035.

[4] M Goeppl et al. "Coplanar waveguide resonators for circuit quantum electrodynamics". In: *Journal of Applied Physics* 104.11 (2008), p. 113904.

[5] Rainee N Simons. *Coplanar waveguide circuits, components, and systems*. Vol. 165. John Wiley & Sons, 2004.

[6] S Doyle et al. "Lumped element resonators for microwave kinetic inductance detectors". In: *Journal of Low Temperature Physics* 151.1 (2008), pp. 467–473.

[7] M Ribó and L Pradell. "Circuit model for mode conversion in coplanar waveguide asymmetric shunt impedances". In: *Electronics Letters* 35.9 (1999), pp. 713–715.

[8] A. Megrant et al. "Planar superconducting resonators with internal quality factors above one million". In: *Applied Physics Letters* 100 (2012), p. 113510. DOI: 10.1063/1.3693409.

[9] Zi-Huai Zhang et al. "Acceptor-Induced Bulk Dielectric Loss in Superconducting Circuits on Silicon". In: *Physical Review X* 14.4 (2024), p. 041022. DOI: 10.1103/PhysRevX.14.041022. URL: https://doi.org/10.1103/PhysRevX.14.041022.

[10] A. Dunsworth et al. "Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits". In: *Applied Physics Letters* 111.2 (2017), p. 022601. DOI: 10.1063/1.4993577.

[11] Zuhawn Sung et al. "Formation of niobium hydride precipitates in superconducting qubit chips". In: *Physical Review Applied* (2024). DOI: 10.1103/PhysRevApplied.21.034034.

[12] C. T. Earnest et al. "Substrate surface engineering for high-quality silicon/aluminum superconducting resonators". In: *Superconductor Science and Technology* 31 (2018), p. 125013. DOI: 10.1088/1361-6668/aae548.

[13] Moritz Singer et al. “Tantalum thin films sputtered on silicon and on different seed layers: material characterization and coplanar waveguide resonator performance”. In: *Proceedings of the 5th IEEE International Conference on Quantum Computing and Engineering (QCE 2024)*. Montreal, Canada: Institute of Electrical and Electronics Engineers Inc., 2024, pp. 1197–1202. DOI: 10.1109/QCE60285.2024.00141.

[14] Xuezeng Lu et al. “Stability, metallicity, and magnetism in niobium silicide nanofilms”. In: *Physical Review Materials* 6 (2022), p. 064402. DOI: 10.1103/PhysRevMaterials.6.064402.

[15] Matt McEwen et al. “Resisting High-Energy Impact Events through Gap Engineering in Superconducting Qubit Arrays”. In: *Physical Review Letters* 133 (2024), p. 240601. DOI: 10.1103/PhysRevLett.133.240601.

[16] C Song et al. “Microwave loss in superconducting resonators due to trapped flux vortices”. In: *Applied Physics Letters* 94.10 (2009), p. 102504.

[17] B. Chiaro et al. “Dielectric surface loss in superconducting resonators with flux-trapping holes”. In: *Superconductor Science and Technology* 29 (2016), p. 104006. DOI: 10.1088/0953-2048/29/10/104006.

[18] M. Virginia P. Altoé et al. “Localization and Mitigation of Loss in Niobium Superconducting Circuits”. In: *PRX Quantum* 3.2 (2022), p. 020312. DOI: 10.1103/PRXQuantum.3.020312. URL: https://doi.org/10.1103/PRXQuantum.3.020312.

[19] Kevin D. Crowley et al. “Disentangling Losses in Tantalum Superconducting Circuits”. In: *Physical Review X* 13 (2023), p. 041005. DOI: 10.1103/PhysRevX.13.041005.

[20] Ray D. Chang et al. “Eliminating Surface Oxides of Superconducting Circuits with Noble Metal Encapsulation”. In: *Physical Review Letters* (2024). DOI: 10.1103/PhysRevLett.134.096501. arXiv: 2408.13051.

[21] M. Virginia P. Altoé et al. “Localization and mitigation of loss in niobium superconducting circuits”. In: *PRX Quantum* 3.2 (2022), p. 020312. DOI: 10.1103/PRXQuantum.3.020312.

[22] Shantam S Ganjam. “Suppression of parasitic modes in superconducting circuits”. PhD thesis. New Haven, CT: Yale University, 2021. URL: https://elischolar.library.yale.edu/gsas_dissertations/256/.

[23] S Probst et al. “Efficient and robust analysis of complex scattering data under noise in microwave resonators”. In: *Review of Scientific Instruments* 86.2 (2015), p. 024706. DOI: 10.1063/1.4907935.

[24] Janka Biznárová. "Development of high-coherence superconducting devices for quantum computing: Fabrication process development, materials analysis, and device characterization". PhD thesis. Gothenburg, Sweden: Chalmers University of Technology, Sept. 2024. URL: https://research.chalmers.se/en/publication/542528.

[25] Jiansong Gao. "The Physics of Superconducting Microwave Resonators". PhD thesis. California Institute of Technology, 2008. DOI: 10.7907/RAT0-VM75. URL: https://resolver.caltech.edu/CaltechETD:etd-06092008-235549.

[26] Alexander PM Place et al. "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds". In: *Nature Communications* 12.1 (2021), pp. 1–9. DOI: 10.1038/s41467-021-21993-y.

[27] Mustafa Bal et al. "Systematic Improvements in Transmon Qubit Coherence Enabled by Niobium Surface Encapsulation". In: *npj Quantum Information* 10 (2024), p. 43. DOI: 10.1038/s41534-024-00840-x.

[28] *Data for the thesis titled: Engineering Noise-Resilient Superconducting Qubits: Fabrication and control of fluxonium and cos(2phi) systems.* https://doi.org/45267100-d50e-4ed5-a91f-53e4a899238f. 2026.

4

JUNCTIONS

"Every great work of art is a monument of the human spirit."

Hans-Georg Gadamer.

In this chapter, we discuss the Josephson junctions designed for fluxonium qubits. We examine how variations in fabrication parameters influence the resulting junction characteristics, providing insights into optimizing their performance. The results summarize the advancements achieved in refining and improving the overall junction fabrication process.

4.1. INTRODUCTION

The Josephson junction is a fundamental building block of circuit quantum electrodynamics (cQED) and lies at the heart of construction of superconducting quantum circuits. It introduces the essential nonlinearity in the circuit, required to generate anharmonic energy potentials, enabling the formation of well-defined qubit states. The operation of the junction is governed by the Josephson effect[1], which arises from the coherent tunneling of Cooper pairs across a thin insulating barrier separating two superconductors. When two superconducting regions are placed in close proximity to each other with a weak link in between, without the application of any voltage, a dissipationless current is observed through this weak link, given by:

$$I = I_c \sin(\Delta\varphi) \tag{4.1}$$

where I_c is the critical current of the weak link and φ is the phase differnece between the two superoncducitng regions.This effect is called the DC Jospeshson effect . The superoncducting regions can be superoncducting materials or a material proximitized by a supercondcutor leading to superoncductivity. The weak link composes of either an insulating barrier, normal metal or constrictions. Each combination will have the realtion change based on the material and geomterical properties. The effect happens due to GL wavefunction overlap in the weak link.

Similarly, if a voltage is applied across the junction a phase difference build as:

$$\frac{d\varphi}{dt} = \frac{2\pi}{\Phi_0} V \tag{4.2}$$

realtion 4.1 and 4.2 leads to the Josephson energy as

$$E(\varphi) = \int I(\varphi) V \, dt = \int I_c \sin(\varphi) \frac{\Phi_0}{2\pi} d\varphi = -\frac{I_c \Phi_0}{2\pi} \cos\varphi = -E_J \cos\varphi \tag{4.3}$$

where, $E_J = I_c \Phi_o / 2\pi$

ALUMINUM AND SIS JUNCTIONS:

When a Josephson junction is constructed from two superconducting electrodes separated by a thin insulating barrier, it is referred to as an SIS (superconductor-insulator-superconductor) junction. The choice of superconducting material for the electrodes is governed by the loss properties of the material, but the critical consideration is the nature of the insulating barrier. An ideal tunnel barrier must satisfy several criteria: (1) it must be thin enough to permit a tunneling supercurrent, (2) its fabrication must be sufficiently reproducible to yield consistent thin-film characteristics across devices, and (3) from a materials perspective, the

barrier should be as low-loss as possible and remain chemically stable over extended periods.

Aluminum (Al) is a strong candidate that satisfies all of these criteria. It forms a self-terminating native oxide in an oxygen environment, a property that naturally produces a reproducible barrier with well-controlled thickness and stoichiometry. The self-terminating growth mechanism also ensures that the tunnel resistance of the junction remains stable over long timescales [2]. Furthermore, aluminum can be deposited by electron-beam evaporation with relative ease, and the double-angle evaporation technique allows a complete $Al/AlO_x/Al$ junction to be fabricated in a single deposition step. Finally, the thin, stoichiometrically stable AlO_x barrier exhibits a low dielectric loss tangent compared to alternative barrier materials such as niobium oxides or silicon oxide, which host a higher density of TLS defects and lead to degraded qubit performance [3, 4]. For these reasons, the $Al/AlO_x/Al$ material stack is used in the majority of superconducting qubit implementations to realise the SIS Josephson junction.

For an SIS junction, the critical current I_c is related to the normal-state tunnel resistance R_N through the Ambegaokar–Baratoff relation [5],

$$I_c(T)\,R_N = \frac{\pi\Delta(T)}{2e}\tanh\left(\frac{\Delta(T)}{2k_BT}\right), \tag{4.4}$$

which at $T = 0$ simplifies to

$$I_c(0) = \frac{\pi\Delta(0)}{2eR_N}. \tag{4.5}$$

This relation also provides a direct estimate of the Josephson energy,

$$\frac{E_J}{h} = \frac{\Phi_0\,\Delta(0)}{4eh} \times \frac{1}{R_N}, \tag{4.6}$$

where $\Phi_0\,\Delta(0)/4eh \approx 140\,\mathrm{GHz\,k\Omega}$ and Δ is the superconducting energy gap of the electrode material. This relation is particularly useful in practice, as it allows E_J to be estimated from a room-temperature resistance measurement of the junction, without requiring a cryogenic measurement.

4.2. FABRICATION OF THE JOSEPSHOSN JUNCTIONS:

To fabricate the Josephson junction with the $Al/AlO_x/Al$ SIS structure, a double-angle shadow evaporation technique is used. In this technique, one electrode of the junction is deposited first while the other is shadowed by the resist, followed by an in-situ oxidation step to form the tunnel barrier, and finally the second aluminium electrode is deposited with the first arm now shadowed. Two common implementations of this shadow evaporation approach are the Dolan

bridge method [6] and the Manhattan technique [7]. These shadow techniques is one of the easy ways the junction can be made with one lithography step and in our lab we mainly use Manhattan technique for deposition. The complete junction fabrication process is divided into three stages: pre-processing, deposition, and post-processing.

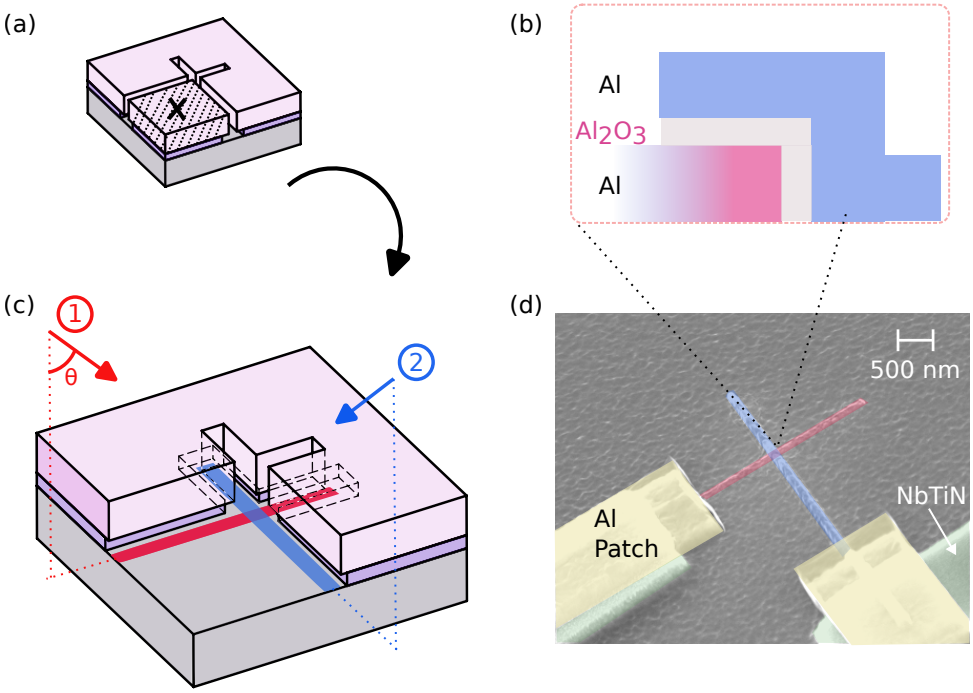

Figure 4.1: Double angled deposition: (a) Substrate with a two layer resist stack, with a developed manhattan junction pattern developed. The cross shows the area not shown when zoomed in. (b) The cross-section view of the SIS junction (c) Deposition angles with the direction and the deposition and the order of the junction arms deposited. Undercut geometry is shown in dotted lines for the full extenet of the junction. (d) False colored SEM of the junction with the galvanic patches is shown.

(1) Pre-processing: The process begins when the base layer has been patterned and the surface has been cleaned of native oxides. The first step of the Manhattan process is spinning of a multilayer resist stack for the junction structures, either a bi-layer or tri-layer configuration depending on the geometry requirements. After resist spinning, two perpendicularly intersecting rectangular areas are exposed using an EBPG 5200 electron-beam lithography system, as shown in Fig. 4.1(a), and subsequently developed in an IPA:H_2O (3:1) mixture. Following development, residual organic contaminants are removed using an oxygen plasma descum process, typically performed in the F1 or Tepla sys-

tem. This is a critical step to ensure clean surfaces prior to metal deposition. The native silicon oxide on the substrate is then removed by a brief immersion of approximately 40 seconds in a 1:7 BOE solution to expose clean silicon and minimise dielectric loss from TLS defects associated with SiO_x [8]. Immediately after, the sample is loaded into the load-lock of the Plassys evaporator.

(2) Deposition: After loading, the Plassys load lock is pumped to high vacuum (5×10^{-7} mBar). The sample is transferred into the main deposition chamber, where the angled evaporation of aluminium takes place. The sample is oriented such that the first deposition is directed along the first rectangular groove, as shown in red in Fig. 4.1(c), with the deposition angle θ defined accordingly. This bottom electrode is 30 nm thick. During this step, a small amount of aluminium is also deposited on the wall of the perpendicular groove with width d_2 (see Fig. 4.4(e)), and at the end of the first electrode, the deposition extends into the undercut region beneath the resist, as shown in Fig. 4.4(c). The significance of these deposited widths is discussed in Section 4.4.1.

After the first deposition, the evaporation chamber is sealed and the oxidation chamber is filled with oxygen to a pressure of a 1-5 mBar. The exact pressure and duration of this oxidation step determine the tunnel resistance of the junction. Following evacuation, the second electrode is deposited with the sample oriented perpendicular to the first deposition direction, as shown by the blue arrow in Fig. 4.1(c). The thickness of this top electrode is 110 nm. Immediately after this deposition, a second oxidation step is performed to form a protective capping oxide on top of the full structure.

(3) Post-processing: After deposition is complete, the sample is unloaded from the load-lock and the aluminium is lifted off in a heated NMP solution at 80 °C for a minimum of 2 hours. It has been observed that longer immersion times result in cleaner lift-off. The sequential execution of these three stages yields a complete Josephson junction. Over time, the individual sub-processes and complete recipe have been optimised to improve junction cleanliness and parameter targeting, as discussed in the results section.

GALVANIC CONNECTION WITH THE AL PATCHES:

In our fabrication process, additional aluminium is deposited at the intersection points of the aluminium electrodes with the NbTiN base layer, as shown in Fig. 4.1(b). During the junction deposition, the oxidation step performed between the two electrode depositions also oxidizes the exposed NbTiN base layer. This results in a niobium oxide barrier forming between the aluminium and the base layer, which introduces an unwanted lossy interface. To mitigate this, aluminium patches are deposited to provide a barrier-free galvanic connection between the junction electrodes and the base layer [9, 10]. These patches are de-

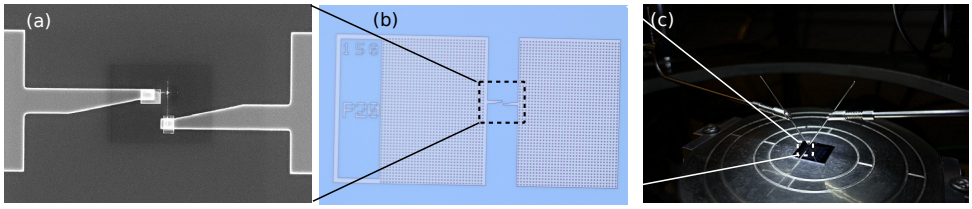

Figure 4.2: Junction probing: (a) SEM top-view image of the junction arm extending onto the NbTiN base layer. (b) Zoomed-out optical microscope image of a test structure with two measurement pads. (c) Probe needles in contact with a test junction sample on the probing stage.

posited using the following recipe.

Process Step	Details
Spinning	MMA EL8 @ 4000 rpm
	PMMA 950 A4 @ 2000 rpm
Baking	5 min @ 185°C each layer
	2 min break between layers
Writing	Beam aperture: 200μm
	Beam current: 98 pA
	Spot size: 2.1 nm
Development	H_2O:IPA (1:3) @ 6 °C for 90 s
	Rinse: 30 s IPA

Table 4.1: Processing parameters for patches

4.3. MEASUREMENT OF R_N

The critical current and Josephson energy of the junction can be calculated from the room-temperature normal-state resistance R_N, measured using a two-probe setup as shown in Fig. 4.2(c), via the relations given in Eqs. 4.6 and 4.4. This room-temperature probing provides a convenient and rapid characterisation of junction parameters without requiring cryogenic measurements. The measurement electronics are described in [11].

4.3.1. TRENDS IN R_N

The Josephson energy E_J is the critical parameter determining the transition frequency of the superconducting qubit. While the charging energy E_C also plays a role, it is comparatively straightforward to control through circuit geometry in the fabrication process. The dependence $E_J \propto 1/R_N$ makes it important to understand and control the junction resistance during fabrication. R_N is governed

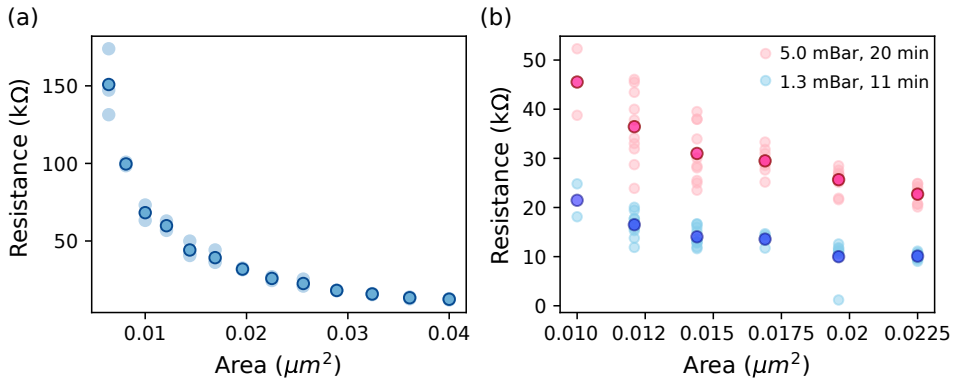

Figure 4.3: Junction resistance vs. area: (a) Measured resistance as a function of junction area for a test device with calibration structures. Light blue points correspond to individual resistance measurements; outlined circles indicate the mean values. (b) Resistance vs. area curves for two different oxidation pressures.

by two factors,

$$R_N \propto \frac{1}{A} \cdot \exp\left(2d\sqrt{\frac{2m\phi}{\hbar^2}}\right),$$

which are discussed below.

1. **Junction area:** $R_N \propto 1/A$, where A is the area of the junction. For a fixed oxidation pressure and time, it follows from Eq. 4.3.1 that $E_J \propto A$, meaning a larger junction area yields a larger E_J. In Fig. 4.3(a), the measured R_N is plotted as a function of junction area $A = w^2$, where w is the junction width, confirming this scaling behaviour.

2. **Tunnelling across the barrier:** $R_N \propto e^{\kappa d}$, where d is the barrier thickness and $\kappa = \sqrt{2m\phi/\hbar^2}$, with ϕ the effective barrier height of the aluminium oxide. For a fixed deposition recipe using the Plassys system, ϕ remains constant [12]. The barrier thickness d can be controlled through the oxidation parameters: both a higher oxygen pressure and a longer exposure time produce a thicker oxide, and consequently a higher R_N. This is demonstrated in Fig. 4.3(b), where an oxidation at 5 mBar for 20 min results in higher resistance values across all junction widths compared to an oxidation at 1.3 mBar for 11 min.

4.4. JUNCTIONS FOR FLUXONIUMS:

The design of the fluxonium [13] qubit requires two main junction-based components: (1) a single Josephson junction in the circuit, and (2) a chain of N junctions in series which act as a linear inductor, yielding an effective inductive energy $E_L = E_{J,\text{arr}}/N$, where $E_{J,\text{arr}}$ is the Josephson energy of each junction in the array. Within each component, the E_J requirements are distinct. The energy parameters listed in Table 4.2 were selected based on numerical simulations optimized for fluxonium designs for improving readout fidelity [14] and improved two-qubit gate performance [15]. Here we use $N = 100$ junctions in the inductor array, which sets E_L to be suppressed by a factor of 100 relative to $E_{J,\text{arr}}$. These energy scales generally yield qubit frequencies in the range of 50–300 MHz, depending on the precise fabricated values. In practice, the majority of our devices have exhibited qubit frequencies in the 100–200 MHz range. It is important to note that E_C consistently decreased from its target value during the fabrication, typically settling around 0.9 GHz, attributed to over-etching of the capacitor pads. This systematic reduction in E_C shifted the resulting qubit frequency away from the intended target, consequently requiring a compensatory adjustment of the junction energies E_J and E_L.

E_C/h (GHz)	E_J/h (GHz)	E_L/h (GHz)
1	4–5	0.5–0.9

Table 4.2: Energy parameters for the fluxonium design.

SINGLE JUNCTION

The single Josephson junction desired E_J in our devices lies in the range of 4–5 GHz. Using the established fabrication recipe in our lab, achieving this E_J range requires junction normal-state resistances of the order of 30–40 kΩ, corresponding to junction areas of approximately $(100\text{–}150\,\text{nm})^2$, as illustrated in Fig. 4.3(b).

Junctions in this size regime are inherently challenging to fabricate with high accuracy, as reproducibility and targeting are limited by lithographic resolution and process variability. This is quantified in Fig. 4.5, which plots the coefficient of variation (C_V) for junction resistance as a function of junction width. The C_V data clearly demonstrates that smaller junctions are increasingly difficult to target accurately, directly translating to greater uncertainty in the achieved E_J. An additional constraint arises from the fact that the junction array is deposited in the same lithographic step as the single junction, imposing more stringent requirements on the lithography and deposition parameters, as discussed in Section 4.4.1.

ARRAY JUNCTIONS

The inductor array consists of N Josephson junctions connected in series, each with Josephson energy $E_{J,\mathrm{arr}}$. In the limit of sufficiently large N, this chain approximates a linear inductor with an effective inductive energy $E_L \approx E_{J,\mathrm{arr}}/N$ [16]. In our design, we employ $N = 100$ junctions, each with a relatively large $E_{J,\mathrm{arr}} >$ 75 GHz, corresponding to junction dimensions exceeding 350 nm. Operating in this larger junction size regime offers a significant fabrication advantage: larger junctions are considerably easier to target reproducibly, resulting in better control over E_L compared to the single junction. This deliberate asymmetry in junction sizes between the single junction and the array makes the fabrication difficult.

4.4.1. FABRICATION OF THE SINGLE JUNCTION AND ARRAY

The required E_J and E_L values for the fluxonium qubit directly constrain the physical junction widths used in fabrication. For the single Josephson junction, the target width lies in the range of 100–160 nm, while the individual array junctions are designed to be considerably larger, in the range of 350–450 nm. These size constraints in turn determine the resist stack geometry required for double-angle shadow evaporation using the Manhattan technique, as illustrated in Fig. 4.4.

The geometric constraints governing the first deposition direction are derived from the cross-sections shown in Fig. 4.4(e) and (f). To prevent direct aluminum contact between the deposited film and the substrate edges, the total resist height H_t must satisfy

$$h_t = w \tan\theta, \tag{4.7}$$

with the additional requirement that

$$h_t < h_{\mathrm{PMMA}}, \tag{4.8}$$

where h_{PMMA} is the PMMA layer thickness, ideal recommendation is $h_{\mathrm{PMMA}} = 2 * h_t$, for more stable resist during deposition and lift-off of aluminium [11]. From the cross-sections in Fig. 4.4(c) and (d), the undercut width U at the junction end is given by

$$U = h_1 \cot\theta. \tag{4.9}$$

To ensure a clean junction overlap without shorts, the designed end clearance d must satisfy $d > U$.

An additional complication arises from the descum ashing step, which etches the resist isotropically in all directions. As shown in Fig. 4.4(h), this modifies

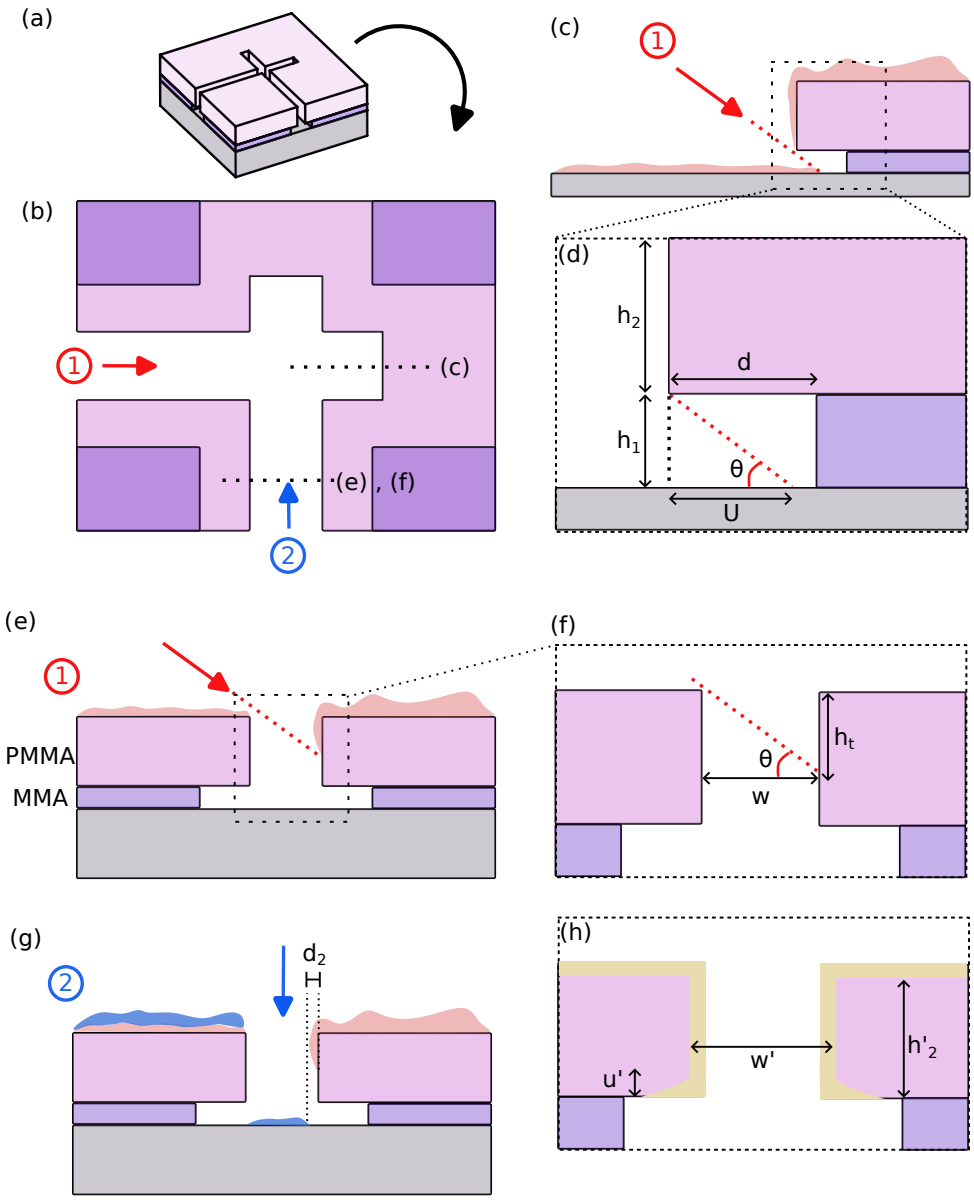

Figure 4.4: Geometry of JJ deposition: (a) Three-dimensional schematic of the bilayer electron-beam resist (PMMA/MMA) geometry prior to metal deposition. (b) Top-down view of the resist mask; colored arrows indicate the directional orientation for the two evaporation steps: red (1) denotes the first deposition at an angle θ, and blue (2) denotes the second deposition. (c) Cross-sectional side profile of the end section as indicated by the dashed line in (b), showing the undercut profile. (d) Magnified view of the undercut geometry from (c), defining the resist heights (h_1, h_2), the evaporation angle θ, and the undercut width U. (e) Side profile during the first aluminum deposition step, where the shadow of the PMMA bridge defines the electrode footprint. (f) Zoomed-in view of (e) illustrating the relationship between the aperture width w and the effective deposition area. (g) Side profile after the second deposition step (blue); the overlap between the first oxidized layer and the second layer forms the Josephson tunnel junction. (h) Cross-sectional view of the resist profile following an ashing process, showing the resulting alterations in the undercut dimensions (u') and width (w').

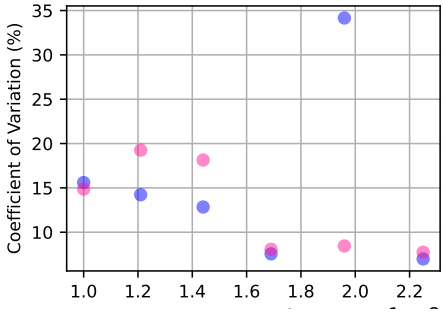

Figure 4.5: Coefficient of variation: The plot illustrates the relative dispersion of Resistance for different sized junctions

the effective junction width relative to the designed pattern, reduces the top resist height h_2 to h_2', widens the undercut to u', and introduces curved sidewalls beneath the top resist layer. Collectively, these geometrical changes allow aluminum to penetrate further into the undercut than predicted by the nominal geometry, directly affecting the realized junction dimensions. Accurate undercut calculations must therefore account for these ashing-induced offsets, and a dedicated dose test is required to calibrate the extent of resist erosion and its impact on the final junction geometry.

4.4.2. HOW TO IMPROVE THE JUNCTIONS FOR THE FLUXONIUMS DEVICES?

During characterization of the fluxonium devices, several recurring signatures of non-optimal device performance have been identified. In qubit measurements, these manifest as follows:

1. **Frequency targeting:** Measured qubit frequencies are consistently lower than designed. This occurs when the fabricated E_J of the single junction exceeds the target value, pushing the device deeper into the heavy fluxonium regime for a fixed E_C. In terms of resistance, this corresponds to a junction normal-state resistance lower than intended. This is a systematic challenge for junctions in the smaller dimension regime, where the larger coefficient of variation C_V leads to reduced targeting accuracy, as discussed above.

2. **Frequency targeting in multi-qubit devices:** Precise frequency targeting becomes considerably more demanding in multi-qubit devices. In a multi-qubit architecture , each qubit must be placed at its designated spectral operating point to enable high-fidelity gate operations. Any deviation in

E_J across the junctions directly compromises the ability to execute the intended gate protocols. This challenge is discussed in further detail in next chapter.

3. **Limited coherence times (T_1 and T_2):** Over the course of the project, the measured T_1 and T_2 times have remained below 150 μs and 50 μs, respectively, across the majority of devices. While the precise mechanisms limiting coherence are difficult to isolate, the consistently moderate average values of T_1 and T_2 across devices suggest that systematic, fabrication-related loss channels are likely contributing.

The underlying causes of these observations are not straightforward to attribute to individual fabrication steps. Much like the dependence of resonator quality factors on process parameters, the qubit performance metrics are a compound result of multiple interdependent fabrication variables, making it difficult to isolate the contribution of any single step. The following sections therefore examine in detail the ways in which junction fabrication can deviate from the optimum and the mechanisms by which such deviations translate into degraded device performance.

To address the two categories of problems identified above, it is necessary to simultaneously reduce both the C_V and the systematic offset in junction resistance targeting, and to understand the microscopic origins of junction-limited coherence. Disentangling these factors requires a careful investigation of the junction growth process, as the quality of the tunnel barrier and the surrounding oxide environment are expected to play a central role in determining both the junction resistance reproducibility and the coherence properties of the qubit.

INFLUENCE OF THE GROWTH PROCESS ON THE MICROSTRUCTURE OF THE JUNCTION:

How does the microstructure of the JJ affects the R_N? The fabrication control of Josephson junctions can be broadly divided into two dominant mechanism that govern the behavior and uniformity of the junction: one related to the area control of the junction electrodes, and the other involving the more intricate growth processes at the microscopic level within the SIS stack.

The first factor pertains to the structural definition of the junction area, which can be engineered through precise lithographic patterning of the electrodes. If lithographic and associated processing tolerances are tightly controlled, features within the 5–10 nm scale can be defined with reasonable accuracy. The electrode structure, though susceptible to some inevitable variations stemming from resist behavior, etch non-uniformity, and proximity effects, remains relatively

more controllable and is thus considered the engineering limit of junction definition[17, 18].

The second factor addresses the complex growth behavior of the SIS stack, which encompasses three critical regions: (1) the bottom aluminum electrode, (2) the intermediate oxide barrier, and (3) the top aluminum electrode. The surface morphology of the bottom electrode is influenced by the underlying silicon topography, which inherits roughness from dry etching. Subsequently, the angle of metal deposition and the deposition rate—typically fluctuating around 10 Å/s with an uncertainty of ±1.5 Å/s (due to PID-based feedback loops), introduce further structural variability [19]. In addition, physical vapor deposition (PVD) conditions, such as those governed by the electron beam evaporation source, impact the charge distribution and thermal radiation incident on the molten aluminum [20]. These factors modulate adatom mobility and nucleation behavior, both of which critically shape the microstructure of the deposited film. Notably, in the Plassys used for deposition, the growth environment is highly subjective to temporal fluctuations due to frequent material consumption and source refills, further complicating reproducibility in growth.

Additionally, one more factor which deteriorates the growth of bottom electrode is the presence of any organic residues from the pre-processing steps of the fabrication flow[21]. This will inevitably lead to the aluminum grains having a polymer influenced growth and degraded contact of the junction with the silicon underneath. This can be spatially varying leading to different properties of spatially located junctions.

Once the bottom electrode is deposited, its surface properties directly influence the subsequent formation of the tunnel barrier. The aluminum layer undergoes oxidation in an oxygen-rich environment where the oxygen adatoms migrate and settle onto the aluminum surface [22, 23, 24]. The degree and nature of oxidation—whether static or dynamic—along with pressure and exposure time, dictate the resultant stoichiometry and structural disorder of the oxide. The barrier is generally amorphous and disordered, and is best described as AlO_x rather than stoichiometric Al_2O_3, due to variable coordination environments and oxygen content[23, 25]. This oxide layer, in turn, serves as the growth surface for the top aluminum electrode, which nucleates based on the seed properties of the oxide. As such, the crystallization and grain structure of the top electrode are indirectly inherited from the characteristics of the oxide [26, 27].

Critically, the tunneling behavior in Josephson junctions is governed not by the average barrier thickness but rather by its thinnest regions. Empirical observations suggest that the majority of electron tunneling occurs through roughly 10% of the barrier area, where the oxide is at its minimum thickness[25, 26]. As such, although variations in deposition rates during the bottom and top elec-

trode growth (steps 1 and 3) may affect the overall morphology, they do not significantly alter the critical current, since the thin regions remain largely consistent. In contrast, the oxidation process (step 2) plays a dominant role in defining junction resistance. It determines the extent of oxygen incorporation, adatom mobility, and coordination complexity within the barrier, thereby exerting the greatest influence on electrical performance.

Thus, while junction behavior is often assessed via averaged quantities such as resistance and critical current, it is essential to recognize that these reflect the characteristics of the most conductive paths within the barrier. Understanding and engineering the growth parameters, during bottom electrode growth and oxidation, remains critical for reliable and reproducible JJ fabrication.

4.5. RESULTS

In this section we discuss the improvements made to the junction fabrication recipes, leading to more precise resistance targeting and cleaner fabrication processes. Additionally, a post-fabrication annealing setup was developed using the voltage-bias annealing method introduced in [28].

4.5.1. MEASUREMENT OFFSETS FROM SHUNT RESISTANCE:

During initial calibration measurements, for calculating the resistance of the junction, the role of the substrate was not taken into account. Since the substrate is Hi-Res silicon and has a photo-conductive effect, the resistance of the substrate lowers. This creates a path for the current to travel in parallel to the junction structures. In Fig. 4.6 (b) n junctions, each with a resistance (R_J) are connected in series, the combined resistance of the circuit can be given by-0

$$R_{AB} = \left(\frac{1}{R_{Sub.}} + \frac{1}{(R_J^1 + R_J^2 + ...R_J^n)}\right)^{-1} \tag{4.10}$$

As the number of junctions increase in the array, the total resistance of the series component increase. If the $R_{sub.}$ is smaller to this series resistance, the total resistance measured saturates. This is specially the case when the substrate is illuminated by photons, leading to deceased c, this effect is seen with the measured resistances in blue, for different number of junctions, the resistance saturates. When the LED is off, the $R_{sub.}$ is very high, leading to a expected linear relation between the resistance and the number of junction in array.

The same effect also creates an offset in the resistance measurements of the single junctions as well. When the resistance of the small junctions reaches values comparable to the one using LED offset, it will saturate the value of the single junction resistance. This initially led to offsets in E_J targeting for the fluxonium targeting.

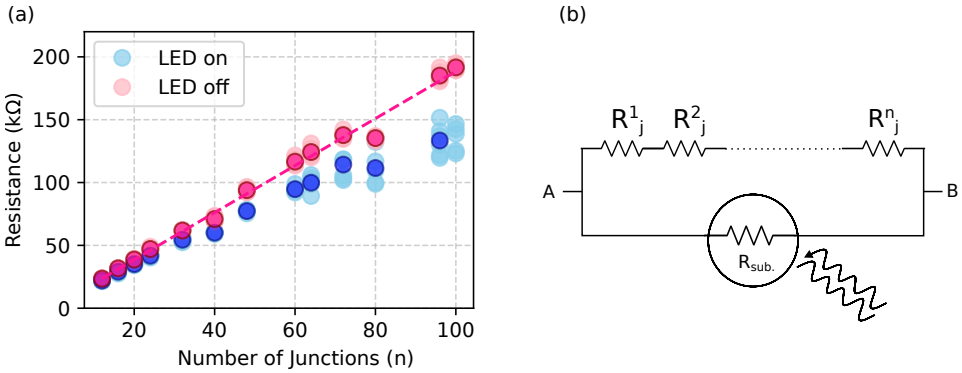

Figure 4.6: Effect of ambient lighting: (a) trend of resistance of n number of junctions in an array with increasing number of junctions. (b) Circuit diagram with the resistances of array junctions in series, parallel to the substrate resistance.

4.5.2. STACK ENGINEERING WITH LITHOGRAPHIC IMPROVEMENTS:

To realize the fabrication of the fluxonium using manhattan style (sec 4.4.1), three iterations of the multi layer resist stacks were optimized. Each one of the stack had a different combination of the PMMA and MMA based resists, targeting slightly different height of the resists and width for the junctions. The stacks are listed here with further details-

Process Step	**Stack A**	**Stack B**
Spinning	MMA EL8 @ 4000 rpm	MMA EL8 @ 3000 rpm
	PMMA 950 A4 @ 2000 rpm	PMMA 495 A8 @ 1000 rpm
		PMMA 950 A3 @ 3000 rpm
Baking	5 min @ 185°C each layer	10 min @ 180°C each layer
	2 min break between layers	2 min break between layers
Writing	Dose: 1250$\mu C/mm^2$	Dose: 175$\mu C/mm^2$
	Beam aperture: 200μm	Beam aperture: 400μm
	Beam current: 98 pA	Beam current: 1 nA
	Spot size: 2.1 nm	Spot size: 3.6 nm
Development	H_2O:IPA (1:3) @ 6 °C for 90 s	H_2O:IPA (1:3) @ 21 °C for 3 min 10 s
	Rinse: 30 s IPA	Rinse: 10 s IPA, then 1 min IPA
	Manual stirring	Sonnication

Table 4.3: Comparison of processing parameters for Stack A and Stack B.

STACK A

The design philosophy behind resist stack A was to use a bilayer resist stack consisting of 300 nm MMA as the bottom layer and 320 nm PMMA 950K as the top layer. As shown in Fig. 4.7(a), deposition is performed at an angle of 35°, satisfying the geometric undercut condition required to fabricate fluxonium junctions. An important assumption in using this stack was that exposing it at a sufficiently high dose would produce a well-developed undercut profile after development. However, subsequent SEM cross-section characterisation of the developed resist revealed that the expected undercut did not form. As shown in Fig. 4.7(c), the cross-section image shows a distorted profile in the middle region, which is attributed to prolonged electron-beam heating of the resist during SEM imaging. Rather than an undercut, the resist profile resembles a trench geometry, leaving insufficient clearance between the resist sidewall and the deposited aluminium film.

Since junction fabrication relies on lift-off, this reduced clearance between the aluminium and the resist degrades the quality of the resulting junction. As shown in Fig. 4.7(e), resist residues are visible across the junction area following lift-off. Furthermore, the proximity of the resist to the substrate influences where the deposited aluminium lands on the silicon surface, resulting in the formation of residual standing aluminium features, also visible in the SEM image.

STACK B

To improve junction cleanliness and eliminate resist residues, several targeted modifications were made to the deposition process over successive fabrication iterations, culminating in an optimised process referred to as Stack B. This recipe employs a trilayer resist stack comprising MMA (300 nm) / PMMA 495k (1 μm) / PMMA 950k (150 nm), with a double-angle shadow evaporation at 40°. The key changes relative to Stack A are described below.

1. **Undercut profile.** A continuous all-around undercut was introduced on all sides of the aluminium electrodes. The trilayer resist stack was chosen to produce an inward, wedge-shaped trench profile, providing greater clearance between the deposited aluminium and the resist sidewalls. In addition, a low-dose MMA exposure step was incorporated to generate a more pronounced inverted-T cross-sectional profile at the junction. The combined effect of the wedge-shaped undercut and the inverted-T geometry is illustrated in Fig. 4.7(d).

2. **Ashing.** Prior to the Stack B optimisation, the descum step was performed using a Tepla oxygen plasma system. The cylindrical geometry of the Tepla reactor produces a radially isotropic ashing profile, which can result in

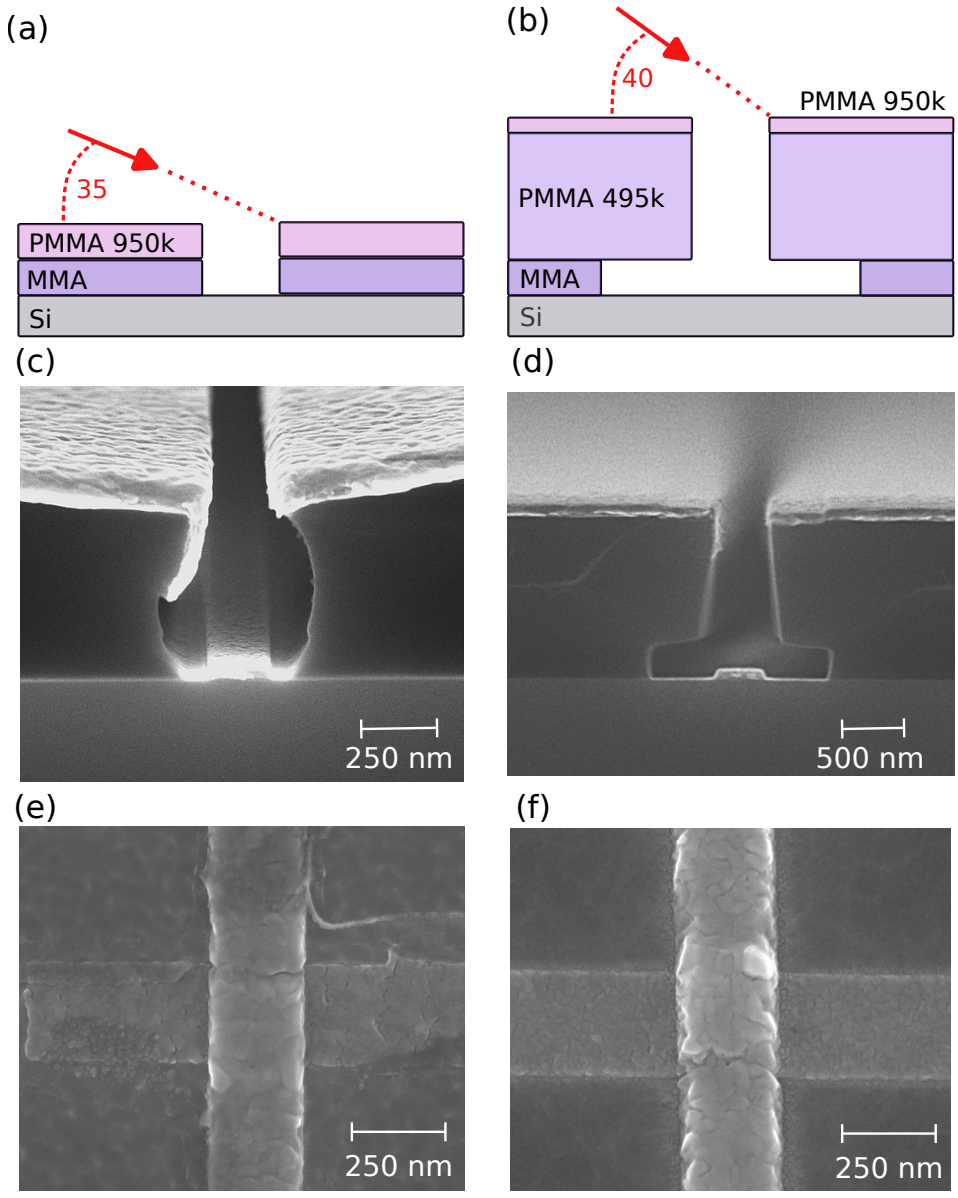

Figure 4.7: Illustration of undercut engineering: (a),(b) Two layer and three layer resist stack for junction fabrication respectively. (c),(d) Cross section SEMs of the resist profile after aluminium deposition, correspondiong to the two layer and three layer stack respectively. (e),(f) Top view SEMs of junction after lift-off, correspondiong to the two layer and three layer stack respectively.

non-uniform resist erosion and sub-optimal descumming [29] . To achieve a more directional and controlled oxygen flow, the process was transferred to an F1 Sentech plasma system, which provides improved uniformity and directionality of the oxygen plasma.

3. **Development.** Compared to Stack A, the development step was modified to use a room-temperature H_2O : IPA solution applied in conjunction with sonication. This replaces manual agitation, eliminating operator-dependent variability in the development process. The sonication additionally promotes more complete dissolution of the exposed resist, yielding a cleaner silicon surface in the junction area. It is important, however, to place the sample horizontally inside the sonication bath and to vary its position periodically, so as to avoid standing-wave hot-spots associated with the sonication field.

4. **Lift-off.** In Stack A, lift-off was performed by immersion in NMP for a minimum of 2 hours. However, this was found to leave residual resist on the chip surface in some cases. In Stack B, a more thorough lift-off procedure was introduced, consisting of sequential rinses in heated NMP baths to ensure complete removal of resist and produce a cleaner final surface.

The combined effect of these modifications in the Stack B recipe results in significantly cleaner junctions. As shown in the top-view SEM image in Fig. 4.7(f), the junctions fabricated with Stack B exhibit no visible resist residues, no spurious standing aluminium structures, and a cleaner aluminium grain morphology compared to Stack A.

Despite these improvements, the thicker Stack B resist profile introduces several residual challenges:

1. The junction width offsets are larger than intended due to the increased resist thickness, resulting in a deviation of the fabricated junction dimensions from the designed values.

2. Accurate resistance targeting becomes more difficult for smaller junctions, which exhibit a proportionally larger width offset. This was partially mitigated by using a higher oxidation pressure during the barrier formation step, which increases the tunnel barrier resistance and partially compensates for the enlarged junction area.

3. The coefficient of variation C_V of the junction resistance remains elevated for smaller junction dimensions, as shown in Fig. 4.5(b), consistent with the general trend of increasing targeting difficulty at reduced junction sizes discussed in Section 4.4.

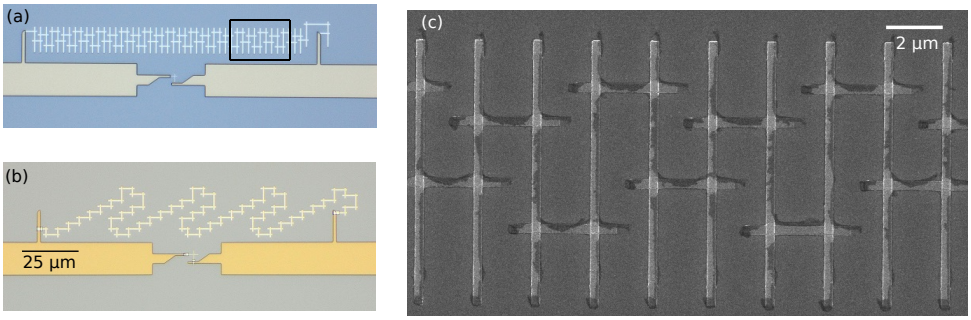

Figure 4.8: (a) Microscope image of a Josephson junction array consisting of 100 junctions connected in series. (b) Scanning Electron Microscopy (SEM) image of the junction array. (c) Microscope image of the staircase junction array

4.5.3. ARRAY IMPROVEMENTS

LAYOUT

The junction array for the fluxonium inductor was initially designed in a ladder configuration, as shown in Fig. 4.8(a), where the junctions are arranged in a staggered pattern with uniform junction dimensions. Undercuts were defined at the ends of the junction leads to prevent aluminium build-up during lift-off. While functional, this layout resulted in arrays containing significant resist residues, as visible in Fig. 4.8(c), motivating a redesign of both the fabrication process and the array geometry.

The revised fabrication conditions require an all-around undercut profile on all sides of the aluminium electrodes, which imposes a minimum separation of 1 μm between any two adjacent undercut regions in order to maintain the mechanical stability of the resist structures during processing. In the ladder layout, the junction leads are too closely spaced to satisfy this constraint, necessitating a new array geometry. The redesigned layout spaces the junction leads further apart, resulting in the staircase-type geometry shown in Fig. 4.8(b). This new layout was developed subject to two simultaneous constraints: (1) maintaining the required minimum separation between all adjacent undercut regions, and (2) preserving the total footprint length of the array to remain comparable to that of the original ladder design.

JUNCTION PARAMETERS

Two mechanisms limit the dephasing time T_2 of fluxonium qubits through the junction array. First, the total perimeter of the array loops and the width of the junction leads influence T_2 through $1/f$ flux noise, which arises from magnetic spin defects at the metal–oxide interfaces surrounding the superconducting wires [30]. Wider leads and a smaller total array perimeter directly reduce the

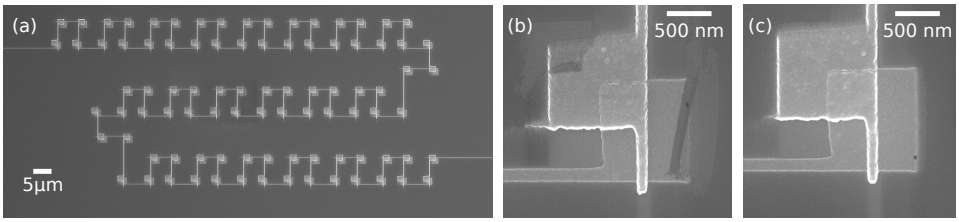

Figure 4.9: New junction array layout (a) SEM image of the array layout showing the staggered junction arrangement. (b) SEM image of individual junction showing the overlap geometry; resist residues are visible due to a non-optimised undercut profile. (c) Improved junction with a cleaner surroundings following undercut optimisation.

flux noise amplitude and thereby improve T_2.

Second, coherent quantum phase slips (CQPS) in the array provide an additional dephasing channel [31]. Phase slips can occur at each individual array junction at a rate that increases with the junction impedance. These phase-slip events coherently interfere via an Aharonov–Casher phase that depends on the offset charges of the array islands, perturbatively modifying the qubit transition frequency. Fluctuations of these offset charges, driven by charge-parity switching within the array, broaden the qubit transition and limit T_2. The single-junction phase-slip amplitude scales as

$$\gamma_{\mathrm{cqps}} \propto \exp\left(-\sqrt{8\,E_{J,\mathrm{arr}}/E_{C,\mathrm{arr}}}\right), \tag{4.11}$$

where $E_{J,\mathrm{arr}}$ and $E_{C,\mathrm{arr}} = e^2/(2C_{\mathrm{arr}})$ are the Josephson and charging energies of a single array junction, respectively. Since $E_{J,\mathrm{arr}} \propto a_A$ and $E_{C,\mathrm{arr}} \propto 1/a_A$, where a_A is the junction area, increasing a_A exponentially suppresses the phase-slip amplitude and therefore exponentially improves the CQPS-limited T_2.

Motivated by these two constraints, we developed a new array junction layout, shown in Fig. 4.9, designed to address both decoherence channels simultaneously. The layout, shown schematically in Fig. 4.9(a), uses staggered square overlap junctions as visible in Fig. 4.9(b), with Fig. 4.9(c) showing a cleaner realisation following undercut optimisation. The key design considerations are as follows: the junction leads are optimised to reduce the total array perimeter while satisfying the minimum spacing requirements imposed by the resist undercut geometry; the individual junction area a_A is increased relative to the previous design to exponentially suppress the CQPS rate; and the total enclosed loop area is kept as small as possible to minimise flux noise pickup and maintain good flux loop control. This layout was implemented using the Manhattan junction process to first establish a working baseline and understand the dominant mechanisms before further optimisation. Extensive measurements on devices

with this new array design are ongoing and will be needed to fully characterise the improvements to coherence.

4.5.4. POST-FAB ANNEALING

During extensive fabrication runs across multiple device generations, a significant fraction of devices could not be targeted to the desired single junction E_J. This was attributed to run-to-run variability in the Plassys deposition system — whether from drift in deposition conditions or other process parameters that could not always be tracked or identified in real time. Since such deviations directly shift the qubit frequency away from the intended operating point, the ability to tune the junction resistance after fabrication became a strong motivation for exploring post-fabrication resistance adjustment techniques.

The approach adopted here is based on thermal annealing of the junction tunnel barrier [28]. The key insight enabling selective tuning of the single junction is the large resistance asymmetry between the two junction types in the fluxonium circuit: the single junction has a significantly lower resistance than the array junctions, and consequently carries the majority of the current when a bias is applied across the circuit. This means that the annealing effect is localised predominantly at the single junction, leaving the resistance of the array junctions largely unaffected [28]. This selectivity is particularly valuable for the fluxonium, since independent post-fabrication control of the single junction E_J — without disturbing E_L — allows in-principle correction of frequency targeting errors without requiring a full fabrication re-run.

To investigate this approach, a dedicated experimental setup was developed as discussed in Appendix 4.6. Initial measurements confirmed a measurable change in the single junction resistance following annealing. While these preliminary results are encouraging, a more systematic study is needed to fully characterise the dependence of the resistance shift on annealing temperature, duration, and the initial junction parameters.

4.6. CONCLUSION AND OUTLOOK

This chapter has described the fabrication, characterisation, and optimisation of Josephson junctions for fluxonium qubits, covering the progression from initial process development through to targeted improvements motivated by coherence considerations.

The fabrication of Al/AlO_x/Al SIS junctions using the Manhattan double-angle shadow evaporation technique was established as the core process. The room-temperature normal-state resistance R_N was shown to provide a reliable proxy for the Josephson energy E_J via the Ambegaokar–Baratoff relation, enabling rapid characterisation without cryogenic measurements. An important

early finding was the role of substrate photo-conductivity in introducing systematic offsets in resistance measurements, which, once identified and controlled, significantly improved the accuracy of E_J targeting.

The resist stack engineering effort produced two distinct process generations. Stack A, based on a bilayer resist, was found to produce inadequate undercut profiles and resulted in junctions contaminated with resist residues. Stack B, employing a trilayer resist with all-around undercut geometry, combined with improvements to the ashing system, development process, and lift-off procedure, yielded significantly cleaner junctions. The trade-off introduced by the thicker Stack B resist — namely, larger junction width offsets and elevated C_V for small junctions — was partially compensated by increasing the oxidation pressure, though targeting accuracy for the smallest junction dimensions remains a fundamental challenge.

The junction array layout was redesigned from the original ladder geometry to a staircase configuration to accommodate the all-around undercut requirements of Stack B, while preserving the total array footprint. More recently, a further redesign of the array junction shape and layout was motivated by two identified decoherence channels: $1/f$ flux noise and coherent quantum phase slip (CQPS) dephasing. The new layout targets both simultaneously by increasing a_A, widening the junction leads, and minimising the enclosed loop area, while remaining within the constraints imposed by the resist undercut geometry. Initial fabrication runs using this design have been completed, though extensive device-level measurements are still required to quantify the coherence improvements.

Finally, a post-fabrication resistance tuning technique based on voltage-bias annealing [28] was explored as a route to correct E_J targeting errors that arise from untracked process variability. Initial measurements confirmed a measurable resistance shift, establishing proof of concept for this approach.

OUTLOOK

On the fabrication side, the C_V of the single junction resistance at the 100–150 nm scale remains the primary limiting factor for frequency targeting accuracy, and further improvements in lithographic dose control, ashing calibration, and undercut uniformity are needed to reduce this variability. A systematic long term control study of effects from lithography and deposition based processes on junction formation is necessary to minimize errors.

For the junction array, the new staggered overlap geometry requires full device-level characterisation, including T_1 and T_2 measurements as a function of junction area and lead width, to experimentally validate the predicted suppression of both flux noise and CQPS dephasing.

For post-fabrication annealing, a systematic study varying annealing voltage, pulse duration, and initial junction resistance is needed to establish a reliable and predictable tuning protocol. Understanding whether the annealing-induced resistance change is stable over time and that it only targets the single junction is necessary to validate the use case post fabrication. If demonstrated to be reliable, voltage-bias annealing could significantly improve the yield of frequency-targeted multi-qubit fluxonium devices, reducing the number of fabrication iterations required to produce devices with the correct spectral configuration.

APPENDIX:A

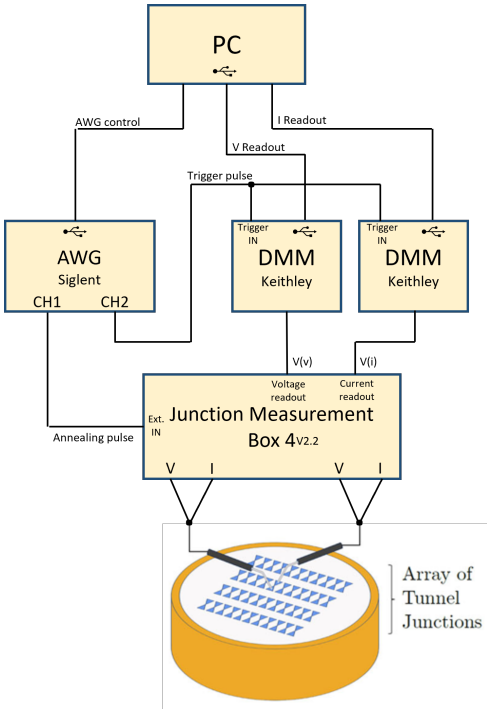

Figure 4.10: Schematic of the room-temperature setup used for resistance measurement and post-fabrication annealing of Josephson junctions. A PC controls a Siglent AWG and two Keithley DMMs via USB. The AWG applies voltage or current pulses through a custom Junction Measurement Box (v2.2), which implements a four-wire resistance measurement via dedicated voltage and current probe lines. The sample is contacted directly by probe needles on a chip carrier.

DATA AVAILABILITY

The data for the experiments along with the scripts to anlayse the data can be found at [32] with an open CC by 4.0 license.

REFERENCES

[1] B. D. Josephson. “Possible new effects in superconductive tunnelling”. In: *Physics Letters* 1.7 (1962), pp. 251–253. DOI: 10.1016/0031-9163(62)91369-0.

[2] N. Cabrera and N. F. Mott. “Theory of the oxidation of metals”. In: *Reports on Progress in Physics* 12.1 (1948), p. 163.

[3] C Deng, MR Otto, and A Lupascu. "Characterization of quantum-regime dielectric loss of aluminum oxide using superconducting LC resonators". In: *Applied Physics Letters* 104.4 (2014), p. 043506.

[4] B. Abdisatarov et al. *Direct Measurement of Microwave Loss in Nb Films for Superconducting Qubits*. 2024. arXiv: 2407.08856 [cond-mat.supr-con]. URL: https://arxiv.org/abs/2407.08856.

[5] Vinay Ambegaokar and Alexis Baratoff. "Tunneling Between Superconductors". In: *Physical Review Letters* 10.11 (June 1963), pp. 486–489. DOI: 10.1103/PhysRevLett.10.486. URL: https://link.aps.org/doi/10.1103/PhysRevLett.10.486.

[6] G. J. Dolan. "Offset masks for lift-off photoprocessing". In: *Applied Physics Letters* 31.5 (1977), pp. 337–339. DOI: 10.1063/1.89690.

[7] A. Potts et al. "CMOS compatible fabrication methods for submicron Josephson junction qubits". In: *IEE Proceedings-Science, Measurement and Technology* 148.5 (2001), pp. 225–228. DOI: 10.1049/ip-smt:20010603.

[8] Dante Colao Zanuz et al. "Mitigating losses of superconducting qubits strongly coupled to defect modes". In: *Phys. Rev. Appl.* 23 (4 Apr. 2025), p. 044054. DOI: 10.1103/PhysRevApplied.23.044054. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.23.044054.

[9] A. Dunsworth et al. "Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits". In: *Applied Physics Letters* 111.2 (2017), p. 022601. DOI: 10.1063/1.4993577.

[10] Ani Nersisyan et al. "Manufacturing low dissipation superconducting quantum processors". In: *Physical Review Applied* 11 (2019), p. 044030. DOI: 10.1103/PhysRevApplied.11.044030.

[11] Nandini Muthusubramanian. "Qubit frequency targeting in scalable superconducting quantum processors". PhD thesis. Delft, Netherlands: Delft University of Technology, 2024. DOI: 10.4233/uuid:b2b1f6e0-8c4a-4e4a-9e4a-8c4a4e4a9e4a. URL: https://research.tudelft.nl/en/publications/qubit-frequency-targeting-in-scalable-superconducting-quantum-pro.

[12] M. Kaden, S. Schulze, and D. Drung. *Systematic Characterization of Josephson Junction Parameters: Separating Barrier Height and Thickness Variations*. 2025. arXiv: 2502.13072 [cond-mat.supr-con]. URL: https://arxiv.org/abs/2502.13072.

[13] Figen Yilmaz et al. “Energy-participation-ratio analysis for very anharmonic superconducting circuits”. In: *Physical Review Applied* (2025). DOI: 10.1103/1rbn-c4xf. arXiv: 2411.15039.

[14] Taryn V. Stefanski et al. “Improved fluxonium readout through dynamic flux pulsing”. In: *APL Quantum* 2.2 (2025), p. 026103. DOI: 10.1063/5.0255892.

[15] Siddharth Singh et al. “Fast microwave-driven two-qubit gates between fluxonium qubits with a transmon coupler”. In: *Physical Review Applied* 25 (2026), p. 024020. DOI: 10.1103/yxf3-jtx5. arXiv: 2504.13718.

[16] Vladimir E Manucharyan. “Superinductance and the Fluxonium Qubit”. PhD thesis. New Haven, CT: Yale University, 2012. URL: https://rsl.yale.edu/sites/default/files/files/vladimir-manucharyan-thesis.pdf.

[17] Dmitry O. Moskalev et al. “Optimization of shadow evaporation and oxidation for reproducible quantum Josephson junction circuits”. In: *Scientific Reports* 13 (2023), p. 4174. DOI: 10.1038/s41598-023-31003-1.

[18] Anastasiya A. Pishchimova et al. “Improving Josephson junction reproducibility for superconducting quantum circuits: Junction area fluctuation”. In: *Scientific Reports* 13 (2023), p. 6772. DOI: 10.1038/s41598-023-34051-9.

[19] Jin-Su Oh et al. “Correlating aluminum layer deposition rates, Josephson junction microstructure, and superconducting qubits’ performance”. In: *Acta Materialia* 284 (2025), p. 120631. DOI: 10.1016/j.actamat.2024.120631.

[20] Frank Volmer et al. “How to solve problems in micro- and nanofabrication caused by the emission of electrons and charged metal atoms during e-beam evaporation”. In: *Journal of Physics D: Applied Physics* 54.22 (Mar. 2021), p. 225301. DOI: 10.1088/1361-6463/abe89b. URL: https://doi.org/10.1088/1361-6463/abe89b.

[21] A. Savenko et al. “On-chip stencil lithography for superconducting qubits”. In: *arXiv preprint* (2025). arXiv: 2507.17005.

[22] S. Fritz et al. “Correlating the nanostructure of Al-oxide with deposition conditions and dielectric contributions of two-level systems in perspective of superconducting quantum circuits”. In: *Scientific Reports* 8 (2018), p. 7956. DOI: 10.1038/s41598-018-26066-4.

[23] S. Fritz et al. “Structural and nanochemical properties of AlO_x layers in $Al/AlO_x/Al$-layer systems for Josephson junctions”. In: *Physical Review Materials* 3 (2019), p. 114805. DOI: 10.1103/PhysRevMaterials.3.114805.

[24] Timothy C. DuBois et al. “Constructing *ab initio* models of ultra-thin Al–AlO_x–Al barriers”. In: *Molecular Simulation* 42.7 (2016), pp. 542–548. DOI: 10.1080/08927022.2015.1068941.

[25] Lunjie Zeng et al. “Atomic structure and oxygen deficiency of the ultrathin aluminium oxide barrier in $Al/AlO_x/Al$ Josephson junctions”. In: *Scientific Reports* 6 (2016), p. 29679. DOI: 10.1038/srep29679.

[26] M. J. Cyster et al. “Effect of atomic structure on the electrical response of aluminum oxide tunnel junctions”. In: *Physical Review Research* 2 (2020), p. 013110. DOI: 10.1103/PhysRevResearch.2.013110.

[27] M. J. Cyster et al. “Simulating the fabrication of aluminium oxide tunnel junctions”. In: *npj Quantum Information* 7 (2021), p. 12. DOI: 10.1038/s41534-020-00360-4.

[28] David P. Pappas et al. “Alternating-bias assisted annealing of amorphous oxide tunnel junctions”. In: *Communications Materials* 5 (2024), p. 150. DOI: 10.1038/s43246-024-00596-z.

[29] John Mark Kreikebaum. “Superconducting Qubit Enabled Single Microwave Photon Detection”. PhD thesis. University of California, Berkeley, 2020. URL: https://escholarship.org/uc/item/0834k8bw.

[30] Jochen Braumüller et al. “Characterizing Next-Generation Superconducting Qubits via Coherence-Benchmarking”. In: *Physical Review Applied* 13 (5 May 2020), p. 054079. DOI: 10.1103/PhysRevApplied.13.054079. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.13.054079.

[31] Mallika T. Randeria et al. “Mitigation of Two-Level System Loss in Tantalum-Based Superconducting Qubits”. In: *PRX Quantum* 5 (3 Sept. 2024), p. 030341. DOI: 10.1103/PRXQuantum.5.030341. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.030341.

[32] *Data for the thesis titled: Engineering Noise-Resilient Superconducting Qubits: Fabrication and control of fluxonium and cos(2phi) systems*. https://doi.org/45267100-d50e-4ed5-a91f-53e4a899238f. 2026.

5

FAST MICROWAVE-DRIVEN TWO-QUBIT GATES BETWEEN FLUXONIUM QUBITS WITH A TRANSMON COUPLER

Two qubit gates constitute fundamental building blocks in the realization of large-scale quantum devices. Using superconducting circuits, two-qubit gates have previously been implemented in different ways with each method aiming to maximize gate fidelity. Another important goal of a new gate scheme is to minimize the complexity of gate calibration. In this work, we demonstrate a high-fidelity two-qubit gate between two fluxonium qubits enabled by an intermediate capacitively coupled transmon. The coupling strengths between the qubits and the coupler are designed to minimize residual crosstalk while still allowing for fast gate operations. The gate is based on frequency selectively exciting the coupler using a microwave drive to complete a 2π rotation, conditional on the state of the fluxonium qubits. When successful, this drive scheme implements a conditional phase gate. Using analytically derived pulse shapes, we minimize unwanted excitations of the coupler and obtain gate errors of 10^{-2} for gate times below 60 ns. At longer durations, our gate is limited by relaxation of the coupler. Our results show how carefully designed control pulses can speed up frequency selective entangling gates.

This chapter has been published as: S. Singh, E. Y. Huang, J. Hu, F. Yilmaz, M. F. S. Zwanenburg, P. Kumaravadivel, S. Wang, T. V. Stefanski, and C. K. Andersen, *Fast microwave-driven two-qubit gates between fluxonium qubits with a transmon coupler*, Physical Review Applied **25**, 024020 (2026).

5.1. INTRODUCTION

Superconducting qubits are a leading candidate for building large quantum processors in recent years, demonstrating significant progress towards quantum simulations [1, 2, 3, 4, 5, 6, 7, 8] and quantum error correction [9, 10, 11, 12]. Within these experiments, single and two-qubit gates are fundamental building blocks that dictate overall performance. The most widespread type of qubit used in these applications is the transmon qubit [13]. Two-qubit gates for transmon qubits were originally implemented with high fidelities using dynamic flux pulses to tune the qubits in and out of resonance [14, 15, 16] or by applying a strong microwave drive to activate an effective two-qubit interaction [17, 18, 19, 20, 21]. However, residual crosstalk often limits such architectures [22, 23]. More recently, tunable couplers have consistently reached high gate fidelities alleviating the crosstalk issue to a great extent [24, 25, 26, 27, 28, 29, 30].

Despite transmon based gates reaching higher fidelities, transmons have a small anharmonicity which eventually limits the speed of gates and can lead to leakage out of the computational subspace [31, 32, 33]. The small anharmonicity further imposes strong requirements on qubit frequency targeting to avoid frequency crowding [34, 35]. Additionally, dielectric losses are known to limit the coherence times of transmons [36, 37].

The fluxonium qubit is one appealing alternative which aim to mitigate these shortcomings as it is less prone to dielectric loss because of its lower qubit frequency and while also featuring a much larger anharmonicity [38, 39]. Coherence times in the millisecond range have been demonstrated in recent years [40, 41, 42]. Fast single-qubit and two-qubit gates have also been shown with high fidelities [43, 44, 45]. One challenge in implementing fast two-qubit gates between fluxonium qubits comes from their small charge dipole moment, suppressing the interaction between the computational states when directly capacitively coupled. Thus, direct capacitive coupling schemes for fluxonium qubits may require very large capacitors or very strong driving [46, 45, 47]; however, stronger coupling schemes are often accompanied by an increased residual interaction rate. A possible remedy is to use inductive coupling [48, 49]. However, inductive coupling may be accompanied with additional flux loops and require closer placement of qubits. Similarly to architectures for transmon qubits, an apparent solution is to introduce a tunable coupler to mediate strong interaction between the fluxonium qubits while simultaneously suppressing any residual coupling between the computational states [44, 48, 50]. In particular, two-qubit gates with high fidelities have been implemented by applying a microwave drive to the coupler [44, 51]. While high fidelities have been achieved with this gate scheme, the maximum reported fidelity utilized a reinforcement learning algorithm for gate calibration [44], a potentially time costly procedure. More-

over, while reinforcement learning algorithms can adapt to unknown situations, these black-box algorithms often lead to calibrated pulse shapes which are hard to interpret. Further, the uncertainty in the pulse shape makes it hard to understand the underlying physical mechanisms for drifts in the performance of the gate with time.

In this work, we implement a two-qubit gate between two fluxonium qubits coupled via a transmon coupler circuit, see Fig. 5.1(a). The gate is implemented by driving a coupler transition that is conditional on the state of the two fluxonium qubits, see Fig. 5.1(b). The transmon coupler interacts strongly with the higher levels of the fluxonium qubits which lead to distinct coupler transition frequencies for each computational state. We expect that a large frequency selectivity among these transitions lead to faster two-qubit gates. On the other hand, for faster gates, the spectral overlap between the pulse and unwanted coupler transitions lead to coherent errors in the gate. Here, we employ simple analytical pulse shaping techniques [52, 53] to achieve fast two-qubit gates. Given the analytical form of these pulse shapes, we can interpret how they minimize coherent errors. Specifically, the pulses are designed to minimize Fourier amplitude at the unwanted transitions. We experimentally obtain two-qubit gate fidelities of 99% for gate times as low as 60 ns.

5.2. DEVICE DESCRIPTION

The device used in this work is composed of two fluxonium qubits, see Fig. 5.1(a). These two fluxonium qubits are both capacitively coupled to a transmon coupler circuit in a geometry that accommodates up to four couplers connected to each fluxonium qubit as required e.g. for implementing the surface code [54, 55, 56, 57]. Throughout this work, we refer to the two fluxonium qubits as F_1 and F_2 and to the transmon coupler as C. The fluxonium qubits and the transmon couplers feature Josephson junctions fabricated using a Manhattan style process, see further details about the fabrication in Appendix 5.8 and in Ref. [58]. The fluxonium qubits include a flux loop each composed of the central Josephson junction with Josephson energy E_J and a Josephson array consisting of 100 junctions in series with a combined inductive energy of E_L. The flux loops of the fluxonium qubits as well as the symmetric SQUID-loop of the transmon coupler can be biased using their individual flux lines. Note that the coupler circuit is also coupled weakly to each drive line, such that the coupler can be readily driven by either. The fluxonium qubits are read out using dispersive readout enabled by their individual readout resonators. Similarly, the transmon coupler is also coupled to each readout resonator. Circuit layout of the three qubits is shown in Fig. 5.1(b). The parameters of the device are further described in Appendix 5.7.

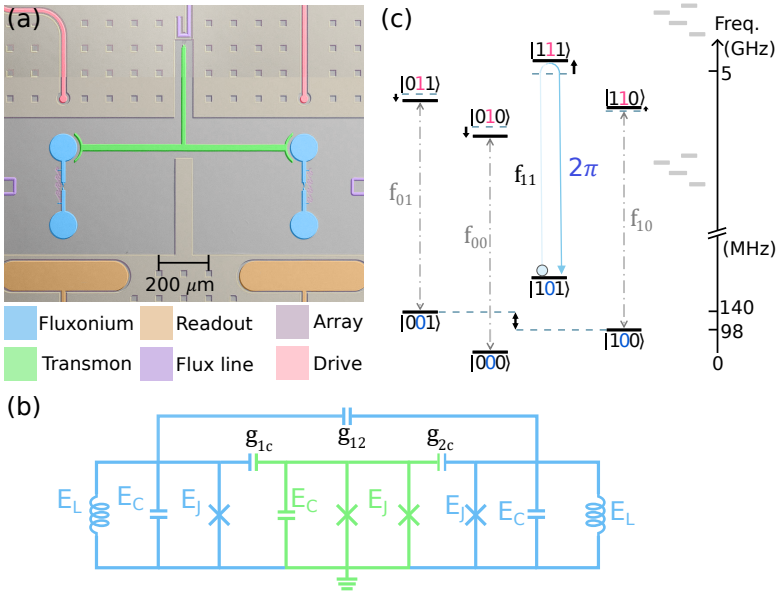

Figure 5.1: (a) False colored microscope image of the device zoomed in to a single set of the combined qubit system. The fluxonium qubits are indicated in blue and are coupled by a transmon displayed in green. All the qubits are flux biased by currents applied to the flux lines (shown in purple), and the fluxonium qubits are driven capacitively by the drive lines (shown in pink). The readout resonators (shown in orange) have large capacitive elements near the qubits and beyond this image, the resonators extend into $\lambda/2$ resonators. (b) Circuit representation of the two fluxonium qubit in blue and transmon in green. Energies corresponding to each component of the qubit is indicated along with the coupling between the qubits. (c) Energy diagram of the combined system with different levels. We use the notation $|F_1 C F_2\rangle$ to indicate the state of fluxonium 1, coupler and fluxonium 2, respectively. For clarity, we mark the system states corresponding to the coupler ground state in blue, constituting the computational subspace. The states with the coupler in the excited state are represented in red. We indicate the transition frequencies of the coupler with $f_{ij} = \omega_{ij}/(2\pi)$ conditioned on the fluxonium qubits' state $|ij\rangle$. Our gate scheme is operated by driving a 2π rotation on the coupler when the fluxonium qubits are in the $|11\rangle$ state indicated with the blue array. The gray energy levels on the right-hand side indicate the higher excited fluxonium states. These states shift the dressed coupler levels (bold black levels) away from the bare coupler levels (dash-dotted levels) allowing for each coupler transition to be selectively addressed.

We write the total system Hamiltonian as

$$\hat{H} = \hat{H}_{F_1} + \hat{H}_{F_2} + \hat{H}_C + \hat{H}_I \tag{5.1}$$

where each fluxonium qubit is described by their individual Hamiltonian

$$\hat{H}_{F_i} = 4E_{Ci}\hat{n}_{F_i}^2 + \frac{1}{2}E_{Li}\hat{\phi}_{F_i}^2 - E_{Ji}\cos(\hat{\phi}_{F_i} - \phi_{i,ext}), \tag{5.2}$$

where E_{Ci} is the charging energy,E_{Li} is the inductive energy and E_{Ji} is the Josephson energy of the ith fluxonium and $\phi_{i,ext}$ is the external flux-bias of each qubit. Similarly, the Hamiltonian for the transmon coupler is

$$\hat{H}_C = 4E_{C,C}\hat{n}_C^2 + E_{J,C}|\cos(\phi_{C,ext})|\cos(\hat{\phi}_C), \tag{5.3}$$

where $E_{C,C}$ is the charging energy of the coupler and $E_{J,C}$ is the total Josephson energy of the SQUID. The transmon is tuned by the external flux $\phi_{C,ext} = \pi\Phi_{C,ext}/\Phi_0$ where Φ_0 is the magnetic flux quantum. Since we have capacitive coupling between each fluxonium and the coupler as well as directly between the two fluxonium qubits, the interaction Hamiltonian is

$$\hat{H}_I = \hbar g_{12}\hat{n}_{F_1}\hat{n}_{F_2} + \sum_{n=1,2}\hbar g_{ic}\hat{n}_{F_i}\hat{n}_C \tag{5.4}$$

We describe the quantum state of this coupled system using the basis state F_1CF_2, referring to the eigenstates of the coupled system unless otherwise mentioned. As indicated in Fig. 5.1(b), the higher excited states of each fluxonium qubit couple strongly with the transmon coupler leading to level repulsions between the transmon levels and the fluxonium levels. In particular, the transition frequency of the coupler will depend on the state of the fluxonium qubits. We will refer to the transition frequencies of the coupler as $\omega_{ij} = (E_{i1j} - E_{i0j})/\hbar$ where E_x is the eigen-energy of state x. Since $\omega_{11} \neq \omega_{00}, \omega_{01}, \omega_{10}$, we can drive a selective 2π rotation on the ω_{11} transition such that the state 101 picks up a phase of π and thereby implementing a conditional phase gate.

The specific fluxonium qubits that we study in this work are operated with qubit frequencies at 98 and 140 MHz for F_1 and F_2, respectively. The transmon coupler has a transition frequency around 4.7 GHz when operated at zero flux bias and with the fluxonium qubits in their ground state. With these qubit frequencies, the device is in a regime where the residual ZZ coupling between is predicted to be below 0.5 kHz, see Fig. 5.2(a) where we extract the theoretical ZZ-coupling based on a numerical diagonalization of the full Hamiltonian. The extracted values of the coupling constant are $g_{12}/2\pi$ = 40 MHz and $g_{ic}/2\pi$ = 236

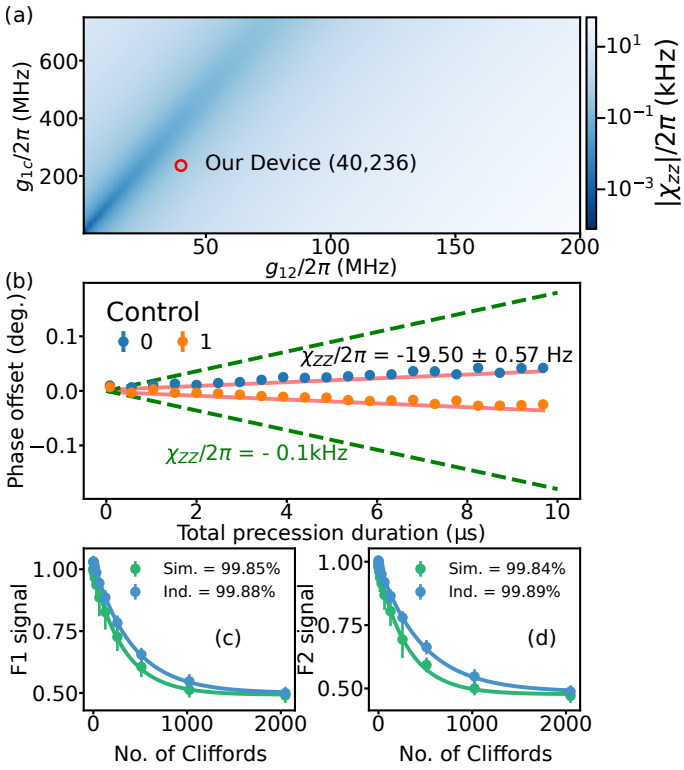

Figure 5.2: (a) The residual ZZ rate for a range of coupling strengths between the two fluxonium qubits, g_{12}, and between each fluxonium and the coupler, g_{ic}. Parameter values corresponding to our device are indicated in red circle. (b) Measured accumulated phase on F_2 dependent on the state of F_1. We extract the residual ZZ rate, χ_{ZZ}, from the slope of the signal. For reference, the theoretical value of $\chi_{ZZ} = 0.1$ kHz is plotted in green. (c, d) Randomized benchmarking for single qubit gates on fluxonium qubit F_1 and F_2, respectively, performed on each individual qubit (blue) and simultaneously on both qubits (green). The inset text displays the fitted average fidelity per gate.

MHz compared to the designed value of $g_{12}/2\pi = 47$ MHz and $g_{ic}/2\pi = 236$ MHz. Here, we define the residual ZZ coupling as

$$\chi_{ZZ} = \omega_{101} - \omega_{100} - \omega_{001} + \omega_{000}. \tag{5.5}$$

To verify this residual coupling experimentally, we perform a modified Ramsey sequence to precisely measure the accumulated phase of F_2 with F_1 in either state 0 or in 1, see also Appendix 5.10. As seen in Fig. 5.2(b), we extract a residual ZZ coupling of around -20 Hz, well below the predicted value of 0.5 kHz in magnitude. Given the small value of measured ZZ, the calculation is highly susceptible to uncertainties in the model parameters. For example, the extracted transition frequencies from the fit have errors on the order of MHz, meaning a relative error of around 0.1% but orders of magnitude larger than the residual ZZ coupling. Similarly, there may also be additional coupling terms that are not included in our model that may affect the coupling strength at the sub-kHz regime. Finally, it is well known, see e.g.[58, 59], that a simple lumped capacitive model of the system is insufficient to predict higher order dispersive features in superconducting qubit systems. However, the Hilbert space here is too big to perform a highly accurate high-frequency simulations of the full device. Overall, we expect the residual ZZ coupling to cause little correlated crosstalk errors.

To further characterize the crosstalk between the two qubits, we employ simultaneous randomized benchmarking [60, 61], see Fig. 5.2(c) and (d). For these results, we decomposed the Clifford group into combinations of $\pi/2$-pulses and virtual-Z gates [62]. Note that the readout signal is slightly shifted for the simultaneous protocol due to a shift in the readout signal arising from the microwave reset scheme of the fluxonium qubits, see also Appendix 5.12. We observe that the single-qubit gate fidelity is only slightly affected by simultaneous driving. Due to the small residual ZZ coupling, we attribute the small degradation in the single qubit fidelities to microwave crosstalk between the fluxonium qubits.

5.3. TWO-QUBIT GATE IMPLEMENTATION

Our gate scheme involves driving the microwave transitions of the transmon coupler. To verify that we can selectively drive the coupler based on the state of the fluxonium qubits, we perform conditional spectroscopy of the coupler, see also Appendix 5.11. As shown in Fig. 5.3, we see four distinct frequency transitions between 4.5 GHz and 4.9 GHz for two different flux bias points. Due to fabrication uncertainty, the coupler frequency in the device lies close to the resonance frequency of the readout resonator of F_1 at 4.993 GHz. As a consequence, the coupler is partly hybridized with the readout resonator and we observe an overall degradation of qubit coherence. To mitigate this detrimental effect, we use the frequency flexible aspect of our design and flux bias the coupler

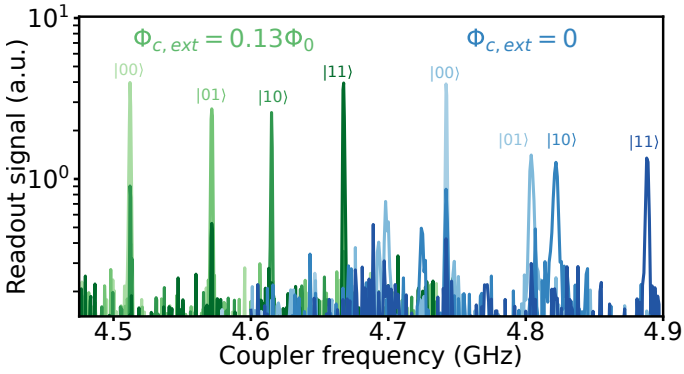

Figure 5.3: Conditional spectroscopy of the coupler state when fluxoniums F_1, F_2 are prepared in different states $|F_1, F_2\rangle$, labeled above each peak. Two sets of peaks are observed for different flux operating points of the coupler, corresponding to $\Phi_{c,ext} = 0$ shown in shades of blue and $\Phi_{c,ext} = 0.13\Phi_0$ shown in shades of green.

to $0.13\Phi_0$ for the remainder of this work. At this operating point, the measured coupler lifetimes are on the order of 1.2 μs. Further, at the chosen flux bias point, four distinct frequency transitions of the coupler are observed, see Fig. 5.3. A detailed description regarding the lifetimes and frequency selectivity is given in Appendix 5.7. We verify that we can coherently drive conditional Rabi oscillations on the coupler transition by applying a microwave pulse to the driveline of F_2 which also capacitively couples to the transmon coupler. As shown in Fig. 5.4, depending on which state the fluxonium qubits are prepared in, we see a Rabi oscillation at different frequencies. Note that while the transmon coupler does not have an individual readout resonator, we can combine the readout signal of F_1 and F_2 to monitor the coupler dynamics, see also Appendix 5.11. Importantly, we notice that, compared to the other states, the fastest Rabi rate is observed when preparing the 11-state, thus, we expect that fastest possible gate to be achievable by driving this transition. Additionally, the coupler frequency corresponding to the prepared 11 state of the fluxonium qubits is most detuned from the $1-2$ transition of F_1, and the coupler has the longest lifetime in this state, see Appendix A. Specifically, we apply a drive leading to the Hamiltonian:

$$\hat{H}_d(t) = [\hbar\Omega_I(t)\cos(\omega_d t) - \hbar\Omega_Q(t)\sin(\omega_d t)]\,\hat{n}_{F_2}, \tag{5.6}$$

where $\Omega_I(t)$ and $\Omega_Q(t)$ are the pulse envelope of the in-phase and quadrature components of the drive, respectively and ω_d is the drive frequency. To achieve a fast high-fidelity gate, we must appropriately design the pulse shape of the drive as well as optimize the driving frequency. A black-box approach to finding the optimal driving parameters can be employed using the optimized randomized

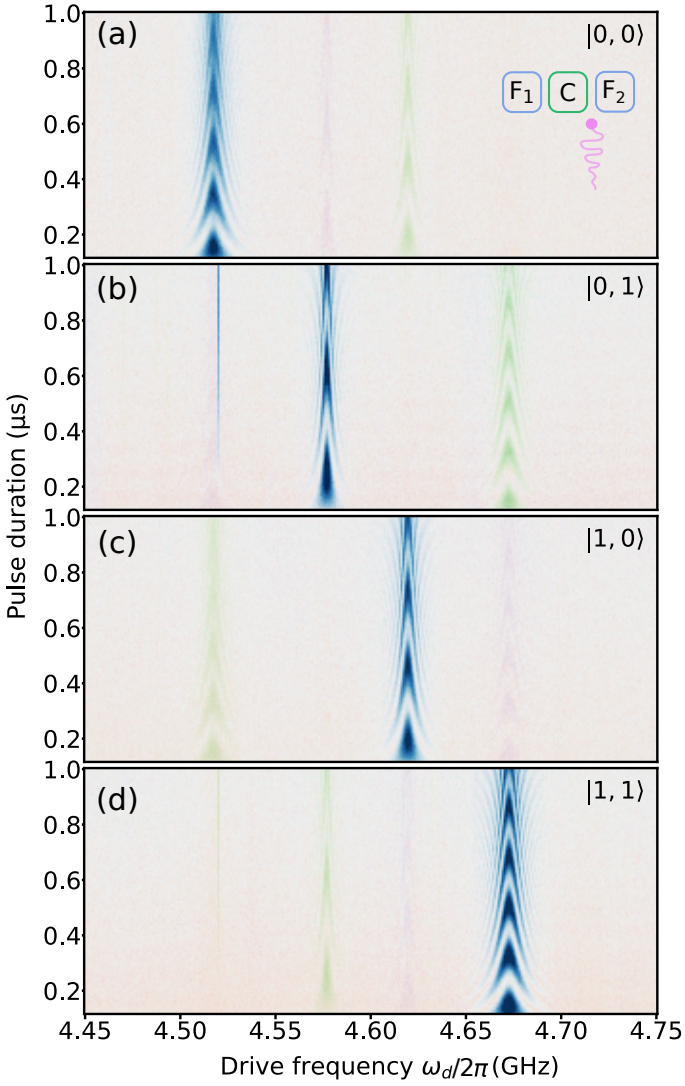

Figure 5.4: Rabi oscillations when driving near the coupler frequencies are shown with different initialization of the fluxonium state, 00, 01, 10 and 11 in panels (a), (b), (c) and (d), respectively. The readout signal of both fluxonium qubits are analyzed using principal component analysis and we visualize the main components in blue-to-green and red-to-purple, see Appendix E. The coupler is driven through the drive line of F_2 as indicated in the inset of (a).

benchmarking for immediate tune-up (ORBIT) protocol [63] or using reinforcement learning algorithms [44]. However, the iterative feedback-loop between the experiment and the classical optimizer may lead to a time-costly calibration procedure. Here, we instead use a straightforward calibration approach with no iteration steps which allows for a fast and deterministic calibration procedure, see Appendix 5.13. In our gate scheme we choose to drive the population from the 101-state to the 111-state and back again to the 101-state to obtain the conditional phase (CZ) gate. To ensure that we recover all the population back to the 101-state, we first fix the pulse shape, i.e., $\Omega_I(t)$ and $\Omega_Q(t)$, and then we measure the coupler response as a function of the drive frequency and the absolute pulse amplitude of the drive. Thus, for each pulse frequency, we deterministically obtain the pulse amplitude that maximally recovers the population back into the 101-state, see also Appendix 5.13. Next, we measure the conditional phase from the application of 1, 3 or 5 gates. For a well-calibrated conditional phase gate, each of these should yield a conditional phase of exactly π. We measure the conditional phase for a range of frequencies and we pick the drive frequency where the conditional phase is equal to π. Finally, we measure the single qubit phases of both F_1 and F_2 subject to repeated application of the gate. We correct for these single qubit phases using virtual Z gates [62].

Using this calibration routine, we first explore a simple pulse shape where the in-phase drive is a raised cosine

$$\Omega_I(t) = \frac{\Omega}{2}[1 - \cos(2\pi t/t_g)], \tag{5.7}$$

where t_g is the total duration of the gate. Additionally, we set $\Omega_Q(t) = 0$. Using interleaved randomized benchmarking [60, 61], we characterize the gate fidelity of the CZ gate, see Fig. 5.5(a). For a cosine pulse, we find a maximal gate fidelity of $98.9\% \pm 0.1\%$ for a gate duration of 68 ns. For longer gate times, the fidelity becomes limited by the coupler coherence times, see the extracted CZ errors in Fig. 5.5(c). For gate durations below 60 ns, the gate error increases significantly, which we interpret as resulting from driving other spurious transitions such as the transmon coupler frequencies associated with different fluxonium qubit states. A well-known technique to eliminate residual driving of weakly off-resonant transitions is the derivative removal by adiabatic gate (DRAG) technique [52]. Using DRAG, we can eliminate the spurious driving at a detuning of Δ by defining the quadrature drive as:

$$\Omega_Q(t) = -\frac{1}{\Delta}\frac{d\Omega_I(t)}{dt}. \tag{5.8}$$

In our case the nearest transition that we are interested in suppressing is ω_{10} (see also Fig. 5.3). Thus, we can set $\Delta = \omega_{10} - \omega_d$ to obtain the pulse shape of

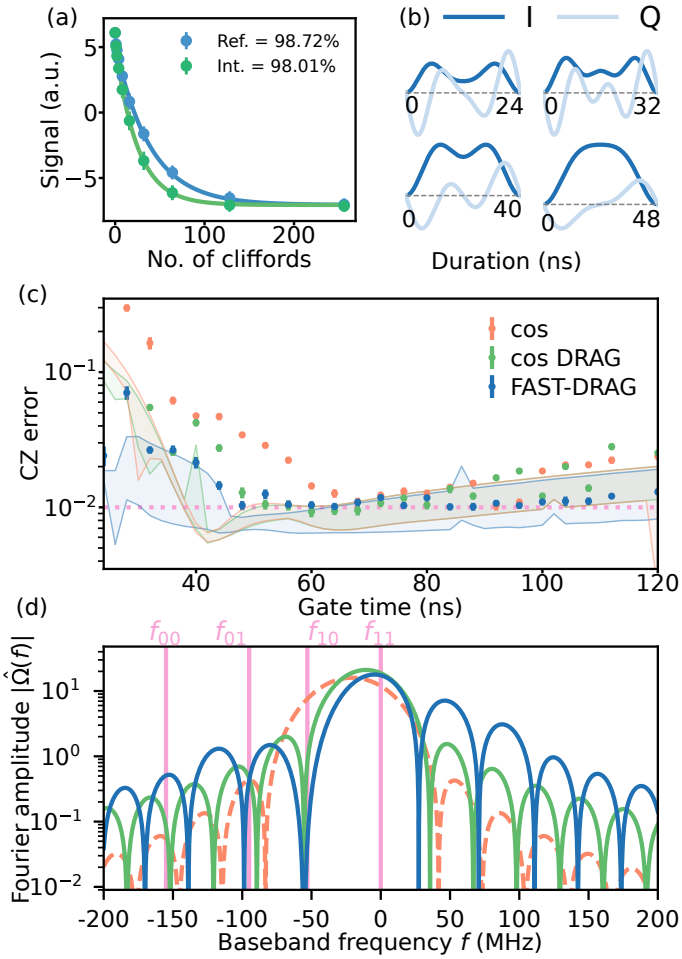

Figure 5.5: (a) Clifford interleaved randomized benchmarking for a two-qubit CZ gate. The readout signal for the reference sequence is shown in blue and in green we show the when interleaving the CZ gate. The extracted fidelities per Clifford is shown in the legend. (b) Pulse shapes used in the FAST-DRAG scheme are shown. Normalized I and Q components are shown in dark blue and light blue, respectively. (c) Average CZ gate infidelity for the simple cosine pulse shape (orange), for the DRAG scheme (green) and for the FAST-DRAG pulse shape (blue). Error bars correspond to 1 standard deviation in the least-squares parameter fit. The horizontal dashed line indicates a fidelity of 99%. The shaded areas indicate the numerical simulations, see details in the main text. (d) Fourier spectrum with different pulse shapes is shown for a gate duration of 32 ns. Four vertical lines have been marked at the top corresponding to the transmon coupler frequencies f_{00}, f_{01}, f_{10} and f_{11}.

the drive. Using this pulse shape, we see in Fig. 5.5(c) that we can obtain a gate fidelity around 99% for a larger range of gate times compared to the simpler cosine pulse shape. However, the DRAG scheme can only be used to suppress one of the three residual transitions.

To extend the degree to which we suppress the driving of residual transitions, we adopt the Fourier *ansatz* spectrum tuning derivative removal by adiabatic gate (FAST-DRAG) method as introduced in Ref. [53]. Specifically, we use a pulse envelope given by

$$\Omega_I(t) = A\sum_n c_n\Big(1-\cos\frac{2\pi n t}{t_g}\Big), \tag{5.9}$$

where t_g is the gate duration, A is the overall amplitude and c_n the coefficients for the cos series. Additionally, we add the quadrature drive following Eq. (5.8) with the $\Delta = \omega_{10} - \omega_d$ as before. As detailed in Appendix 5.14, we engineer the Fourier spectrum of the pulse by solving a convex minimization problem with three cos components in the series. In particular, we aim to minimize the Fourier components of the pulse in narrow frequency windows around f_{01}, f_{10} and f_{00}. Examples of the resulting pulse shapes are displayed in Fig. 5.5(b) and the resulting average gate fidelities are shown in Fig. 5.5(c). We find a minimal gate error of $99.0\% \pm 0.1\%$ for a gate time of 64 ns with both the DRAG scheme and with the FAST-DRAG scheme. It is interesting to note that while the overall lowest error is not improved compared to the DRAG scheme, the FAST-DRAG provides good performance over a broader range of gate duration showing that we more consistently accumulate less coherent errors from residual driving of spurious transitions. This is particularly observed for gates with gate times shorter than 50 ns. We visualize the impact of Fourier engineering by plotting the resulting pulse spectrum, see Fig. 5.5(d). Generally, we find that the simple cosine pulse has large frequency components around all spurious transmon transitions. In contrast, the DRAG scheme minimizes the component at f_{10} as expected while the FAST-DRAG makes the window around f_{10} broader while also suppressing the frequency components around f_{01}. Given the over-constrained optimization problem, we find that the frequency component of f_{00} remains less suppressed.

To gain further insights into the errors, we simulate the gate numerically. We model the fluxonium qubits and the transmon coupler with the parameters discussed in Appendix 5.7 and implement a time dependent drive using the qiskit dynamics package [64] . In the simulations, we include four energy-levels for each fluxonium qubit and three levels for the transmon coupler. From the resulting unitary evolution, we truncate the evolution to the computational subspace of the two fluxonium qubits to obtain the matrix U from which we calculate the process fidelity as $F_p = |\mathrm{Tr}(U^\dagger_{ideal}U)|^2/d^2$, where $d = 4$ is the Hilbert

space dimension. For each type of pulse shape and duration, we numerically optimize the drive amplitude and the drive frequency in order to maximize the gate fidelity. In the simulations, we include a drive on fluxonium qubit F_2 and residual relative driving on F_1 and on the transmon coupler. Specifically, we use a drive Hamiltonian $\hat{H}_D \propto \hat{n}_{F_2} + \eta_C \hat{n}_C + \eta_{F_1} \hat{n}_{F_1}$ to describe the relative capacitive crosstalk, η_i from the driveline of F_2 to the transmon coupler C and fluxonium qubit F_1. We fix $\eta_{F_1} = 0.028$ based on capacitance matrix simulations of the device design. We find the simulated gate fidelities to be very sensitive to crosstalk, so to capture uncertainties in the precise crosstalk, we vary η_C between 0.2 and 1.0. The size of the full Hilbert space of two fluxonium qubits and the transmon coupler including the higher excited states prohibits efficient simulations as an open system. Thus, we include transmon T_1 in an ad-hoc manner. Specifically, we calculate the process fidelity of each transmon transition as

$$F_{p,ij} = \frac{1}{2}(1 + e^{\int_0^{tg} \gamma_{1,ij} p_{i1j}(t) dt}), \tag{5.10}$$

where $p_{i1j}(t)$ is the probability to be in the state $i1j$ and where $T_{1,ij} = 1/\gamma_{1,ij}$ is the lifetime of the coupler when the fluxonium qubits are prepared in the state ij. We calculate the total process fidelity as $F_{P,total} = F_P \prod_{ij} F_{ij}$ and finally we extract the average gate fidelity as $F_g = (dF_{P,total} + 1)/(d + 1)$. As we increase η_C, the residual population in the transmon increases and, thus, the CZ errors increase as well. In Fig. 5.5(c), the shaded area indicate the span from the minimum to the maximal gate errors as we vary η_C and we specifically observe that the cross-driving of the coupler can lead to an additional error of up to 1%. For pulse durations longer than 60 ns, we see that the overall behavior of simulations matches well with the experimental data, thus, the coupler lifetime remains the main limiting factor for longer gate times. We also note that FAST-DRAG can potentially lead to smaller errors for longer gates since this drive scheme will cause less residual driving of the transmon and, thus, less potential errors due to the coupler lifetime. This effect is also partly visible in the experimental data. For the gate times smaller than 50 ns, we expect a lower error from the simulations, thus the experiment is clearly limited by residual driving of other transitions. From the simulations, we can identify residual driving of the 1-2 transition of F_1 through the transmon coupler as the main cause of coherent errors, see Fig. 5.8(a). Additionally, for the gate times smaller then 50 ns, the FAST-DRAG gives experimentally lower error compared to the conventional DRAG scheme. In fact, when there is large crosstalk to the transmon coupler, the simple cosine pulse is expected to outperform the FAST-DRAG scheme since the cosine pulse is overall more localized in frequency. Additionally, the simulations do not include crosstalk to the readout resonator detuned only 320 MHz from the drive

frequency. For the short gate times, residual driving of the readout resonator may lead to additional dephasing of the qubits.

Finally, the simulations indicate that the simple cosine and the DRAG scheme should perform similarly. However, in the experiment we clearly see an improvement from using DRAG. The reason is likely that the numerical optimization scheme simultaneously optimizes the amplitude and detuning of the gate and finds a good compromise where a larger detuning is fine-tuned to get the precise conditional phase. In the experiments, we optimize the population recovery independent of the conditional phase, as described in Appendix 5.13. Thus, our simulations points to a clear direction for future devices, namely that we must minimize microwave crosstalk between the fluxonium qubits and improve the frequency targeting of the resonators and the coupler in a fluxonium-transmon-fluxonium architecture.

5.4. CONCLUSION

In this work we have explored a system consisting of two fluxonium qubits coupled directly and through a transmon coupler circuit. The device parameters were designed such that the residual ZZ crosstalk between the two fluxonium qubits is suppressed. In this architecture, we can selectively drive the transmon coupler based on the fluxonium qubits' state. Thus, by driving a full 2π rotation conditioned on the fluxonium qubits being in the 11 state, we were able to implement a conditional phase (CZ) gate. Using a simple cosine pulse shape for the drive, we obtained a fidelity of $98.9\% \pm 0.1\%$ with a gate duration of 68 ns. By employing pulse shaping techniques, we optimized the performance further for shorter pulse durations. Specifically, we use the DRAG and the FAST-DRAG techniques to obtain fidelties around $99.0\% \pm 0.1\%$ over a larger range of gate durations. These pulse shaping techniques use simple analytical formulas to suppress the driving of residual coupler transitions and thereby minimizing the coherent errors from the drive.

5.5. ACKNOWLEDGMENTS

S.S., F.Y. and S.W. carried out the numerical simulations and calculations for the device energy parameters. S.S., F.Y. and P.K. fabricated the device. S.S., E.Y.H., J.H., M.F.S.Z. and T.V.S. characterized the device and performed the measurements for the experiments. S.S, E.Y.H., J.H., F.Y., M.F.S.Z. and T.V.S. contributed to the measurement setup. S.S wrote the manuscript with feedback from all coauthors. C.K.A. supervised the execution of the project and the writing of the manuscript.

The authors thankfully acknowledge support from the Kavli Nanolab Delft

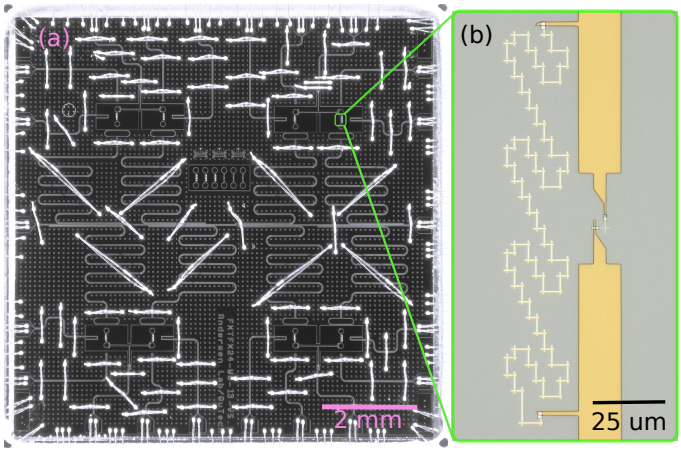

Figure 5.6: (a) The device is shown with all four sets of the three qubit system. The device can be seen with aluminum wirebonds on top to connect ground planes across control lines and Coplanar waveguide resonators. (b) A single fluxonium qubit zoomed in to show the array design and the single junction. The array comprises of total 100 junctions with a dimension of 450×450 nm^2, and the single junction of 150×150nm^2.

cleanroom staff members including B. van Asten, E.J.M. Straver, E.F.D. Pot, B.P. van den Bulk, A. van Run, C.R. de Boer, M. Fischer, A.K. van Langen-Suurling and M. Zuiddam. We also acknowledge fruitful discussions with A. Kamlapure, M. Finkel, H.M. Veen and D.J. Thoen. The authors also acknowledge support from J. Bauer and L.J. Splitthoff on the junction development process.

All authors acknowledge receiving support from the Dutch Research Council (NWO). E.Y.H. acknowledges funding support from Holland High Tech (TKI), J.H. acknowledges funding support from NWO Open Competition Science M and T.V.S. acknowledges support from the Engineering and Physical Sciences Research Council (EP- SRC) under EP/SO23607/1.

5.6. DATA AVAILABILITY

The data for the experiments can be found at [65] with an open CC by 4.0 license. Scripts to analyze the data can be found in the gitlab repository at [66]

5.7. APPENDIX A: DEVICE MEASUREMENTS

In this section, we describe the device studied in this work and our approach to extract the device parameters. The device incorporates four pairs of fluxonium qubits each with a transmon coupler, as shown in Fig. 5.6(a). In this work, we focus on a single set of fluxonium qubits. Qubit spectroscopy for one of the fluxonium qubits is shown in Fig. 5.7, where pulsed spectroscopy peaks (blue

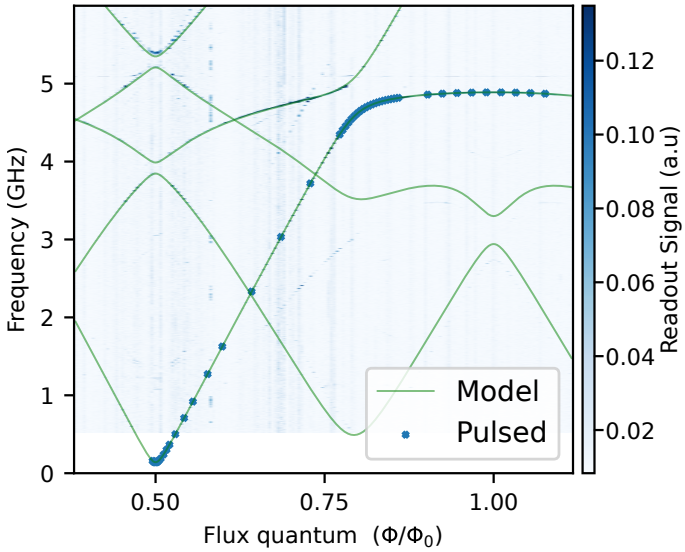

Figure 5.7: Two-tone continuous wave spectroscopy for fluxonium F_2 as a function of external flux and drive frequency shown in the background with spectroscopy peaks as blue crosses. A single fluxonium spectrum model is overlaid as solid green lines, with energy parameters fitted to the pulsed spectroscopy data.

markers) were fitted to a single fluxonium spectrum model. This fitting gives us the individual fluxonium energy parameters listed in Table 5.1. The transmon capacitive charging energy is estimated by its measured anharmonicity and its Josephson energy is inferred from the measured coupler frequency. At the operational point, we measure the fluxonium qubits and the couplers lifetimes and coherence times using standard time-domain techniques, see Table 5.1.

We use the the scqubits package [67] to analyze the energy spectrum of the system in more detail. The energy spectrum is further used to calculate the coupling values (g_{ic}, g_{ij}). The fitting is done with the shift in coupler frequencies when the coupler is in either 0 or 1, corresponding to different initialized fluxonium states. In Fig. 5.8(a), we see how the transmon transition frequencies are positioned between the 1 to 2 transitions and the 0 to 3 transitions of each fluxonium. In particular, we notice a strong avoided crossing between the 0 to 1 transition of C and the 1 to 2 transition of F_1 as we increase the external flux bias of the transmon coupler. Moreover, it is essential to ensure that the coupler is sufficiently detuned from the fluxonium resonators to avoid Purcell decay of the coupler transmon into the readout resonator. At an external flux of $\Phi/\Phi_0 = 0$, the transmon coupler transitions approaches the readout resonator frequencies. Biasing the transmon coupler away from zero external flux reduces the Purcell

Table 5.1: Measured device parameters with F_1, F_2 being the fluxoniums and C being the coupler. The transition frequency and lifetimes of the coupler are measured with the fluxonium states prepared in $|00\rangle$, $|01\rangle$, $|10\rangle$ & $|11\rangle$, written in this order respectively. T_2^E of the coupler for the $|00\rangle$ fluxonium state could not be measured, marked –.

Measured	F_1	C	F_2
T_1	72.3 μs	0.54 μs, 1.55 μs 1.05 μs, 1.19 μs	89.6 μs
T_2^R	17.86 μs	0.77 μs, 1.23 μs 1.19 μs, 1.19 μs	18.11 μs
T_2^E	21.34 μs	–, 2.38 μs 2.45 μs, 1.93 μs	25.50 μs
$\omega_{01}/2\pi$	98.95 MHz	4.517 GHz, 4.577 GHz 4.619 GHz, 4.672 GHz′	144 MHz
$\omega_{\text{res}}/2\pi$	4.993 GHz	–	5.082 GHz
$g_{ic}/2\pi$	236 MHz	–	236 MHz
$g_{12}/2\pi$	40 MHz	–	40 MHz
E_J/h (GHz)	4.9928	16.87	4.3350
E_L/h (GHz)	0.5008	–	0.4921
E_C/h (GHz)	0.8805	0.1861	0.8829
$\chi/2\pi$	1.029 MHz	–	1.319 MHz

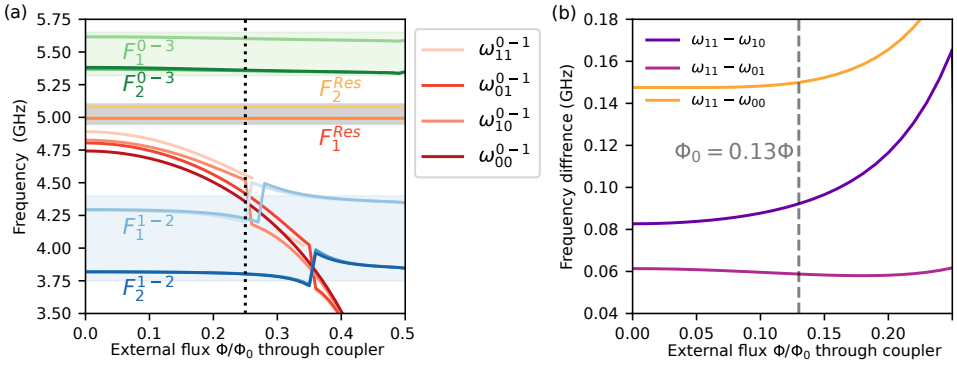

Figure 5.8: (a) Energy levels of the fluxonium-transmon-fluxonium system as a function of the external flux of the transmon. The transition frequency from state n to m for fluxonium i is indicated with F_i^{n-m}. In the blue and green regions we see the 12 and 03 transitions of the fluxoniums respectively. The grey region contains the two readout resonators. The coupler transitions are labeled ω_{ij}^{0-1} corresponding to the fluxonium qubits in state ij. (b) Differences between the undesired transition of the transmon (with the fluxoniums in $00, 01, 10$) and the desired transition 11 as a function of the external flux of the transmon.. ω_{ij} represent F_1 prepared in the ith state and F_2 in jth state, respectively. A black vertical dashed line is marked to show value of $\Phi = 0.13\Phi_0$ corresponding to the operational point used in the main text.

decay rate, although the coupler becomes more sensitive to flux noise. At the external flux of $\Phi/\Phi_0 = 0.13$, the coupler is detuned by $\Delta/2\pi = 321$ MHz, along with a measured $\kappa/2\pi = 3.276$ MHz for the resonator of qubit F_1 and $g/2\pi = 17$ MHz for coupling between the coupler and resonator of qubit F_1, giving a Purcell limited lifetime of 17.3 μs

Furthermore, as seen in Fig. 5.8(b), when we bias the transmon coupler away from zero external flux, we also increase the frequency selectivity of the desired transitions frequency ω_{11} relative to the ω_{00} and ω_{10}. The dispersive shift (χ) corresponding to the resonators of fluxonium qubits F_1 and F_2 is measured and for the readout of both fluxoniums, a readout pulse length of 2.24 μs was used.

5.8. APPENDIX B: FABRICATION

The device used in this work has been fabricated using a recipe similar to those described in [58, 68], we will here describe further details as several improvements have been implemented in the process. The process begins with cleaning a high-resistivity silicon wafer (>20 kΩ·cm) with a (100) orientation, sourced from Topsil. The wafer is first immersed in a nitric acid bath for 7 minutes, accompanied by sonnication, to remove organic residues from the surface. Next, it is sequentially rinsed in two deionized water baths for 60 seconds each before being dried with nitrogen gun. The wafer undergoes an oxide removal step us-

ing a 40% HF solution for 6 minutes, followed by another deionized water rinse as described above. To further passivate the surface and delay oxide formation, the wafer is coated with HMDS vapor [69]. Since this passivation has a limited effective duration, the wafer is immediately loaded into a sputter system. For deposition of NbTiN, we utilized a sputtering system with a confocal target geometry. The deposition is performed using an NbTi target with a chemical composition of Nb:Ti (70:30). We deposit 200 nm of NbTiN using a deposition rate of 28.3 nm/s. To further process the wafer, we dice it into smaller pieces of 18×18 mm^2. To remove the protective resist from the dicing step, the device is put in NMP for 2 hours at 80 °C and additionally left in NMP at room temperature over night. To remove potential resist residues, we follow the overnight cleaning step with IPA for a 1 min, a 65 °C PRS for 20 minutes and finally IPA for 1 min.

The next step is to pattern the NbTiN using electron beam lithography. To prepare the sample for the lithography step, an anisotropic oxygen plasma descum, with 20sccm O_2 at 0.1mbar and 20W for 2 minutes, is used to remove organic residues. To improve resist adhesion, we bake the sample for 10 minutes at 150°C and spin a HMDS layer. We use AR.P6200.18 as resist, which we spin at 2500 rpm and bake at 160°C for 3 min. We finally pattern the coplanar waveguide and capacitive structures of the device using a Raith 5200 EBPG system using a beam current of 250 nA and a dose of 350 μC/cm^2. After exposure, the resist os developed in a multi step process: (i) Amyl acetate development for 30 s with sonnication, (ii) Amyl acetate development for 30 s without sonnication in a separate beaker, (iii) MIBK for 10 s, (iv) IPA for 60 s, (v) MIBK for 20 s with sonnication, (vi) 2 times: IPA for 15 s in separate beakers. To remove resist residues an additional 2-minutes oxygen descum is performed as in the previous descum step. The pattern is etched using reactive ion etching system in a two-step recipe: The first step uses SF_6:O_2 13.5:4 at 70 W and a pressure of 0.1 mbar until a signal-drop from the end-point detection system. The second step step uses SF_6:O_2 4: 16 at 50 W and a pressure of 0.08 mbar. We perform a final descum step with a 220 sccm flow of O_2 at 200 W for 30 s.

To remove the resist, we use the resist remover PRS-3000 first for 1 minute with sonnication followed by 3 hours without sonnication in a separate beaker. To further remove any organic residues from the surface, we clean for 30 s using nitric acid, followed by 25 minutes cleaning using 7: 1 BOE to remove any native oxides [70].

The Josephson junctions are patterned with E-beam lithography. This process uses a trilayer stack of MMA El8; PMMA 495 A8 and PMMA 950 A3 which we spin at 3000, 1000 and 3000 rpm, respectively. Each spinning step involves baking at 185°C for 10 minutes and cooling down for 2 minutes. The junctions are defined by exposing the resist using a 5nm resolution beam and a current of 1nA.

The exposure is designed to define a tapered undercut profile [71] using a selective exposure of the MMA resist. The undercuts are defined using a dose of 170 $\mu C/cm^2$ and we use a ratio of the junction dose to the undercut dose of 6:1. The development of the resist is performed in a H2O:IPA (1:3) mixture at room temperature with sonnication for 3 min 10s. Subsequently, the device is cleaned in IPA 2 times for 10 s and 1 minute, respectively. To clean resist residues, we do an oxygen descum, followed by cleaning with a 7:1 BOE mixture [72]. Immediately afterwards, we load the device into a Plassys MEB550 system for the electron-beam evaporation process. The junction deposition starts with 30 nm Al deposition followed by a 5 mbar 20 min oxidation step. Then we deposit 110 nm of Al after rotating the sample 90° followed with another oxidation at 1.3 mbar for 11 mins.

The chip is cleaned using NMP at 80°C for 2 hours and then left overnight at room temperature. Next, the device is cleaned again using NMP while removing the residual aluminum film with a pipette. We additionally clean the device in 2 more times with NMP and 3 additional times with IPA for 5 minutes each.

To create a better galvanic connection between the junctions and the base layer, we deposit additional aluminum bandages [73]. We spin a resist stack of MMA EL8 and PMMA 950 A4 at 4000 and 2000 rpm, respectively. Each resist layer is baked at 180°C for 5 minutes. The resist is exposed using electron-beam lithography with a beam current of 250 nA and a dose of 1200 $\mu C/cm^2$ and developed with s H2O:IPA (1:3) mixture at 6°C for 1 min 30 s followed by 30 s IPA.

Prior to the Al deposition a ion-milling step is performed to remove native oxides and resist residues and we deposit 150nm of Al film followed by 1.3-mbar 11 min oxidation. The lift-off is done similar to the junction lift-off step. The final result is illustrated in Fig 5.6(b), where we see both the smaller Josephson junction as well as the larger Junctions in the array that were fabricated simultaneously in this process. The whole device is finally diced again and we wirebond the device to a PCB as shown in Fig 5.6(a).

5.9. APPENDIX C: EXPERIMENTAL SETUP

The PCB holding the device is attached to a gold-plated copper mount and covered with two rectangular aluminum shields. The assembly is installed in a Blue Fors LD-400 system. The sample holder is further covered with three coaxial shields, consisting of two mu-metal cans and one gold-plated copper can. The PCB is connected to room-temperature electronics with coaxial microwave cables as shown in Fig. 5.9.

We use a Zurich instruments (ZI) UHFQA to generate the input and output signals for readout. The signals are up- and down-converted using a ZI HDIQ IQ modulator and a custom built IQ modular, respectively. The LO is provided

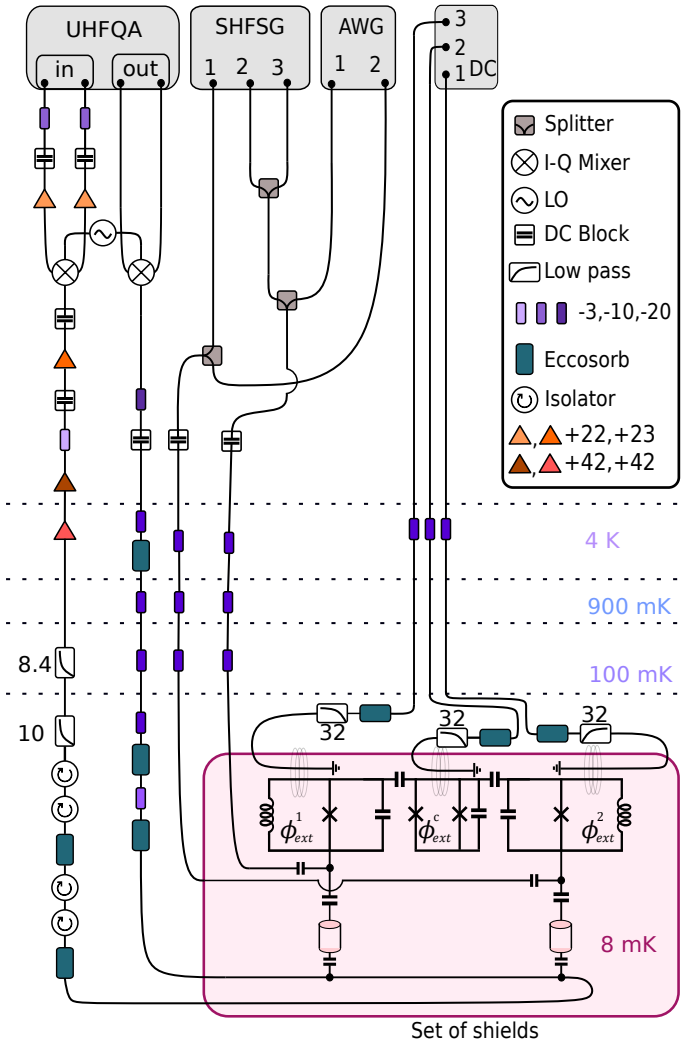

Figure 5.9: Schematic of the wiring in the dilution fridge and electronic setup used during the experiment. Inset shows different components, the numbers shown represent the attenuation and amplification in dB. Filters used on the output line of the transmission chain are written in GHz, while the filters connected on flux lines are represented in MHz.

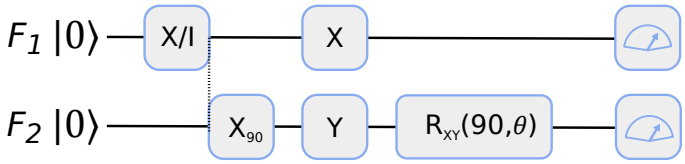

Figure 5.10: Circuit for measuring the residual ZZ interaction rate.

by a signal generator from from Anapico. The output signal is amplified at the 4K stage using a HEMT amplifier from Low-Noise factory and at room temperature using a series of amplifiers as shown in Fig. 5.9. For driving the individual fluxonium qubits, we a ZI HDAWG which generates signals directly at the qubit frequencies. Additionally, we use a ZI SHFSQ to drive the transmon coupler as well as for qubit reset, see also Appendix 5.12. These signals are combined with power combiners at room temperature. For flux bias of the qubits, we use in-house built current source modules (S4G). The flux lines on each qubit have a 32 MHz low pass filters to filter out noise at the fluxonium qubit frequency.

5.10. APPENDIX D: ZZ MEASUREMENT

The residual ZZ rate, as defined in Eq. (5.5), was measured using the circuit in Fig. 5.10. One fluxonium is designated as the control qubit and is initialized in either the $|0\rangle$ or $|1\rangle$ state. The other fluxonium is initialized in an equal superposition state. Both qubits freely evolve for a duration $\tau/2$ and we then apply refocusing π pulses on both qubit simultaneously. Both qubits then freely evolve again for a duration $\tau/2$. Finally, a $\pi/2$ recovery rotation is applying for a set of different rotation angles θ. During the free evolution time, the target qubit rotates around the Z-axis by $+\chi_{ZZ}\tau/2$ ($-\chi_{ZZ}\tau/2$) when the control qubit is prepared in the $|0\rangle$ ($|1\rangle$) state. This phase offset can be extracted from a cosine fit to the target qubit population as a function of θ.

5.11. APPENDIX E: COUPLER READOUT AND RABI

Since the transmon coupler does not have a dedicated readout resonator, we measure its state using a Rabi assisted readout method [74]. In this scheme, both fluxoniums are initialized and kept in a known $|0\rangle$ or $|1\rangle$ state until the transmon is read out, at which point a π-pulse is applied on both fluxonium qubits which will only be resonant if the coupler is in its ground state. If the coupler is in its excited state, the fluxonium qubit frequencies are shifted by more than sixty megahertz, see Fig. 5.4, hence the π-pulse will not flip the state of the fluxoniums. In other words, if the fluxonium qubits have flipped state, then the coupler is in its excited state.

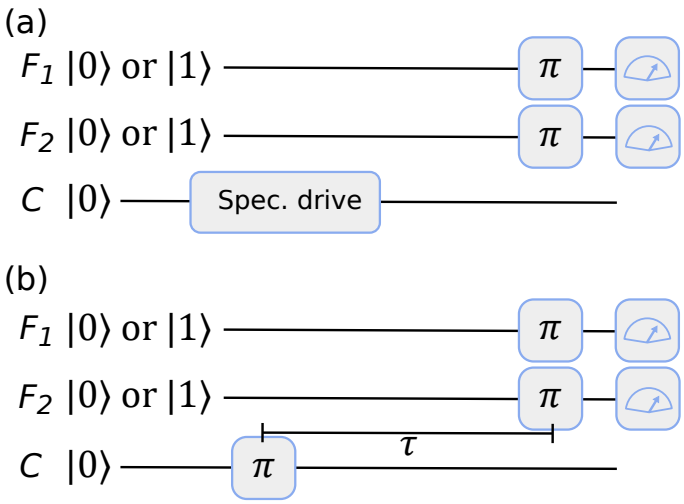

Figure 5.11: (a) Circuit to measure the conditional spectroscopy for the coupler and fluxoniums is shown. (b) Circuit to measure the coherence times of the coupler.

To avoid calibrating the state discrimination boundary in Rabi assisted readout, we work with the complex integrated readout traces directly. Multiplexed readout produces a complex value for each fluxonium, resulting in a total of four real values. To observe Rabi oscillations in the coupler, we perform principal component analysis on the four-dimensional data points. The first principal component should capture the largest variations, corresponding to oscillations in the final coupler state. The second and subsequent principal components should have only small variations if the system is truly confined to a two state manifold.

To characterize the lifetime of the coupler, we use the circuit shown in Figure 5.11(b), where the coupler is excited with a calibrated π pulse. In this measurement, we initialize each fluxonium qubit in either 0 or 1 to extract any state-dependent lifetime of the coupler.

5.12. APPENDIX F: RESET

We reset the qubits using a two-tone driving scheme similar to [75]. The qubit reset protocol uses two drive pulses at the angular frequencies ω_1 and ω_2. The drive frequencies correspond to transition frequency between $|1,0\rangle$ and $|2,0\rangle$ and to the transition frequency between $|2,0\rangle$ and $|0,1\rangle$, respectively. When the system is in the state $0,1$, it will decay back into $|0,0\rangle$ by the dissipation of one photon through the resonator, see Figure 5.12(a). The reset uses a square bichromatic drive with a duration of 32 μs followed by an idle time of 8 μs to ensure that the

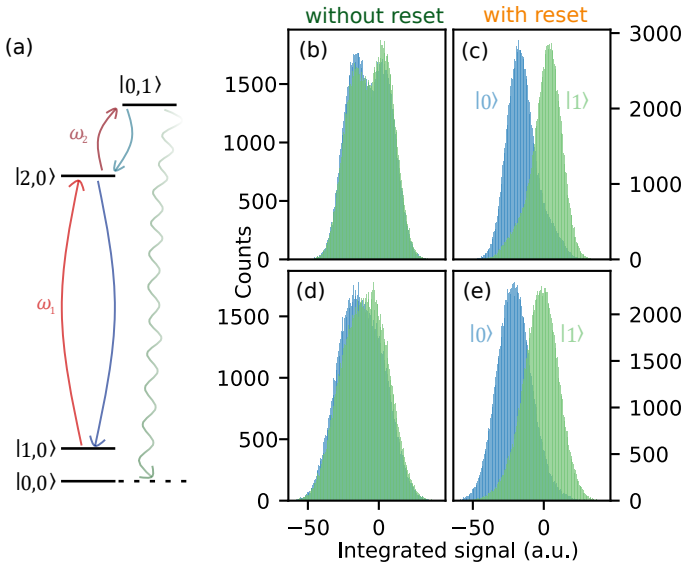

Figure 5.12: (a) Reset scheme used for initializing the qubit in the ground state. The energy levels are labeled with $|Q, R\rangle$ where Q is the qubit state and R is the resonator state shown with simultaneous driving at the angular frequencies ω_1 and ω_2 and with relaxation back into the qubit ground state. (b),(c) Single shot readout histograms for qubit F_1. Panel (b) is the outcome without performing the reset, while (c) is the result after we apply the reset. Similarly panel (d),(e) represent data for qubit F_2.

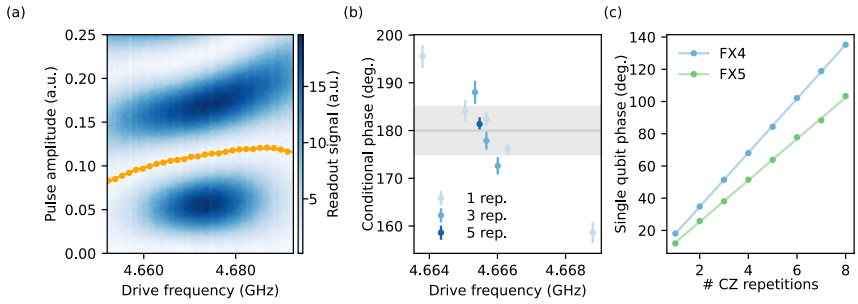

Figure 5.13: (a) Coupler response to varying drive frequency and pulse amplitude for a cosine shaped pulse with a duration of 64 ns. The heatmap corresponds to the magnitude of the difference in the Rabi assisted readout signal when a drive is or is not played, see main text. The yellow markers indicate the 2π rotation amplitude for each drive frequency and the solid yellow line is a spline fit through the markers. (b) Conditional phase vs. drive frequency with varying gate repetitions. The gray band represents a conditional phase error of ± 5 degrees. (c) Single qubit phase errors amplified with increasing number of gate repetitions.

resonator is fully dissipated. When the qubits (F_1 and F_2) are not reset, they thermalize into a mixed state with almost equal populations in 0 and 1. As a result, the single shot readout histogram does not clearly distinguish between the qubit states, see Fig. 5.12(b) and (d). After applying the reset, the qubit is initialized in 0 states and we now clearly see two well-separated histogram peaks in Fig. 5.12(c) and (e). The readout assignment fidelity for the fluxonium qubit 1 and fluxonium qubit 2 is 80% and 77% respectively.

5.13. APPENDIX G: CZ GATE TUNE-UP

In this work, to calibrate the gate, we first measure Rabi oscillations on the coupler when the fluxoniums are prepared in the $|11\rangle$ state, see Fig. 5.13(a) where the pulse duration and pulse shape is kept fixed. From this measurement, we extract the pulse amplitude for each drive frequency that maximally returns the transmon back to its initial state, see orange markers in Fig. 5.13(a) which we fit with a spline function to establish a map from drive frequency to drive amplitude.

Using the extracted drive amplitudes, we determine the conditional phase from the CZ gate as a function of the drive frequency, see Fig 5.13(b). Specifically, we perform a measurement of the conditional phase for 1, 3 and 5 gate repetitions. To efficiently sample the frequency axis, we use a noise aware root finding algorithm to locate the point of π phase. Once the phase error is within ± 5 degrees, we increase the number of gate repetitions and repeat the procedure and we terminate when 5 gate repetitions yield a phase error less than 5 degrees.

Once the pulse amplitude and frequency is calibrated to generate the conditional phase gate, we calibrate the residual single qubit phases for both fluxonium qubits. The single qubit phases are measured by amplifying the single qubit phase errors with repeated applications of the CZ gate, see Fig. 5.13(c). We linearly fit the accumulated phase data and use the slope as a measure of the acquired phase per gate. The unwanted single qubit phases are compensated by using virtual-Z gates [62].

5.14. APPENDIX H: FAST-DRAG CALCULATIONS

To eliminate the undesired frequency components in the in-phase pulse envelope (Ω_I) of the drive used for the gate, we implemented the FAST-DRAG method [53] to derive analytical solutions for the shape of the pulse. This approach works by attenuating the frequency amplitude across n frequency spans (f_n) each spanning a range from $f_{n,l}$ to $f_{n,h}$. Additionally, the amplitude of each frequency span is suppressed by an amount proportional to a specified weight (w_n) assigned to that span. The method works by defining the pulse envelope, Ω_I, as a cosine

series as described in Eq. (5.9)

$$\Omega_I = A \sum_{n=1}^{N} c_n g_n(t), \tag{5.11}$$

where $g_n = (1 - \cos \frac{2\pi n t}{t_g})$ is cosine pulse with n period for the duration of t_g. The pulse amplitude is given by A and each cosine component is weighted with a coefficient c_n. To solve for the coefficients c_n such that the combined pulse shape minimizes the frequency components at the undesired frequencies, we choose $N = 3$ basis functions and solve the following quadratic optimization problem

$$\text{minimize} \sum_{n=1}^{3} w_n \int_{f_{n,l}}^{f_{n,h}} |\hat{\Omega}_I(f)|^2 \, df, \tag{5.12}$$

$$s.t \sum_{n=1}^{3} c_n t_g = \pi, \tag{5.13}$$

where $\hat{\Omega}_I(f)$ is the Fourier transform of the pulse envelope. The condition of π comes from the desired angle of rotation for the gate. To provide an experimental basis to the mathematical relations, we define the center frequencies of each frequency span to be the conditional transitions of the coupler $\{f_{00}, f_{01}, f_{10}\}$ not corresponding to the $|11\rangle$ state of the fluxonioums. Additionally, we assign the weights of each spurious frequency as the square of the Rabi rate of the coupler, corresponding to the transition, see Fig. 5.4. Further, we use a frequency span of 5 MHz centered at each transition.

To find the appropriate minimum as given in Eq. (5.12), we first consider the Fourier transform of basis function given by

$$\begin{aligned} \hat{g}_n = t_g (& e^{-i\pi t_g f} \operatorname{sinc}(\pi t_g f) \\ & - \frac{1}{2} e^{i\pi(n/t_g - f)t_g} \operatorname{sinc}(\pi(n/t_g - f)t_g) \\ & + \frac{1}{2} e^{i\pi(n/t_g + f)t_g} \operatorname{sinc}(\pi(n/t_g + f)t_g) \end{aligned} \tag{5.14}$$

such that we now need to minimize

$$\sum_{n=1}^{3} w_n \int_{f_{l,n}}^{f_{h,n}} \Big| \sum_{i=1}^{3} c_n \hat{g}_n(f) \Big|^2. \tag{5.15}$$

We can re-write Eq. (5.15) as

$$\sum_{n,m} c_n c_m A_{nm} = c^T A c \tag{5.16}$$

where $A_{nm} = \sum_{i=1}^{3} w_i \int_{f_{i,l}}^{f_{i,h}} \hat{g}_n(f)\hat{g}_n^*(f)df$ and c is a vector containing the coefficients c_n. The minimization problem now become equivalent to solving the matrix equation

$$\begin{bmatrix} A+A^T & -b \\ b^T & 0 \end{bmatrix} \times (c^T, \mu)^T = (0^T, \pi)^T \tag{5.17}$$

where b = $(1,.....1)^T \in \Re^{N\times 1}$, where we have introduced the Lagrange multiplier μ. The pulse shapes presented in the main text are found using the solutions to this matrix equation.

REFERENCES

[1] Trond I Andersen et al. "Thermalization and criticality on an analog-digital quantum simulator". In: *arXiv* (2024). DOI: 10.48550/arXiv.2405.17385.

[2] J. I. Colless et al. "Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm". In: *Phys. Rev. X* 8 (1 Feb. 2018), p. 011021. DOI: 10.1103/PhysRevX.8.011021. URL: https://link.aps.org/doi/10.1103/PhysRevX.8.011021.

[3] Alexander Miessen et al. "Benchmarking Digital Quantum Simulations Above Hundreds of Qubits Using Quantum Critical Dynamics". In: *PRX Quantum* 5 (4 Nov. 2024), p. 040320. DOI: 10.1103/PRXQuantum.5.040320. URL: https://link.aps.org/doi/10.1103/PRXQuantum.5.040320.

[4] Yu Zhang et al. "Exploring parity magnetic effects through quantum simulation with superconducting qubits". In: *Phys. Rev. Appl.* 21 (3 Mar. 2024), p. 034052. DOI: 10.1103/PhysRevApplied.21.034052. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.21.034052.

[5] Eunjong Kim et al. "Quantum Electrodynamics in a Topological Waveguide". In: *Phys. Rev. X* 11 (1 Jan. 2021), p. 011015. DOI: 10.1103/PhysRevX.11.011015. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.011015.

[6] Tuure Orell et al. "Collective bosonic effects in an array of transmon devices". In: *Phys. Rev. A* 105 (6 June 2022), p. 063701. DOI: 10.1103/PhysRevA.105.063701. URL: https://link.aps.org/doi/10.1103/PhysRevA.105.063701.

[7] Ilan T Rosen et al. "Flat-band (de) localization emulated with a superconducting qubit array". In: *arXiv* (2024). DOI: 10.48550/arXiv.2410.07878.

[8] Ilan T. Rosen et al. “A synthetic magnetic vector potential in a 2D superconducting qubit array”. In: *Nature Physics* 20.12 (Dec. 2024), pp. 1881–1887. DOI: 10.1038/s41567-024-02661-3.

[9] Rajeev Acharya et al. “Quantum error correction below the surface code threshold”. In: *Nature* 638.8052 (2025), pp. 920–926. DOI: 10.1038/s41586-024-08449-y.

[10] Sebastian Krinner et al. “Realizing repeated quantum error correction in a distance-three surface code”. In: *Nature* 605.7911 (May 2022), pp. 669–674. DOI: 10.1038/s41586-022-04566-8.

[11] J. F. Marques et al. “Logical-qubit operations in an error-detecting surface code”. In: *Nature Physics* 18.1 (Jan. 2022), pp. 80–86. DOI: 10.1038/s41567-021-01423-9.

[12] Harald Putterman et al. “Hardware-efficient quantum error correction via concatenated bosonic qubits”. In: *Nature* 638.8052 (2025), pp. 927–932. DOI: 10.1038/s41586-025-08642-7.

[13] Jens Koch et al. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Phys. Rev. A* 76 (4 Oct. 2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319. URL: https://link.aps.org/doi/10.1103/PhysRevA.76.042319.

[14] Rami Barends et al. “Diabatic Gates for Frequency-Tunable Superconducting Qubits”. In: *Physical Review Letters* 123.21 (2019), p. 210501. DOI: 10.1103/PhysRevLett.123.210501.

[15] L. DiCarlo et al. “Demonstration of two-qubit algorithms with a superconducting quantum processor”. In: *Nature* 460.7252 (July 2009), pp. 240–244. DOI: 10.1038/nature08121.

[16] S. A. Caldwell et al. “Parametrically Activated Entangling Gates Using Transmon Qubits”. In: *Phys. Rev. Appl.* 10 (3 Sept. 2018), p. 034050. DOI: 10.1103/PhysRevApplied.10.034050. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.10.034050.

[17] Jerry M Chow et al. “Microwave-activated conditional-phase gate for superconducting qubits”. In: *New Journal of Physics* 15.11 (Nov. 2013), p. 115012. DOI: 10.1088/1367-2630/15/11/115012. URL: https://dx.doi.org/10.1088/1367-2630/15/11/115012.

[18] Jerry M. Chow et al. “Simple all-microwave entangling gate for fixed-frequency superconducting qubits”. In: *Physical Review Letters* 107.8 (2011), p. 080502. DOI: 10.1103/PhysRevLett.107.080502.

[19] Bradley K. Mitchell et al. “Hardware-Efficient Microwave-Activated Tunable Coupling between Superconducting Qubits”. In: *Physical Review Letters* 127.20 (2021), p. 200502. DOI: 10.1103/PhysRevLett.127.200502.

[20] Sarah Sheldon et al. “Procedure for systematically tuning up cross-talk in the cross-resonance gate”. In: *Physical Review A* 93.6 (2016), p. 060302. DOI: 10.1103/PhysRevA.93.060302.

[21] Chad Rigetti and Michel Devoret. “Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies”. In: *Phys. Rev. B* 81 (13 Apr. 2010), p. 134507. DOI: 10.1103/PhysRevB.81.134507. URL: https://link.aps.org/doi/10.1103/PhysRevB.81.134507.

[22] Sebastian Krinner et al. “Benchmarking coherent errors in controlled-phase gates due to spectator qubits”. In: *Physical Review Applied* 14.2 (2020), p. 024042. DOI: 10.1103/PhysRevApplied.14.024042.

[23] Neereja Sundaresan et al. “Reducing unitary and spectator errors in cross resonance with optimized rotary echoes”. In: *PRX Quantum* 1.2 (2020), p. 020318. DOI: 10.1103/PRXQuantum.1.020318.

[24] Michele C Collodo et al. “Implementation of conditional phase gates based on tunable zz interactions”. In: *Physical Review Letters* 125.24 (2020), p. 240502. DOI: 10.1103/PhysRevLett.125.240502.

[25] Yangsen Ye et al. “Realization of high-fidelity controlled-phase gates in extensible superconducting qubits design with a tunable coupler”. In: *Chinese Physics Letters* 38.10 (2021), p. 100301. DOI: 10.1088/0256-307X/38/10/100301.

[26] Fabian Marxer et al. “Long-Distance Transmon Coupler with cz-Gate Fidelity above 99.8%”. In: *PRX Quantum* 4.1 (2023), p. 010314. DOI: 10.1103/PRXQuantum.4.010314.

[27] Ying Li et al. “Tunable Coupling Scheme for Implementing High-Fidelity Two-Qubit Gates”. In: *Physical Review Applied* 10.5 (2018), p. 054062. DOI: 10.1103/PhysRevApplied.10.054062.

[28] Youngkyu Xu et al. “Realization of High-Fidelity CZ and ZZ-Free iSWAP Gates with a Tunable Coupler”. In: *Physical Review X* 11.2 (2021), p. 021058. DOI: 10.1103/PhysRevX.11.021058.

[29] Eyob A Sete et al. “Parametric-resonance entangling gates with a tunable coupler”. In: *Physical Review Applied* 16.2 (2021), p. 024050. DOI: 10.1103/PhysRevApplied.16.024050.

[30] Frank Arute et al. “Quantum Supremacy Using a Programmable Superconducting Processor”. In: *Nature* 574 (2019), pp. 505–510. DOI: 10.1038/s41586-019-1666-5.

[31] Kevin C Miao et al. “Overcoming leakage in quantum error correction”. In: *Nature Physics* 19.12 (2023), pp. 1780–1786. DOI: 10.1038/s41567-023-02264-1.

[32] J. F. Marques et al. “All-Microwave Leakage Reduction Units for Quantum Error Correction with Superconducting Transmon Qubits”. In: *Phys. Rev. Lett.* 130 (25 June 2023), p. 250602. DOI: 10.1103/PhysRevLett.130.250602. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.250602.

[33] Max Werninghaus et al. “Leakage reduction in fast superconducting qubit gates via optimal control”. In: *npj Quantum Information* 7.1 (2021), p. 14. DOI: 10.1038/s41534-020-00346-2.

[34] Kentaro Heya et al. “Floquet analysis of frequency collisions”. In: *Phys. Rev. Appl.* 21 (2 Feb. 2024), p. 024035. DOI: 10.1103/PhysRevApplied.21.024035. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.21.024035.

[35] Alexis Morvan et al. “Optimizing frequency allocation for fixed-frequency superconducting quantum processors”. In: *Physical Review Research* 4.2 (2022), p. 023079. DOI: 10.1103/PhysRevResearch.4.023079.

[36] Suhas Ganjam et al. “Surpassing millisecond coherence in on-chip superconducting quantum memories by optimizing materials and circuit design”. In: *Nature Communications* 15 (2024), p. 3687. DOI: 10.1038/s41467-024-47857-6. URL: https://www.nature.com/articles/s41467-024-47857-6.

[37] Alexander P. Read et al. “Precision Measurement of the Microwave Dielectric Loss of Sapphire in the Quantum Regime with Parts-per-Billion Sensitivity”. In: *Physical Review Applied* 19.3 (Mar. 2023), p. 034064. DOI: 10.1103/PhysRevApplied.19.034064.

[38] Vladimir E Manucharyan et al. “Fluxonium: Single cooper-pair circuit free of charge offsets”. In: *Science* 326.5949 (2009), pp. 113–116. DOI: 10.1126/science.1175552.

[39] Long B Nguyen et al. “High-coherence fluxonium qubit”. In: *Physical Review X* 9.4 (2019), p. 041041. DOI: 10.1103/PhysRevX.9.041041.

[40] Aaron Somoroff et al. “Millisecond coherence in a superconducting qubit”. In: *Physical Review Letters* 130.26 (2023), p. 267001. DOI: 10.1103/PhysRevLett.130.267001.

[41] Fei Wang et al. “High-coherence fluxonium qubits manufactured with a wafer-scale-uniformity process”. In: *Phys. Rev. Appl.* 23 (4 Apr. 2025), p. 044064. DOI: 10.1103/PhysRevApplied.23.044064. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.23.044064.

[42] Nathan Earnest et al. “Realization of a Λ system with metastable states of a capacitively shunted fluxonium”. In: *Physical Review Letters* 120.15 (2018), p. 150504. DOI: 10.1103/PhysRevLett.120.150504.

[43] Martijn F. S. Zwanenburg et al. “Single-qubit gates beyond the rotating-wave approximation for strongly anharmonic low-frequency qubits”. In: *Physical Review Research* 7 (2025), p. 043290. DOI: 10.1103/z62h-kcnh. arXiv: 2503.08238.

[44] Leon Ding et al. “High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler”. In: *Physical Review X* 13.3 (2023), p. 031035. DOI: 10.1103/PhysRevX.13.031035.

[45] Quentin Ficheux et al. “Fast Logic with Slow Qubits: Microwave-Activated Controlled-Z Gate on Low-Frequency Fluxoniums”. In: *Phys. Rev. X* 11 (2 May 2021), p. 021026. DOI: 10.1103/PhysRevX.11.021026. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.021026.

[46] Ebru Dogan et al. “Two-fluxonium cross-resonance gate”. In: *Physical Review Applied* 20.2 (2023), p. 024011. DOI: 10.1103/PhysRevApplied.20.024011.

[47] Konstantin N. Nesterov et al. “Microwave-activated controlled-Z gate for fixed-frequency fluxonium qubits”. In: *Phys. Rev. A* 98 (3 Sept. 2018), p. 030301. DOI: 10.1103/PhysRevA.98.030301. URL: https://link.aps.org/doi/10.1103/PhysRevA.98.030301.

[48] Helin Zhang et al. “Tunable Inductive Coupler for High-Fidelity Gates Between Fluxonium Qubits”. In: *PRX Quantum* 5.2 (2024), p. 020326. DOI: 10.1103/PRXQuantum.5.020326.

[49] Xizheng Ma et al. “Native approach to controlled-Z gates in inductively coupled fluxonium qubits”. In: *Physical Review Letters* 132.6 (2024), p. 060602. DOI: 10.1103/PhysRevLett.132.060602.

[50] Ilya A. Simakov et al. “Coupler microwave-activated controlled-phase gate on fluxonium qubits”. In: *PRX Quantum* 4.4 (2023), p. 040321. DOI: 10.1103/PRXQuantum.4.040321.

[51] Ilya N. Moskalenko et al. “High fidelity two-qubit gates on fluxoniums using a tunable coupler”. In: *npj Quantum Information* 8.1 (2022), p. 130. DOI: 10.1038/s41534-022-00644-x.

[52] Felix Motzoi et al. “Simple pulses for elimination of leakage in weakly nonlinear qubits”. In: *Physical Review Letters* 103.11 (2009), p. 110501. DOI: 10.1103/PhysRevLett.103.110501.

[53] Eric Hyyppä et al. “Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes”. In: *PRX Quantum* 5.3 (2024), p. 030353. DOI: 10.1103/PRXQuantum.5.030353.

[54] Austin G. Fowler et al. “Surface codes: Towards practical large-scale quantum computation”. In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324. URL: https://link.aps.org/doi/10.1103/PhysRevA.86.032324.

[55] R. Versluis et al. “Scalable Quantum Circuit and Control for a Superconducting Surface Code”. In: *Phys. Rev. Appl.* 8 (3 Sept. 2017), p. 034021. DOI: 10.1103/PhysRevApplied.8.034021. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.8.034021.

[56] Christian Kraglund Andersen et al. “Repeated quantum error detection in a surface code”. In: *Nature Physics* 16.8 (2020), pp. 875–880. DOI: 10.1038/s41567-020-0920-y.

[57] Google Quantum AI. “Suppressing quantum errors by scaling a surface code logical qubit”. In: *Nature* 614.7949 (2023), pp. 676–681. DOI: 10.1038/s41586-022-05434-1.

[58] Figen Yilmaz et al. “Energy-participation-ratio analysis for very anharmonic superconducting circuits”. In: *Physical Review Applied* (2025). DOI: 10.1103/1rbn-c4xf. arXiv: 2411.15039.

[59] Zlatko K. Minev et al. “Energy-participation quantization of Josephson circuits”. In: *npj Quantum Information* 7.1 (Aug. 2021), p. 131. DOI: 10.1038/s41534-021-00461-8.

[60] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. “Characterizing quantum gates via randomized benchmarking”. In: *Phys. Rev. A* 85 (4 Apr. 2012), p. 042311. DOI: 10.1103/PhysRevA.85.042311. URL: https://link.aps.org/doi/10.1103/PhysRevA.85.042311.

[61] Ana Silva and Eliska Greplova. “Hands-on Introduction to Randomized Benchmarking”. In: *SciPost Phys. Lect. Notes* 97 (2025). DOI: 10.21468/SciPostPhysLectNotes.97.

[62] David C. McKay et al. “Efficient Z gates for quantum computing”. In: *Phys. Rev. A* 96 (2 Aug. 2017), p. 022330. DOI: 10.1103/PhysRevA.96.022330. URL: https://link.aps.org/doi/10.1103/PhysRevA.96.022330.

[63] J. Kelly et al. “Optimal Quantum Control Using Randomized Benchmarking”. In: *Phys. Rev. Lett.* 112 (24 June 2014), p. 240504. DOI: 10.1103/PhysRevLett.112.240504. URL: https://link.aps.org/doi/10.1103/PhysRevLett.112.240504.

[64] Daniel Puzzuoli et al. “Algorithms for perturbative analysis and simulation of quantum dynamics”. In: *Journal of Computational Physics* 489 (2023), p. 112262. ISSN: 0021-9991. DOI: https://doi.org/10.1016/j.jcp.2023.112262. URL: https://www.sciencedirect.com/science/article/pii/S0021999123003571.

[65] https://doi.org/10.4121/cec2d36b-c31b-4c70-a66f-65d3ef7bd514. 2025.

[66] https://github.com/AndersenQubitLab/fast_two_qubit_gate. 2025.

[67] Peter Groszkowski and Jens Koch. *scqubits Documentation*. https://scqubits.readthedocs.io/en/latest. 2021.

[68] Taryn V. Stefanski et al. “Improved fluxonium readout through dynamic flux pulsing”. In: *APL Quantum* 2.2 (2025), p. 026103. DOI: 10.1063/5.0255892.

[69] A. Bruno et al. “Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates”. In: *Applied Physics Letters* 106.18 (2015), p. 182601. DOI: 10.1063/1.4919761.

[70] M. Virginia P. Altoé et al. “Localization and mitigation of loss in niobium superconducting circuits”. In: *PRX Quantum* 3.2 (2022), p. 020312. DOI: 10.1103/PRXQuantum.3.020312.

[71] Nandini Muthusubramanian et al. “Wafer-scale uniformity of Dolan-bridge and bridgeless Manhattan-style Josephson junctions for superconducting quantum processors”. In: *Quantum Science and Technology* 9.2 (2024), p. 025006. DOI: 10.1088/2058-9565/ad1c3f.

[72] Dante Colao Zanuz et al. “Mitigating Losses of Superconducting Qubits Strongly Coupled to Defect Modes”. In: *Phys. Rev. Applied* 23 (2025), p. 044054. DOI: 10.1103/PhysRevApplied.23.044054.

[73] A. Dunsworth et al. “Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits”. In: *Applied Physics Letters* 111.2 (2017), p. 022601. DOI: 10.1063/1.4993577.

[74] Chi Zhang et al. “Characterization of tunable coupler without a dedicated readout resonator in superconducting circuits”. In: *Applied Physics Letters* 122.2 (2023), p. 022601. DOI: 10.1063/5.0133900.

[75] Helin Zhang et al. “Universal Fast-Flux Control of a Coherent, Low-Frequency Qubit”. In: *Phys. Rev. X* 11 (1 Jan. 2021), p. 011010. DOI: 10.1103/PhysRevX.11.011010. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.011010.

6

COS2ϕ QUBITS

"Never let a passion for the perfect take precedence over pragmatism."

Walter Isaacson

This chapter covers the fundamentals of the cos2ϕ qubits and thier building blocks. We explored the construction of such qubits using two different types of junction material stack: SIS and SNS junctions respectively. Demonstration of protection was shown with the SIS approach, leading to an improvement of one order of magnitude of lifetimes by using it in the desired regime.

Parts of this chapter are in preparation:"Coherence limitations of a Fourier-engineered cos(2phi) transmon qubit" - N. Zhurbina, S. Singh, L. Splitthoff, A.M. Bozkurt, E.Y. Huang, F. Yilmaz and C.K. Andersen.

6.1. INTRODUCITON

In the last two decades, the superconducting quantum circuits community has focused on the development and refinement of different qubits, especially more effort has been put to refine transmon-style qubits. This sustained period of research has yielded significant advancements in device geometry and material science, resulting in several orders of magnitude improvement in coherence times. However, while optimizing external factors such as dielectric loss and surface treatment have been highly successful, one parallel approach to high-fidelity qubits: the implementation of hardware-level intrinsic protection directly within the system Hamiltonian, needs more research. Unlike traditional transmons, which rely on a large shunt capacitance to suppress charge noise, the aim of protected qubits is to engineer a Hamiltonian that is fundamentally desensitized to environmental noise and fluctuations.

In this chapter we explore the realization of such a strategy by manipulating the potential energy landscape of superconducting circuits. We perform this by leveraging a combination of superconducting-normal-superconducting (SNS) junctions in a gatemon device and subsequently a similar potential using superconducting-insulator-superconducting (SIS) junction-based components within the circuit quantum electrodynamics (cQED) framework.

6.2. HAMILTONIAN PROTECTION:

To implement intrinsic noise protection using an appropriate Hamiltonian, one of the simple potentials that can be employed in superconducting qubit circuits, is the replacement of the $\cos(\phi)$ term with a $\cos(2\phi)$ potential in the transmon-like circuit [1]. The resulting Hamiltonian is given by:

$$\hat{H} = 4E_C(\hat{n} - n_g)^2 - E_J \cos(2\hat{\phi}), \tag{6.1}$$

the potential introduces a π periodicity in the phase ϕ. With the physical periodicity of 2π in the phase of the superconducting qubit, this creates two degenerate ground states which are the superpositions of the odd and even cooper pair states. The qubit states thus have a zero overlap in the charge basis, suppressing the tunneling of single cooper pairs. With the suppressed tunneling, the change in parity states becomes difficult, making the qubit intrinsically protected. This type of circuit has been experimentally realized [1, 2].

To realise such a potential in a physical device, it is necessary to consider the underlying physics of Josephson junctions. In a conventional SIS junction, Cooper pair tunnelling occurs through channels of relatively low transparency, such that the current-phase relation is dominated by the first harmonic, proportional to $\cos(\hat{\phi})$. Higher-order harmonic corrections arise from defects and inhomogeneities in the tunnel barrier, but the first harmonic remains the dominant

contribution [3]. Implementing a $\cos(2\hat{\phi})$ potential therefore requires deliberate control over higher harmonics. This can be achieved through two distinct approaches: either by using junctions with intrinsically higher channel transparencies — such as superconductor-normal metal-superconductor (SNS) junctions — or by engineering circuits in which the first harmonic contributions cancel by design. Both approaches are discussed in the following sections.

6.3. MAKING QUBITS USING SNS JUNCTIONS:

The SNS junction [4, 5, 6, 7] in the short junction regime [8] has a nonlinear potential given by

$$V_{\mathrm{SNS}} = -\Delta \sum_i \sqrt{1 - T_i \sin^2(\hat{\phi}/2)}, \tag{6.2}$$

where the sum runs over i transmission channels, each with transmission probability T_i. The transmission T_i can be controlled by the local electrostatic potential; in the case of an InAs nanowire with a voltage-tunable gate, the gate voltage directly influences the total junction potential. To realise a qubit with this type of junction, the nanowire can be integrated into a transmon-like circuit [9, 10, 11], with the junction shunted by a capacitor. The Hamiltonian for such a qubit is

$$\hat{H} = 4E_C \hat{n}^2 - \Delta \sum_i \sqrt{1 - T_i \sin^2(\hat{\phi}/2)}. \tag{6.3}$$

Expanding the potential term using a Taylor series gives

$$V_{\mathrm{SNS}} = E_J \frac{\hat{\phi}^2}{2} - E_J \left(1 - \frac{3\sum T_i^2}{4\sum T_i}\right) \frac{\hat{\phi}^4}{24} + \mathcal{O}(\hat{\phi}^6). \tag{6.4}$$

The potential is visualised in Fig. 6.1(a) for different values of transmission. Substituting this expansion into the Hamiltonian and neglecting higher-order terms gives

$$\hat{H} \approx 4E_C \hat{n}^2 + E_J \frac{\hat{\phi}^2}{2} - E_J \left(1 - \frac{3\sum T_i^2}{4\sum T_i}\right) \frac{\hat{\phi}^4}{24}. \tag{6.5}$$

This has the same form as the transmon Hamiltonian, with a qubit transition frequency $\omega_{01} = \sqrt{8E_C E_J}/\hbar$ and a transmission-dependent anharmonicity [12],

$$\alpha = E_{12} - E_{01} \approx -E_C \left(1 - \frac{3}{4} \frac{\sum T_i^2}{\sum T_i}\right). \tag{6.6}$$

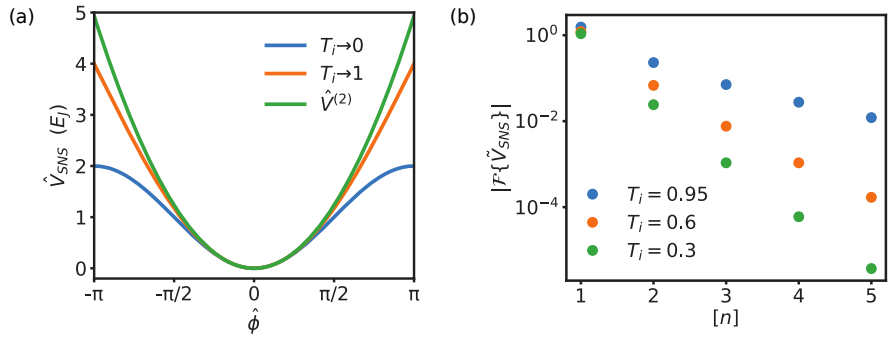

Figure 6.1: Gatemon potential: Figures have been reproduced from [13] (a)SNS potential in different transmission limits (b) Coefficients corresponding to different harmonics are plotted with different trasnmsisions

The anharmonicity is therefore tunable via the gate voltage, which controls the channel transmissions T_i and the device is referred to as a gatemon.

6.3.1. DESIGN OF COS2ϕ DEVICE FROM SNS JUNCTIONS:

The SNS potential in Eq. 6.2 can also be expressed using a Fourier expansion as

$$V_{\text{SNS}} = -\Delta \sum_{n=0}^{\infty} \sum_{i=0}^{\infty} c_{i,n} \cos(n\phi) \tag{6.7}$$

$$\equiv -\Delta \sum_{n=0}^{\infty} C_n \cos(n\phi), \tag{6.8}$$

where $c_{i,n}$ is the coefficient of the n-th harmonic as a function of the channel transparencies, and $C_n = \sum_i c_{i,n}$ is the total coefficient of the n-th harmonic. The coefficients corresponding to different harmonics are plotted in Fig. 6.1(b). The $c_{i,n}$ drop exponentially with increasing n, showing that the higher harmonic coefficients are small compared to the first harmonic.

To create a $\cos(2\hat{\phi})$ potential using SNS junctions, it is necessary to suppress the dominant first harmonic. This can be achieved by combining two junctions in a SQUID loop geometry, as shown in the circuit diagram in Fig. 6.2(e). With an external flux Φ_{ext} threading the loop, the flux quantization condition requires

$$\phi_1 - \phi_2 + 2\pi \frac{\Phi_{\text{ext}}}{\Phi_0} = 2\pi k, \tag{6.9}$$

where ϕ_1 and ϕ_2 are the phase differences across the two SNS junctions respectively. Substituting into Eq. 6.7, the total SQUID potential becomes

$$\hat{V}_{\text{SQUID}} = -\Delta \sum_{n=0}^{\infty} \left(C_n^{(1)} \cos(n\hat{\phi}) + C_n^{(2)} \cos\left(n\left(\hat{\phi} - \frac{2\pi\Phi_{\text{ext}}}{\Phi_0} \right) \right) \right). \tag{6.10}$$

When the external flux is tuned to $\Phi_{\text{ext}} = \Phi_0/2$, the odd harmonics from the two junctions destructively interfere, giving

$$\begin{aligned} \hat{V}\left(\Phi_{\text{ext}} = \frac{\Phi_0}{2} \right) &= -\Delta \sum_{n=0}^{\infty} \left(C_n^{(1)} \cos(n\hat{\phi}) + C_n^{(2)} \cos(n\hat{\phi} - n\pi) \right) \\ &= -\Delta \sum_{n=0}^{\infty} \left(C_n^{(1)} + (-1)^n C_n^{(2)} \right) \cos(n\hat{\phi}). \end{aligned} \tag{6.11}$$

From the harmonic analysis it is clear that $|C_1| > |C_2| > |C_{n\geq 3}|$, with the coefficients diverging further at lower transmission, as seen in Fig. 6.1(b). To achieve a dominant $\cos(2\hat{\phi})$ term, the gate voltages of both junctions must be tuned such that

$$|C_1^{(1)} - C_1^{(2)}|, \quad |C_3^{(1)} - C_3^{(2)}| \ll |C_2^{(1)} + C_2^{(2)}|. \tag{6.12}$$

This requires the corresponding order of harmonic coefficients of the two junctions to be closely matched in magnitude; otherwise, incomplete cancellation of the odd harmonics will prevent the $\cos(2\hat{\phi})$ term from dominating. Once the odd harmonics are suppressed, the potential reduces to

$$\begin{aligned} \hat{V}_{\cos(2\phi)} &\approx -\Delta \left(\left(C_2^{(1)} + C_2^{(2)} \right) \cos(2\hat{\phi}) + \sum_{n=4}^{\infty} \left(C_n^{(1)} + (-1)^n C_n^{(2)} \right) \cos(n\hat{\phi}) \right) \\ &\approx -\Delta \left(C_2^{(1)} + C_2^{(2)} \right) \cos(2\hat{\phi}). \end{aligned} \tag{6.13}$$

PROTECTED $\cos(2\phi)$ HAMILTONIAN

The tuning described above yields the effective Hamiltonian

$$\hat{H}_{\cos(2\phi)} = 4E_C \hat{n}^2 - \Delta C_2^{\Sigma} \cos(2\hat{\phi}), \tag{6.14}$$

where $C_2^{\Sigma} = C_2^{(1)} + C_2^{(2)}$. This assumes $\Delta C_2^{\Sigma} \gg E_C$, analogously to the transmon regime, allowing the offset charge to be neglected. Achieving this condition in practice requires both a large superconducting gap Δ and a large C_2^{Σ}. However, since $C_2/C_1 < 0.1$ in typical devices, the second harmonic coefficient is inherently small, and a larger Δ is needed to compensate. Furthermore, realising a large C_2^{Σ} requires many high-transparency transmission channels, since

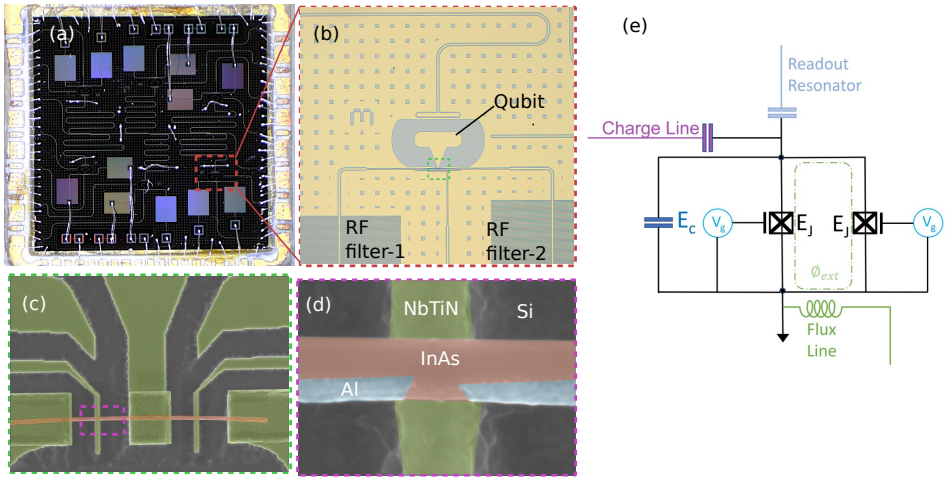

Figure 6.2: Gatemon device. (a) Optical image of the wire-bonded device. (b) Individual qubit pocket with RF lines, flux lines, and readout resonator. (c) Zoomed-in SEM image of the nanowire-based SQUID loop. (d) SEM image of a single SNS junction. (e) Circuit representation of the SNS-based qubit device.

higher harmonics are suppressed at low transparency. The two minima of the $\cos(2\hat{\phi})$ potential form near-degenerate ground states when the charging energy is small. Any deviation from ideal flux tuning introduces residual odd harmonics that break this degeneracy and lift the ground-state protection.

6.4. SNS DEVICE

The qubit layout consists of a single capacitor plate grounded through two SNS junctions, as shown in fig. 6.2(e), with the full device shown in fig. 6.2(b). Each junction is controlled by a dedicated gate line approaching from opposite sides of the qubit. To prevent qubit de-excitation through the gate lines, low-pass RF filters are integrated on each gate line, consisting of an interdigitated capacitor in series with a spiral inductor connected via a wire bond, visible in fig. 6.2(a). A zoomed-in image of the qubit is shown in fig. 6.2(c), with a further close-up of the nanowire junction in fig. 6.2(d).

FABRICATION

The fabrication of the Gatemon devices uses the same superconducting layer of 200 nm thick NbTiN layer and the base layer features are written in a similar process as 2.3.

After the base layer the fabrication steps followed are described in the following table;

Process	Details
PECVD (Gate dielectric)	Clean chamber: CF4 clean (30 min), pump-pre conditioning. **SiN deposition**: 2 min 20 s (28 nm, Recipe: dep. Silicon Nitride 300C HiQual)
Lithography spinning	**Spin primer** AR 300 80 (4000 rpm, 190°C, 2 min). **Spin resist** ARN 7500.18 (4000 rpm, 95°C, 3 min)
Exposure (EBPG 5200)	1500 μC fine field (25 nA beam, 400 μm)
Development	AR 300 47, 2:00 min stirring; H_2O rinse 30 s
Wet etch	BOE bruker 20:1 (3 min); H_2O rinse ×2, dry from H_2O
Strip resist	PRS3000, rinse H_2O, Acetone, IPA
Descum TEPLA	400 W, 200 ml/min O_2
Etching F1	Step 1: 70 W recipe, high SF_6. Step 2: 30 W recipe, higher O_2
Descum TEPLA	400 W, 200 ml/min O_2
Strip resist	NMP
NW Transfer	Nanowire placement with nano-manipulator
Junction Etch	Standard HDMS deposition: prebake 150°C 10 min, HDMS deposition (recipe 3). Spin resist: PMMA 950 A4 (4000 rpm; 185°C; 5 min)
Exposure (EBPG 5200)	1700 μC fine field (25 nA beam, 400 μm)
Development	H_2O:IPA = 1:3, 1 min, stirring
Wet etch	MF322, 50 s; H_2O 5 s; H_2O 30 s; IPA 30 s (no stirring)
NbTiN Contacts	Spin resist: PMMA 950 A4 (4000 rpm; 185°C; 5 min)
Exposure (EBPG 5200)	1900 μC fine field (25 nA beam, 400 μm)
Development	MIBK:IPA = 1:3, 1 min, stirring
Descum TEPLA	400 W, 200 ml/min O_2
Super AJA	Ar milling 120 s (load lock pressure 6×10^{-6} mbar). Pre-sputtering 2 min; Sputtering 4:10 min:s
Lift-off	Acetone, dry from IPA
Dicing	Chip separation

Table 6.1: Details of the individual fabrication steps in the full device fabrication flow.

6.5. RESULTS SNS:

The fabricated gatemon devices were wirebonded using the similar enclosures as ch. 3, and measured. The results of various devices were :

(a) Device fabrication inconsistencies:

Considerable effort was dedicated to ensuring optimal device performance throughout the fabrication process. In the initial stages of nanowire-based gatemon fabrication, defining sub-200 nm gate structures and qubit islands larger than 5 μm required iterative dose testing to reliably reproduce the desired metal features, and multiple rounds of such tests were carried out over the course of the project. Despite this, metal etch residues were consistently observed in the vicinity of small structures. As shown in Fig. 6.3(a), etch residues near the nanowire

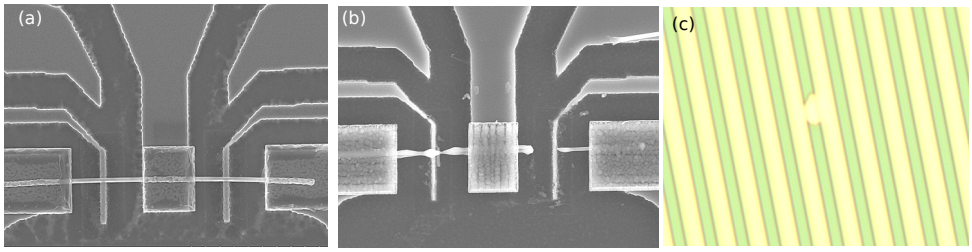

Figure 6.3: Fabrication imperfections. (a), (b) SEM images of the top view of the nanowires. (c) Optical image of the RF filter, zoomed in to show a short circuit.

and gate structures can cause unintended shorts between the gates and the ground plane.

Device yield was further limited by inconsistencies in the dielectric etching step. In particular, the SiN wet-etch process proved difficult to reproduce, frequently yielding unreliable results, as discussed in sec. 2.5.1.

The most critical and yield-limiting fabrication step, however, was the nanowire placement. Nanowires were highly susceptible to electrostatic discharge (ESD), and Fig. 6.3(b) shows an example of a nanowire that was fragmented as a result of an ESD event. Additionally, during the transfer from the nanowire manipulation setup to the cleanroom and during subsequent resist spinning, many nanowires were displaced from their intended positions, drastically reducing the number of functional qubit structures per chip.

For devices that survived these fabrication challenges, further issues were encountered at the measurement stage. Many gate lines were found to be shorted, in part due to insufficient gap widths in the RF filter layout, which happened due to incomplete development of the resist, as shown in Fig. 6.3(c). Wire bonding residues further contributed to shorting of these lines. In several junctions, no measurable tunnel resistance could be obtained, a failure attributed to a combination of factors including thin dielectric layers and residual metallic contamination introduced during wire bonding. The susceptibility of the devices to these failure modes was compounded by the LC filter design, which made it particularly challenging to realise independently tunable junction elements. Achieving two independently controllable junctions within a single SQUID loop proved especially difficult. As discussed further, only one gatemon device functioned successfully during the project — and even in that case, only a single junction remained operational. The measurements from this device are presented below.

(b) Gatemon measurements:

In the one device where the nanowire junction successfully formed a func-

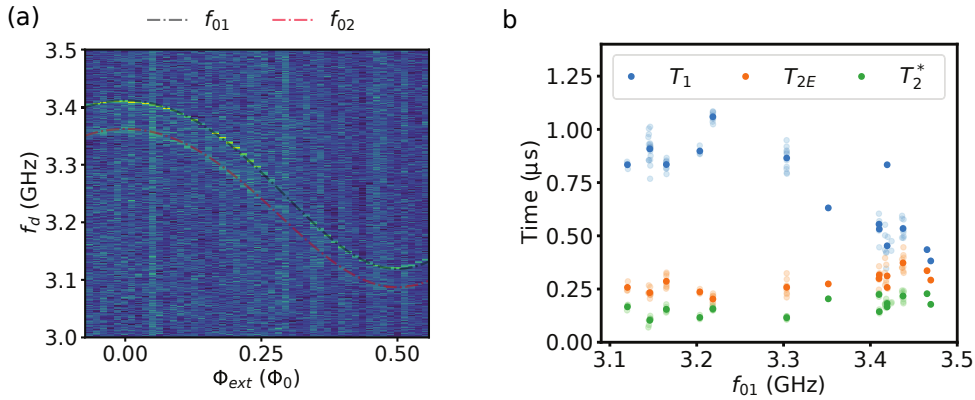

Figure 6.4: Spectroscopy of SNS gatemon qubit: (a) Two-tone spectroscopy of the qubit as a function of the external flux (b) Time domain measurement of the qubit as a function of the qubit frequency.

tional gatemon, the gate voltage on one junction was tunable and allowed control of E_J. The second junction in the SQUID loop did not exhibit any voltage-dependent tunability. The response of the junction to gate voltage was first identified through a shift in the resonator spectroscopy frequency as the gate voltage was swept. Once the resonator power sweep showed nonlinear behaviour indicative of qubit-resonator coupling, two-tone qubit spectroscopy was performed at several gate voltage values. A gate voltage of $V = 3.235\,\text{V}$ was selected for the flux sweep. The measurements confirmed the expected gate voltage and flux tunability, albeit for only one junction in the SQUID loop.

The qubit frequency as a function of the flux threading the loop is shown in Fig. 6.4(a). The data was fitted using a two-junction SNS model, yielding fit parameters of two transmission channels with coefficients $[0.985, 0.984]$ for the first junction, one transmission channel with coefficient $[0.185]$ for the second junction, and an effective superconducting gap of $\Delta = 10.34\,\text{GHz}$. While a reduced effective gap compared to the bulk value is a known feature of SNS junctions [2], the fitted gap value here is unusually low, suggesting that the junction may be highly disordered or that the short-junction approximation underlying the model may not be strictly applicable here. Drawing firm conclusions from these numbers would therefore require additional measurements from multiple devices, which was not possible given the unavailability of further functional devices.

Following the frequency tunability characterisation, time-domain measurements were performed on this single-junction gatemon. The Ramsey coherence time T_2^*, relaxation time T_1, and spin echo time T_{2E} are shown as a function of

qubit frequency in Fig. 6.4(b). Due to the limited dataset from a single device, no conclusive statements can be made regarding the coherence properties. One observation that can be drawn from these measurements is that the effective superconducting gap is significantly lower than the theoretical value of $\Delta = 51$ GHz for aluminium, which is consistent with a highly disordered junction hosting a large number of subgap states.

6.6. MAKING COS2ϕ WITH SIS:

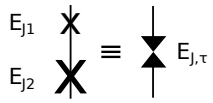

Figure 6.5: Modified SNS junction: Schematic of series combination of two Josephson junctions with Josephson energies E_{J1} and E_{J2} mapped onto an effective non-sinusoidal junction energy $E_{J,\tau}$.

The SIS junction expresses the first harmonic dominantly in its current-phase relation [14]. However, recent theoretical work [15] proposed a scheme to engineer arbitrary energy-phase relations (EPR) using SIS junctions. The key insight is that two SIS junctions connected in series realise an effective SNS junction in the short-junction regime — an equivalence that had not been previously exploited. As shown in Fig. 6.5, two junctions with Josephson energies E_{J1} and E_{J2} in series yield an effective SNS junction with the energy-phase relation [16]

$$E_{\triangleright\triangleleft}(\varphi) = -E_J\sqrt{1-\tau\sin^2(\varphi/2)}, \tag{6.15}$$

where $E_{\triangleright\triangleleft}$ is the energy of the effective SNS junction, and the effective parameters are

$$E_J \equiv E_{J1} + E_{J2}, \qquad \tau \equiv \frac{4E_{J1}E_{J2}}{(E_{J1}+E_{J2})^2}. \tag{6.16}$$

This effective EPR can be further decomposed into a sum of harmonics as

$$E_{\triangleright\triangleleft}(\varphi) = E_J\sum_n c_n(\tau)\cos(n\varphi), \tag{6.17}$$

where $c_n(\tau)$ are the Fourier coefficients of the SNS energy-phase relation, which depend on the effective transparency τ. The coefficients decay exponentially with harmonic order n, with c_1 dominant at low transparency and higher harmonics becoming increasingly accessible as $\tau \to 1$ [16]. We refer to this effective EPR as the series SNS potential hereafter.

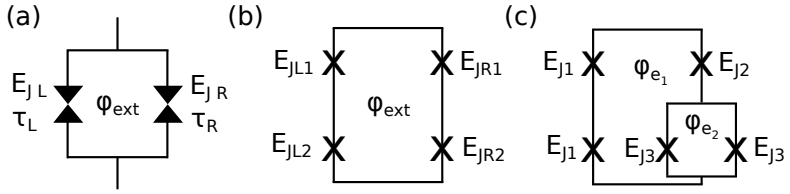

Figure 6.6: Circuit representation: (a) Two series SNS junctions are shown in a SQUID loop geometry with an external flux threading the loop. (b) The implementation of series SNS junction with actual SIS junctions (c) Design of the SIS junctions in the actual device circuit.

Equation 6.15 opens the possibility of realising a charge-insensitive ($E_J/E_C \gg 1$) qubit based on this modified SNS junction, since the effective $E_J = E_{J1} + E_{J2}$ is the sum of the individual SIS junction energies. Because the Josephson energy of an SIS junction can be made arbitrarily large by increasing the junction area, this approach circumvents one of the key limitations of the nanowire SNS-based implementation of the cos(2φ) qubit discussed in Section 6.3.1, where the second harmonic coefficient C_2 is intrinsically small.

DESIGN OF THE SIS COS 2ϕ DEVICE

The series SNS potential, constructed by combining two SIS junctions in series, closely resembles the potential of a single SNS junction, and can be expressed as a Fourier sum of higher harmonics 6.17. An important advantage of this construction is that the transmission coefficient of the junction channel can be tuned by adjusting the E_J values of the constituent SIS junctions, providing an additional degree of freedom. An optimal choice of junction parameters can simultaneously achieve a high E_J/E_C ratio and a large transmission coefficient.

However, the potential of a single series SNS junction is still dominated by the first-order sinusoidal harmonic, and the contribution of higher harmonics remains weak. This is evident from the Fourier coefficients shown in Fig. 6.1(b), where the magnitude of the coefficients decreases exponentially with harmonic order for all values of transmission. Even at transmission close to unity, the ratio of the first to second harmonic coefficient remains approximately 5. The first harmonic therefore cannot be suppressed simply by engineering the transmission of a single series element.

To overcome this limitation, two series SNS junctions are placed in a parallel SQUID configuration, as shown in Fig. 6.6(a), with an external flux Φ_{ext} threading the loop. The flux quantization condition introduces a phase relationship between the two arms, and the flux tunability allows the combined potential to be engineered such that the first harmonics of the two arms cancel, leaving the

second harmonic as the dominant term.

IDEAL CIRCUIT

The combined circuit is shown in Fig. 6.6(a). The left and right arms contain series SNS junctions with total Josephson energies E_{JL} and E_{JR}, and transmissions τ_L and τ_R, respectively. Each series SNS junction is physically realised using two SIS junctions, as shown in Fig. 6.6(b). The left arm comprises junctions E_{JL1} and E_{JL2}, and the right arm comprises E_{JR1} and E_{JR2}, giving total arm energies

$$E_{JL} = E_{JL1} + E_{JL2}, \qquad E_{JR} = E_{JR1} + E_{JR2}, \tag{6.18}$$

and transmissions

$$\tau_L = \frac{4\,E_{JL1}\,E_{JL2}}{(E_{JL1}+E_{JL2})^2}, \qquad \tau_R = \frac{4\,E_{JR1}\,E_{JR2}}{(E_{JR1}+E_{JR2})^2}. \tag{6.19}$$

In the ideal case, when the two junctions within each arm have equal E_J values and both arms are balanced ($E_{JL} = E_{JR}$), the combined potential at $\Phi_{\text{ext}} = 0.5\,\Phi_0$ can be written as

$$U_{\cos 2\varphi} \approx -(E_{JL} - E_{JR})\,c_1 \cos\varphi - (E_{JL} + E_{JR})\,c_2 \cos 2\varphi - \ldots, \tag{6.20}$$

where c_i denotes the Fourier coefficient of the i-th harmonic. When $E_{JL} = E_{JR}$, the first-harmonic term vanishes exactly, and the $\cos(2\varphi)$ term becomes the dominant contribution — realising the desired $\cos(2\varphi)$ potential.

REALISTIC CIRCUIT WITH FABRICATION IMPERFECTIONS

In practice, the four junction energies E_{JL1}, E_{JL2}, E_{JR1}, E_{JR2} cannot be controlled with perfect precision during fabrication, due to the inherent variability in the junction resistance targeting discussed in Section 4.3.1. When the junction energies are unequal, the harmonic coefficients of the two arms differ, and the combined potential at $\Phi_{\text{ext}} = 0.5\,\Phi_0$ takes the more general form

$$\begin{aligned} U(\Phi_{\text{ext}} = 0.5\,\Phi_0) = &-(E_{JL}\,c_{L1} - E_{JR}\,c_{R1})\cos\varphi \\ &-(E_{JL}\,c_{L2} + E_{JR}\,c_{R2})\cos 2\varphi \\ &+(E_{JL}\,c_{L3} - E_{JR}\,c_{R3})\cos 3\varphi + \ldots \end{aligned} \tag{6.21}$$

where c_{Li} and c_{Ri} are the Fourier coefficients of the i-th harmonic for the left and right arm, respectively. The first harmonic is no longer automatically cancelled, and its residual amplitude is proportional to the asymmetry between the two arms.

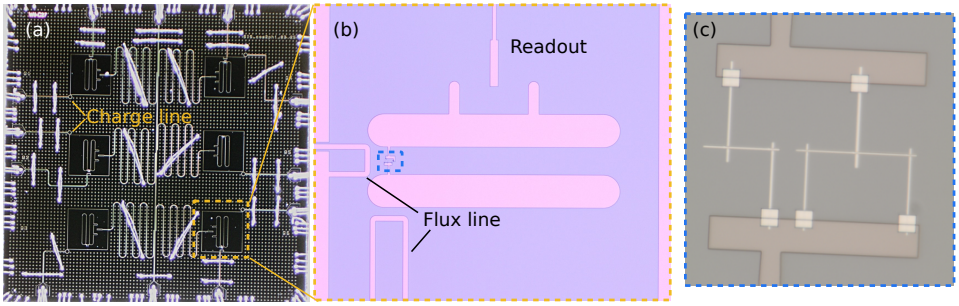

Figure 6.7: SIS device: (a) Microscope image of the wirebonded device is shown with a marking of the charge lines close to the qubit pocket. (b) Individual qubit pocket with marking of the flux line and resonator is imaged with the junctions in the blue inset. (c) The junction geometry is imaged using the optical microscope.

TUNABLE CIRCUIT IMPLEMENTATION

To account for fabrication-induced junction variability and restore the ability to cancel the first harmonic in situ, the circuit is modified to include an additional SQUID loop in the right arm, as shown in Fig. 6.6(c). This renders E_{JR} and consequently τ_R flux-tunable, allowing independent in-situ control of the effective Josephson energy and transmission of the right arm. The cancellation condition for the first harmonic is then

$$E_{JL}\tau_L = E_{JR}\tau_R, \tag{6.22}$$

which can be satisfied over a continuous range by tuning the flux in the additional SQUID loop, even in the presence of junction asymmetry introduced during fabrication.

DEVICE DESIGN AND FABRICATION

With this circuit design, a device was fabricated using the tantalum base-layer process: 200 nm thick Ta on silicon with a Nb seed layer 2.3.2, with junctions fabricated using the Stack B recipe described in Section 4.4. The targeted junction parameters are summarised in Table 6.2.

The device layout consists of six qubit pockets in a transmon-inspired geometry, shown in Fig. 6.7(a) with the charge lines indicated. A zoomed-in view of a single qubit pocket is shown in Fig. 6.7(b), where two floating capacitor pads are visible with the junction circuit between them (blue box). Two dedicated flux lines approach the pocket from opposite sides, enabling simultaneous independent tuning of both flux loops. The physical junction layout implementing the circuit schematic of Fig. 6.6(c) is shown in Fig. 6.7(c).

Table 6.2: Targeted fabrication parameters and resulting ideal effective circuit parameters for the cos(2φ) device.

E_{J1}/h (GHz)	E_{J2}/h (GHz)	E_{J3}/h (GHz)	E_C/h (MHz)
70	70	40	211

E_{JL}/h (GHz)	E_{JR}/h (GHz)	τ_L	τ_R	E_C/h (MHz)
140	0–150	1	0–1	211

6.7. RESULTS: SIS

Once the device was fabricated, it was wire-bonded and measured in the same dilution refrigerator setup as described in Section 3.2. The measurement sequence began with resonator characterisation, followed by flux sweeps of the resonators and a crosstalk calibration of the two flux loops, which is necessary due to the presence of two independent flux lines coupling to the two SQUID loops. After correcting for the crosstalk matrix, qubit spectroscopy was performed and coherence times were measured as a function of flux.

The left column of Fig. 6.8 shows the qubit transition frequency f_{01} as a function of the flux φ_{ext_1} through the larger loop, measured at three different fixed values of the flux φ_{ext_2} through the smaller tunable loop. The corresponding relaxation times T_1 are shown in the right column. Each dataset was fitted using the `scqubits` package [17] to extract the effective right-arm Josephson energy E_{JR} and transmission τ_R at each flux bias point. The extracted parameters are summarised in Table 6.3.

φ_{ext_2}	E_{JR}/h (GHz)	τ_R
0.406	94.37	0.77
0.424	90.5	0.70
0.436	87.1	0.66

Table 6.3: Right-arm junction parameters extracted from `scqubits` fits to the spectroscopy data at three flux bias points of the smaller loop. The left arm has fixed parameters $E_{JL}/h = 140$ GHz and $\tau_L = 1$.

$\varphi_{\text{ext}_2} = 0.406$: At this flux bias, the extracted right-arm parameters are $E_{JR}/h = 94.37$ GHz and $\tau_R = 0.77$, giving a product $E_{JR}\tau_R$ that is close in value to the left arm product $E_{JL}\tau_L$ (with $E_{JL}/h = 140$ GHz and $\tau_L = 1$). This near-cancellation of the first harmonic is visible in Fig. 6.8(a): the qubit frequency drops steeply and reaches a very low minimum at the flux sweet spot ($\varphi_{\text{ext}_1} = 0.5$), characteristic of a potential transitioning from a single-well to a double-well shape as the cos(2φ) term becomes dominant. The frequency spans several GHz over a nar-

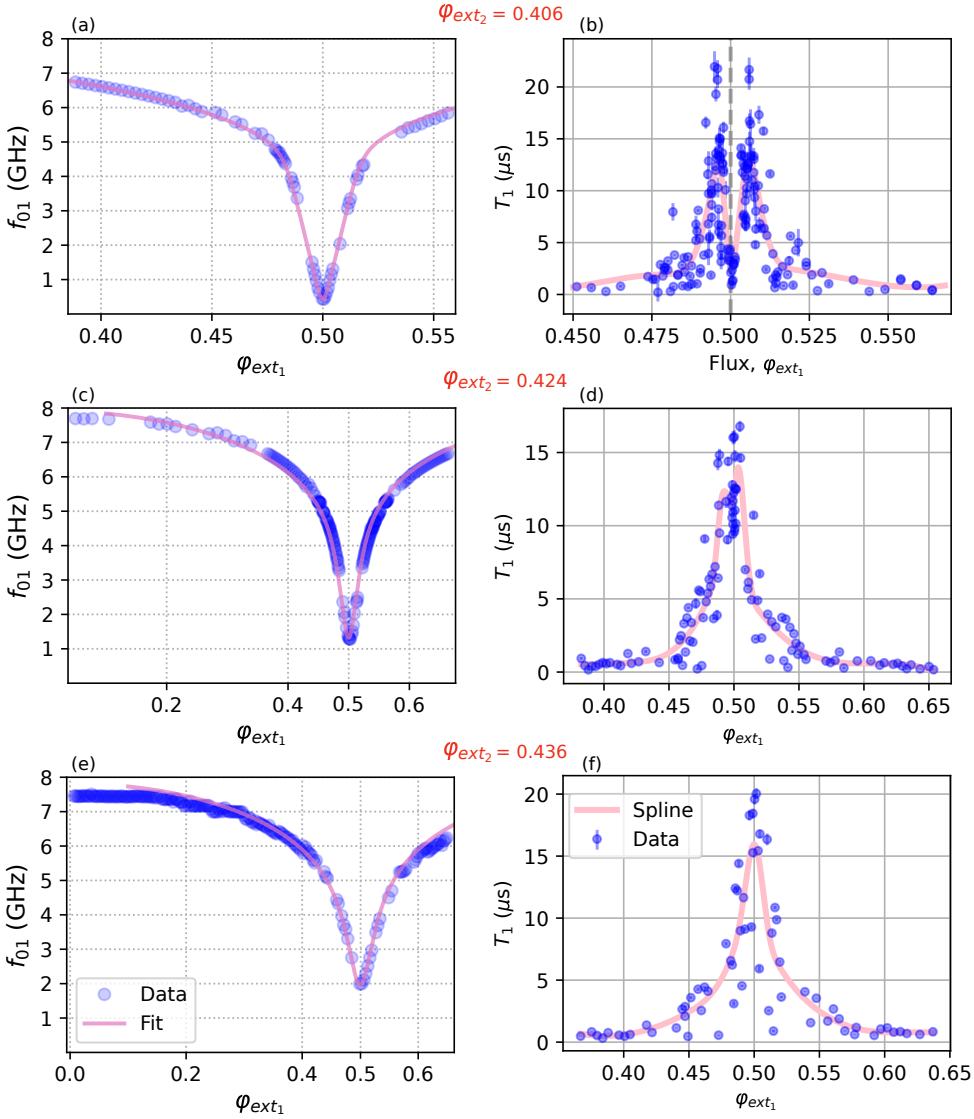

Figure 6.8: Spectroscopy and relaxation time measurements at three bias points of the smaller flux loop (φ_{ext_2}). Left column (a,c,e): qubit transition frequency f_{01} as a function of flux φ_{ext_1} through the larger loop, with the pink line showing the `scqubits` fit. Right column (b,d,f): corresponding T_1 relaxation times as a function of φ_{ext_1}, with a spline overlaid to guide the eye. Each row corresponds to a different value of φ_{ext_2}, decreasing from top to bottom, as the right arm is tuned progressively away from the first-harmonic cancellation condition.

row flux range, consistent with the strongly anharmonic character of the cos(2φ) potential.

The corresponding T_1 data in Fig. 6.8(b), however, reveals a striking feature: while T_1 rises as the flux approaches the sweet spot, it drops sharply precisely at the sweet spot where the qubit frequency is at its lowest. This behaviour suggests that a loss mechanism becomes dominant at the lowest qubit frequencies, making 1/f noise more dominant, though the exact noise channel has not yet been fully identified. This remains an open question under active investigation.

$\varphi_{\mathbf{ext}_2} = 0.424$**:** At this bias point, $E_{JR}/h = 90.5$ GHz and $\tau_R = 0.70$, so the product $E_{JR}\tau_R$ is further from the left arm, meaning the first-harmonic component is less suppressed. As a result, the qubit spectrum in Fig. 6.8(c) shows a shallower frequency dip at the sweet spot, and the overall dispersion starts to resemble that of a transmon-like circuit with a weaker flux dependence near $\varphi_{\mathrm{ext}_1} = 0.5$. Correspondingly, the T_1 data in Fig. 6.8(d) shows generally higher values at the sweet spot compared to the previous bias point, though a residual dip in T_1 at the sweet spot is still visible, suggesting that the same low-frequency loss mechanism continues to play a role.

$\varphi_{\mathbf{ext}_2} = 0.436$**:** At this bias, $E_{JR}/h = 87.1$ GHz and $\tau_R = 0.66$, placing the right arm furthest from the cancellation condition among the three measured points. The qubit spectrum in Fig. 6.8(e) now closely resembles a standard transmon-like spectrum, with a higher minimum frequency at the sweet spot and a broader, smoother flux dependence. The T_1 data in Fig. 6.8(f) shows a monotonically increasing trend as the qubit frequency decreases towards the sweet spot, consistent with dielectric loss, which is known to decrease at lower qubit frequencies [18]. The absence of the sharp T_1 dip at the sweet spot in this regime further supports the interpretation that the anomalous T_1 suppression observed at $\varphi_{\mathrm{ext}_2} = 0.406$ is associated with the low-frequency, double-well character of the potential near the cos(2φ) operating point.

6.8. CONCLUSION & OUTLOOK

This chapter presented the design, fabrication, and characterisation of cos(2φ) qubits pursued through two complementary approaches: nanowire-based SNS junctions and SIS junction-based circuits, both aimed at realising a Hamiltonian with intrinsic first-harmonic suppression.

SNS JUNCTION APPROACH

The SNS-based cos(2φ) device employed two voltage-tunable InAs nanowire junctions in a SQUID loop geometry, exploiting flux quantization to cancel the dominant first harmonic of the combined potential. The design incorporated dedicated low-pass RF filters on each gate line to suppress qubit relaxation through

the voltage control lines. The fabrication process required significant optimisation efforts, including iterative dose testing for sub-200 nm gate features, ESD-protection protocols for the nanowires, and careful integration with the PCB to avoid wire-bonding induced failures.

Despite these efforts, the majority of fabricated devices were non-functional, with failure modes including nanowire ESD damage during transfer and resist spinning, gate line shorts arising from lithographic residues and wire-bonding contamination, and irreproducible SiN wet-etch results. These issues drastically reduced device yield, and only a single gatemon junction — part of a one-junction SQUID — was successfully measured throughout the project.

Spectroscopy of this device confirmed gate voltage and flux tunability of the junction. Fitting the qubit frequency dispersion to an SNS model yielded an effective superconducting gap of $\Delta = 10.34$ GHz, significantly below the theoretical value of 51 GHz for aluminium. While reduced gap values are known to occur in SNS junctions [2], the particularly low fitted value suggests the junction may be highly disordered or that the short-junction limit assumed in the model may not strictly apply. Without a second functional junction in the SQUID loop, first-harmonic cancellation and $\cos(2\varphi)$ operation could not be demonstrated. These findings, combined with the intrinsically small second-harmonic coefficient $C_2/C_1 < 0.1$ of the SNS potential, motivated a change of strategy.

SIS JUNCTION APPROACH

Building on the theoretical proposal of Ref. [15], the approach was changed to using pairs of SIS junctions in series to engineer an effective SNS-like energy-phase relation. This construction offers two key advantages over the nanowire SNS approach: the total effective E_J can be made arbitrarily large by increasing the junction area, enabling operation in the charge-insensitive regime ($E_J/E_C \gg 1$); and the effective transmission τ can be tuned by adjusting the ratio of the constituent SIS junction energies.

A six-qubit device was designed in a transmon-inspired layout, with two series SNS arms placed in a SQUID loop. An additional smaller SQUID loop in the right arm provided in-situ flux tunability of E_{JR} and τ_R, allowing the first-harmonic cancellation condition $E_{JL}\tau_L = E_{JR}\tau_R$ (Eq. 6.22) to be satisfied post-fabrication, compensating for inevitable junction energy variability. The device was fabricated using a tantalum base layer and Stack B junctions.

Spectroscopy and T_1 measurements at three right-arm flux bias points revealed clear signatures of the transition between transmon-like and $\cos(2\varphi)$-like regimes. As $E_{JR}\tau_R$ was tuned towards the left-arm value, the qubit frequency dispersion deepened and narrowed dramatically near the flux sweet spot ($\varphi_{\text{ext}_1} = 0.5$), consistent with the emergence of a double-well $\cos(2\varphi)$ potential. Simul-

taneously, the T_1 behaviour transitioned from a monotonically increasing trend towards the sweet spot — consistent with dielectric loss — to a non-monotonic profile in which T_1 drops sharply at the sweet spot itself. This anomalous T_1 suppression at the lowest qubit frequencies is an active open question, with inductive losses and enhanced $1/f$ noise sensitivity as candidate mechanisms.

OUTLOOK

Several directions remain open for future work on both approaches.

For the SIS $\cos(2\varphi)$ device, the most immediate priority is to understand the loss mechanism responsible for the anomalous T_1 suppression at the sweet spot in the near-cancellation regime. A systematic study of T_1 and T_2 as a function of qubit frequency, flux bias, and device geometry — combined with noise spectroscopy — would help disentangle the contributions of dielectric loss, inductive loss, and $1/f$ flux noise. In parallel, achieving closer matching of the left and right arm parameters through improved junction fabrication targeting (see Chapter 4) would allow the device to be operated more deeply in the $\cos(2\varphi)$ regime, where the protection properties are more pronounced.

For the SNS approach, a redesign of the gate line RF filter to increase the gap widths and reduce shorting susceptibility, combined with improved ESD mitigation during nanowire transfer, could significantly improve device yield. If two functional junctions within a single SQUID loop can be reliably realised, the gate-tunable first-harmonic cancellation scheme becomes experimentally accessible, and the regime of $\cos(2\varphi)$ protection could be directly probed. The anomalously low effective gap observed in the single functional device also warrants further investigation, ideally with a larger dataset from multiple devices, to determine whether it reflects genuine junction disorder or a breakdown of the short-junction approximation.

More broadly, both approaches converge on the same underlying challenge: reliably engineering circuits in which the dominant first harmonic is suppressed while maintaining sufficient E_J/E_C for charge insensitivity and low enough loss rates to permit qubit operation. The SIS approach currently offers the cleaner path to addressing this challenge, and the results presented here provide a first experimental demonstration of the transition into the $\cos(2\varphi)$-dominated potential regime in a fully gate-defined, flux-tunable superconducting circuit.

DATA AVAILABILITY

The data for the experiments along with the scripts to anlayse the data can be found at [19] with an open CC by 4.0 license.

REFERENCES

[1] W. C. Smith et al. “Superconducting circuit protected by two-Cooper-pair tunneling”. In: *npj Quantum Information* 6.1 (2020), p. 8. DOI: 10.1038/s41534-019-0231-2. URL: https://doi.org/10.1038/s41534-019-0231-2.

[2] T. W. Larsen et al. “Parity-Protected Superconductor-Semiconductor Qubit”. In: *Phys. Rev. Lett.* 125 (5 July 2020), p. 056801. DOI: 10.1103/PhysRevLett.125.056801.

[3] P. Willske et al. “Observation of Josephson Harmonics in Tunnel Junctions”. In: *Phys. Rev. Lett.* 131 (21 Nov. 2023), p. 217001. DOI: 10.1103/PhysRevLett.131.217001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.217001.

[4] C. J. Müller, J. M. van Ruitenbeek, and L. J. de Jongh. “Experimental evidence for quantum size effects in metals by contact to a superconductor”. In: *Phys. Rev. Lett.* 69 (1 1992), pp. 140–143. DOI: 10.1103/PhysRevLett.69.140.

[5] H. Takayanagi, T. Akazaki, and J. Nitta. “Observation of Maximum Supercurrent Quantization in a Superconducting Quantum Point Contact”. In: *Phys. Rev. Lett.* 75 (19 1995), pp. 3533–3536. DOI: 10.1103/PhysRevLett.75.3533.

[6] Th. Schäpers et al. “Phase-coherent transport in ballistic Nb/In0.77Ga0.23As/Nb Josephson junctions”. In: *Applied Physics Letters* 70.10 (1997), pp. 1244–1246. DOI: 10.1063/1.119934.

[7] G. E. Blonder, M. Tinkham, and T. M. Klapwijk. “Transition from metallic to tunneling regimes in superconducting microconstrictions: Excess current, charge imbalance, and supercurrent conversion”. In: *Phys. Rev. B* 25 (7 1982), pp. 4515–4532. DOI: 10.1103/PhysRevB.25.4515.

[8] C. W. J. Beenakker. “Universal limit of critical-current fluctuations in mesoscopic Josephson junctions”. In: *Physical Review Letters* 67 (27 1991), pp. 3836–3839. DOI: 10.1103/PhysRevLett.67.3836.

[9] T. W. Larsen et al. “Semiconductor-Nanowire-Based Superconducting Qubit”. In: *Phys. Rev. Lett.* 115 (12 2015), p. 127001. DOI: 10.1103/PhysRevLett.115.127001.

[10] G. de Lange et al. “Realization of Microwave Quantum Circuits from High-Mobility Semiconductor Nanowires”. In: *Phys. Rev. Lett.* 115 (12 2015), p. 127002. DOI: 10.1103/PhysRevLett.115.127002.

[11] B. van Heck, S. Mi, and A. R. Akhmerov. "Single Cooper pair pumping in a semiconductor-nanowire-based transmon". In: *Phys. Rev. B* 90 (15 2014), p. 155450. DOI: 10.1103/PhysRevB.90.155450.

[12] Anders Kringhøj et al. "Anharmonicity of a Gatemon Qubit with a Few-Mode Josephson Junction". In: *Physical Review B* 97 (2018), p. 060508. DOI: 10.1103/PhysRevB.97.060508. arXiv: 1703.05643.

[13] Alexander S. Ivlev. *Developing protected qubits using semicondcuting-superconductin cos(2phi) elements*. Master's thesis, Andersen Lab, QuTech, Delft University of Technology. Unpublished. 2022.

[14] Vinay Ambegaokar and Alexis Baratoff. "Tunneling Between Superconductors". In: *Phys. Rev. Lett.* 10 (11 June 1963), pp. 486–489. DOI: 10.1103/PhysRevLett.10.486. URL: https://link.aps.org/doi/10.1103/PhysRevLett.10.486.

[15] A. Mert Bozkurt et al. "Double-Fourier engineering of Josephson energy-phase relationships applied to diodes". In: *SciPost Phys.* 15 (2023), p. 204. DOI: 10.21468/SciPostPhys.15.5.204. URL: https://scipost.org/10.21468/SciPostPhys.15.5.204.

[16] A. Mert Bozkurt and Valla Fatemi. *Josephson tunnel junction arrays and Andreev weak links: linked by a single energy-phase relation*. 2024. arXiv: 2408.06462 [cond-mat.mes-hall]. URL: https://arxiv.org/abs/2408.06462.

[17] Peter Groszkowski and Jens Koch. "scqubits: a Python package for superconducting qubits". In: *Quantum* 5 (June 2021), p. 483. DOI: 10.22331/q-2021-06-17-483. URL: https://doi.org/10.22331/q-2021-06-17-483.

[18] Clemens Müller, Jared H. Cole, and Jürgen Lisenfeld. "Towards understanding two-level-systems in amorphous solids: insights from quantum circuits". In: *Reports on Progress in Physics* 82.12 (2019), p. 124501. DOI: 10.1088/1361-6633/ab3a7e.

[19] *Data for the thesis titled: Engineering Noise-Resilient Superconducting Qubits: Fabrication and control of fluxonium and cos(2phi) systems*. https://doi.org/45267100-d50e-4ed5-a91f-53e4a899238f. 2026.

7

Conclusion & Outlook

"You can't connect the dots looking forward; you can only connect them looking backwards."

Steve Jobs

This chapter summarises the key findings and contributions of the thesis, drawing together the results from different chapters covered.

Building a better superconducting quantum processor is not a single problem — it is a stack of interconnected problems, each requiring its own careful understanding before the next layer can be addressed. This thesis has followed that approach, building and improving one component at a time. Here is a summary of what was learned at each stage, and where the most promising directions lie for future work.

NANOFABRICATION

At the start of the project, we started with the estalishing nanofabrication processes in the cleanroom starting from substrate selection, cleaning of the substrate, thin-film deposition, lithographic patterning, and etching. The primary superconducting material system pursued was NbTiN on silicon for its superoncducting properties and processing compatibility.

For substrate selection, we also experimented with Sapphire combinations to make resonators. Deposition of NbTiN was established using sputtering process with 200nm thin films as our standard. Tantalum deposition on sapphire was developed in-house using a heated deposition with a Nb seed layer to promote the α-phase. Preliminary SEM evidence suggested α/β phase mixing in some recipes, which needs to be optimised.

The surface pre-treatment protocol was systematically developed, culminating in a sequence comprising solvent cleaning, wet organic clean, and BOE oxide removal immediately before metal deposition. This pre-treatment was found to be critical for reducing interfacial TLS defects at the substrate–metal interface.

For etching, a two-step RIE recipe using SF_6/O_2 was developed for NbTiN on silicon, with the first step achieving high selectivity and the second promoting chemical etching for a smooth substrate surface. A CF_4-based recipe was developed for tantalum on sapphire. Metal island residues were identified as a recurring yield-limiting artefact, attributed to resist residues and loading effects during the RIE process, and were mitigated through improved development sonication, ashing steps. Finally resist residues were removed better with a transition to PRS-3000 as a resist stripper.

RESONATORS

With a working fabrication process in place, the next step was to understand what limits device performance at the level of the basic microwave components. Coplanar waveguide resonators served as the primary diagnostic tool, because they share the same material interfaces and fabrication steps as qubits but are much simpler to fabricate and measure. Improvements in resonator quality factor translate directly into improvements in qubit coherence. Resonators fabri-

cated from NbTiN, tantalum, and niobium were used as experimental proxies to identify and systematically reduce intrinsic loss mechanisms.

Starting from initial NbTiN resonators with quality factors limited to $\sim 10^3$, iterative refinements to the surface treatment protocol yielded an improvement of approximately two orders of magnitude in the internal quality factor Q_i. The dominant loss mechanisms identified were TLS defects at the metal–air (MA) interface, arising from native niobium oxides, and at the substrate–air (SA) interface, arising from residual silicon oxide. The final optimised recipe, consisting of organic cleaning with nitric acid followed by a long 30 min 7:1 BOE dip, was found to reliably remove both oxide species. The best NbTiN resonators on silicon achieved Q_i values approaching 10^6.

Air-bridge integration was explored as a means of suppressing spurious slotline modes in higher-density chip layouts. While functional air bridges were demonstrated, the additional wet processing steps required for their fabrication led to a reduction in Q_i, attributed to surface damage and re-oxidation. A surface cleaning procedure compatible with the presence of aluminium air bridges remains an outstanding challenge.

The transition to tantalum-on-sapphire was motivated by the significantly improved performance of commercially sourced wafers relative to NbTiN. Commercial tantalum-on-sapphire resonators demonstrated Q_i values substantially above those achievable with in-house NbTiN, without requiring extensive process optimisation. Attempts to replicate this performance using in-house-grown tantalum did not match the commercial baseline, with the best in-house results falling below the optimised NbTiN devices. The tantalum-on-silicon material combination was identified as a potentially promising direction warranting further exploration.

Josephson Junctions

When making the junctions we understood the fabrication and characterisation of Al/AlO_x/Al SIS Josephson junctions for fluxonium qubits, covering the physical mechanisms governing junction resistance, the resist stack engineering required to produce clean junctions, and the development of a post-fabrication resistance tuning technique.

An important early finding was the role of substrate photo-conductivity as a systematic shunt resistance, which introduced significant offsets in R_N measurements under ambient lighting and led to E_J targeting errors in early device generations. Once identified and controlled, this artefact was eliminated from the measurement protocol.

Two resist stacks were developed and compared. Stack A, a bilayer MMA/PMMA process, was found to produce inadequate undercut profiles, leaving significant

resist residues on the junction area after lift-off. Stack B, a trilayer MMA/PMMA 495k/PMMA 950k process, introduced an all-around undercut geometry, with improved ashing, sonication-assisted development, and sequential hot NMP lift-off. The combined result was significantly cleaner junctions with improved aluminium grain morphology. The residual limitation of Stack B is a larger junction width offset due to the thicker resist, which increases the coefficient of variation C_V for the smallest junction dimensions. This was partially compensated by using higher oxidation pressure to increase R_N.

The junction array layout was redesigned from the original ladder geometry to a staircase configuration to accommodate the change in stack, while preserving the total array footprint. A further redesign of the array junction geometry was initiated, motivated by the need to simultaneously suppress $1/f$ flux noise and coherent quantum phase slip (CQPS) dephasing. Initial fabrication runs of this new layout have been completed; extensive device-level coherence measurements remain to be performed.

Post-fabrication resistance tuning via voltage-bias annealing was explored as a route to correct E_J targeting errors arising from run-to-run variability in the Plassys deposition system. The resistance asymmetry between the single junction and the array junctions enables selective annealing of the single junction without disturbing E_L. Initial measurements confirmed a measurable resistance shift, providing proof of concept, though systematic characterisation of the temperature, duration, and stability dependence of the annealing effect remains to be completed.

TWO-QUBIT GATES WITH FLUXONIUM QUBITS

With reliable junction fabrication and well-characterised base layer processes, the next step was to put the pieces together into a two-qubit device perform an experimental realisation of fast, high-fidelity two-qubit gates between fluxonium qubits using a transmon coupler. The device comprised two fluxonium qubits (F_1 and F_2) and a transmon coupler (C). The fluxonium qubits operated at frequencies of 98 MHz and 140 MHz respectively, with coherence times of $T_1 = 72.3\,\mu$s and $T_1 = 89.6\,\mu$s.

The gate scheme exploited the state-dependent coupler transition frequencies: since the higher excited states of the fluxonium qubits strongly interact with the transmon coupler, the coupler transition frequency ω_{ij} depends on the computational state $|ij\rangle$ of the two fluxonium qubits. By driving a selective 2π rotation on the ω_{11} transition, a conditional phase (CZ) gate was implemented in which the $|11\rangle$ state acquires a phase of π while all other states are unaffected.

Using a simple cosine drive pulse, a CZ gate fidelity of $98.9\% \pm 0.1\%$ was obtained for a gate duration of 68 ns. Further improvement was achieved by apply-

ing analytical pulse shaping: the DRAG and FAST-DRAG techniques, which suppress coherent errors arising from spectral overlap with neighbouring coupler transitions. The simple analytical form of these pulse shapes provides physical insight into the error-suppression mechanism, in contrast to numerically optimised pulses.

$\cos(2\phi)$ Qubits

We also built the parallel approach to noise protection with pursuing intrinsically noise-protected qubits through Hamiltonian engineering of the $\cos(2\phi)$ potential, pursued via two complementary junction types: SNS junctions and SIS junction-based circuits.

For the SNS approach, two voltage-tunable InAs nanowire junctions were integrated into a SQUID loop geometry, exploiting flux quantization to cancel the dominant first harmonic of the combined potential. The fabrication process required significant optimisation for ESD-protection of the nanowires, gate dielectric definition, and RF filter integration. Despite these efforts, device yield was severely limited by nanowire ESD damage, gate line shorts, and irreproducible SiN wet-etch results. Only a single functional gatemon junction was measured, which showed gate voltage and flux tunability but lacked a second functional junction for first-harmonic cancellation. The fitted effective superconducting gap showed a value far below the theoretical value for aluminium — suggesting a highly disordered junction or a breakdown of the short-junction approximation, though insufficient data from a single device precluded definitive conclusions.

The approach was subsequently changed to SIS-based series junctions, motivated by the theoretical equivalence between two SIS junctions in series and an effective SNS junction in the short-junction limit. This construction allowed the effective E_J to be made arbitrarily large by increasing the junction area, circumventing the intrinsically small $C_2/C_1 < 0.1$ of the SNS potential. A six-qubit device was designed in a transmon-inspired layout, with two series SNS arms in a SQUID loop and an additional smaller SQUID loop in the right arm to provide in-situ flux tunability of E_{JR} and τ_R.

Spectroscopy and T_1 measurements revealed clear signatures of the transition between the transmon-like and $\cos(2\phi)$-like regimes. As the right arm was tuned towards the cancellation condition, the qubit frequency dispersion deepened and narrowed near the flux sweet spot, consistent with the emergence of a double-well $\cos(2\phi)$ potential. Simultaneously, the T_1 behaviour transitioned from a transmon like trend towards the sweet spot to a non-monotonic profile with a sharp T_1 suppression at the sweet spot. This loss at the lowest qubit frequencies, likely inductive in origin or associated with enhanced $1/f$ noise sensitivity in the double-well regime, remains an open question.

OUTLOOK

In **nanofabrication** processes, the systematic development of a process for consistent deposition and subsequent etching of α-phase tantalum on silicon represents an important future direction, given the potential for high-quality tantalum performance on a more process-compatible substrate. More systematic monitoring of compressive film stress during NbTiN deposition, combined with improved RIE chamber conditioning and maintennace, would reduce device-to-device variability. Eliminating HMDS from the pre-treatment flow and replacing BOE with ammonium fluoride solutions are promising routes to further suppress interfacial losses.

For the surface loss optimisation for **resonators**, hydrogen incorporation into niobium films during the extended BOE etch is a suspected loss channel that has not yet been conclusively characterised. Switching to alternative oxide removal chemistries with reduced hydrogen exposure — such as vapour HF or ammonium fluoride — could yield further Q_i improvements. For tantalum, developing an in-house recipe that reliably produces single-phase α-Ta, validated by XRD, is a prerequisite for achieving competitive quality factors. Fabricating NbTiN-compatible air bridges from refractory materials (e.g., NbTiN or Ta) rather than aluminium would decouple air-bridge integration from the surface cleaning constraint.

For Josephson **junctions** the C_V of single junction resistance at 100–150 nm scales remains the primary limiting factor for frequency targeting. A systematic calibration od processes combined with careful tracking of Plassys source condition over its lifetime, would provide the data needed to tighten targeting accuracy.

For the **two qubit gates**, the demonstrated gate fidelity is competitive with the state of the art for fluxonium-based processors and was achieved at gate speeds well below 100 ns. Key directions for improvement include extending the demonstration to multiple qubit pairs on the same chip, and improving the uniformity of junction targeting to reduce qubit frequency spread. Scaling this architecture towards the surface code will require careful management of frequency crowding in larger arrays.

For the **SIS** $\cos(2\phi)$ device, a systematic noise spectroscopy study varying qubit frequency, flux bias, and device geometry is needed to disentangle the competing loss mechanisms near the $\cos(2\phi)$ operating point. The SIS approach currently offers the more practical path forward, and the results presented constitute a first experimental demonstration of the transition into the $\cos(2\phi)$-like potential regime.

ACKNOWLEDGMENTS

This thesis would not have been possible without the immense amount of support I've found through my family, friends, and colleagues at TU Delft. I've shared many conversations and hold deep gratitude for the people mentioned in these acknowledgments.

I'm extremely blessed to have my Mummy and Papa who have instilled in me a deep sense that karma comes first, in the truest sense. **Mummy**, it would not have been possible to keep my sanity through these four years without our everyday 9 am calls, which helped me get my head straight and stay in action mode for the day. **Papa**, I would always end the day tired and call you back. You never seemed tired, even after your long days, and kept me motivated to see what happens when you keep showing up consistently. At the same time, **Bitlu**, thanks for all the weekend conversations, nudging me to not take life too seriously.

I joined QuTech back in 2021, still reeling from Covid, and found my new home in the Andersen lab. **Christian**, I would like to thank you for this wonderful opportunity to pursue a PhD. I was very scared at the start, unsure of my ability to truly understand and perform. I'm blessed to have begun the PhD at the right point in my life. Your guidance and office door have been as open as your approach to science — always open throughout. You have given me the time and resources to make my mistakes in the cleanroom at my own pace. I'll always be indebted to you. Thank you.

The Andersen lab would not have been the same without all the warmth I've received from my E231 café mates. **Figen**, you've been a constant companion from almost the very beginning — from the cleanroom to the wirebonding room. You were determined enough to handle fabrication while simultaneously designing the qubits. Thanks for all the fun conversations and for showing me the spirit of never giving up. **Martijn**, we started in the group at almost the same time, and since then you've always been there for all the troubleshooting — starting early every morning and never getting sick of it. Good luck with your new beginnings in life. **Taryn**, we still haven't reached a consensus on who the funniest person in QuTech is. You literally had my back throughout the four years, and it has been a joy to learn from you. You're always positive, cracking jokes all the while putting in the hard work. **Eugene**, my thesis project would have been partially impossible had it not been for you taking on the measurements. We've been great collaborators and even better travel companions. Let's go to Corsica soon.

Natasha, we spent many long hours inside the cleanroom troubleshooting the nanowires. The cleanroom wasn't always fun, but our conversations made it possible to enjoy the toil. Good luck with writing. **Jinlun**, you started as a master's student when we did wirebonding together, and we've had great learning experiences with the junction fabrication. I hope you're not still running another dose test. Playing tennis has been great — I hope I can join a tournament soon. **YenAn**, it was fun to make the flipmon, especially working on the first 2-inch wafers together. Thanks for all the chocolates, especially after the cleanroom hard work. **Sara**, you were the first and last student during my PhD. You taught me how to become a better teacher. Thank you, and good luck with those long couplers.**Barka, Saqar, Ciara**, we didn't get much overlap during my PhD, I would wish you good luck with your PhDs.

I would also like to thank my mentors in the cleanroom and the lab. **Lukas**, always full of energy and ready to tackle ideas with intensity — you've shown me how to bring that energy to every situation. Thank you. **Marta**, thanks for teaching me how to operate in the cleanroom. You trusted my instincts in there, and it was always great to see how involved you were in all the activities at QuTech. **Jaap**, thanks for all the guidance in the cleanroom during my first year. Our conversations always gave me so much food for thought and offered such thoughtful perspectives. Good luck with the next phase of your life! **Shikhar**, thanks for signing up as my first master's student. I could not have asked for a more hard-working person. You helped me gain real insight inside the cleanroom. Also, thanks to **Alena, Sasha, Siyu, Rohan, Floris, Roald**, and all the members who came along during my PhD — your discussions and contributions helped me grow as a scientist and a teacher.

he Andersen lab is hosted in the Qubit Research division of QuTech. Thank you all for giving me an even broader understanding of quantum physics and different types of qubits. I would like to thank **Christian M., Qing, Ivan, David, Rebecca, Dhananjay, Francesco, Miguel, Praveen, Sameer, and Guan** for the scientific discussions and insights into the dark art of fabrication. **Srijit, Anusua,** and **Michael**, thanks for the open discussions on science and beyond.

Delphine, it was great fun to begin the fabrication work in the cleanroom when you were helping us set up the processes. **Nicos**, I learnt a great deal about cleanroom optimisation from you. **Piran**, you showed such great insight into process control. I would like to thank all of you for teaching me techniques and tools inside the cleanroom.

At QuTech, I would also like to thank the LeoDC group, our superconducting qubit companions. I have learnt a great deal from you about the science of qubits. **Nandini**, you taught me a great deal about cleanroom work and junction fabrication. **Anand**, you've shown me how even the smallest details of resist can

work magic for the junctions. **Matvey**, you gave me so many insights into cleanroom processes. **Sean**, thanks for showing me how to use the automated probing setup. **Santi**, your questions have helped me develop a better understanding of qubit gate operations. **Leo, Kishore, Vishwanath, Ruggero, Tim, Yuejie, Marios**, thanks for all the scientific discussions over the years. You've all refined my understanding of how qubits work.

Jenny, I'm extremely grateful for all the support with administrative tasks I've received from you. You've always welcomed me with a smile in your office. Thank you.

Jason, you've been a treasure trove of insight. From vacuum gauges to mechanical clamps, you've always shared so much information that I've come away with ideas I never knew I needed. **Olaf, Raymond, Hitham**, and everyone else who made it possible to work seamlessly inside QuTech — I've come to realise how important it is to have the right insights and experience to move quickly, and the people at QuTech have provided exactly that. Thank you for your time.

From the Quantum Nanoscience department, I would also like to thank **Robin** for the insights into improving junction fabrication, and **Jana** — it was a pleasure iterating over junction resist together.

I would like to thank the Kavli Nanolab staff for the help they've provided throughout the years in understanding the science and operation of the tools used inside the cleanroom. Your support in troubleshooting issues quickly has been invaluable. Thank you!

I would like to thank my friends. **Mohit**, this PhD would not have been possible without your meals, which always helped me forget what went wrong inside the cleanroom. Thank you for always being so welcoming. **Siddhant**, we started this journey together. The constant support from the very beginning, the conversations, and all the weekend trips we took allowed me to have fun through the years at QuTech. **Pratik**, had it not been for our shared travel plans, I would have gone insane far earlier in my PhD — the fun shall go on. **Onima, Anubhav, Sameer**, it has been great fun these last few years. Finally, I would like to thank the country of **The Netherlands** for providing me with such beautiful opportunities and people in the past four years! Thank You!

CURRICULUM VITÆ

Siddharth SINGH

18-11-1998	Born in Jaipur, India

EDUCATION

2013–2015	High School Step by Step High School, Jaipur, India
2015–2020	Dual Degree with B.E. (Hons.) in Mechanical Engineering & M.Sc. (Hons.) in Physics, Birla Institute of Technology and Science (BITS), Pilani, India
2021–2025	PhD in Applied Physics QuTech, Technische Universiteit Delft *Thesis:* Engineering noise-resilient superconducting qubits: fabrication and control of fluxonium and cos(2phi) systems *Copromotor:* Dr. C. K. Andersen *Promotor:* Prof. dr. L. DiCarlo

List of Publications

4. **S. Singh**, E.Y. Huang, J. Hu, F. Yilmaz, M.F.S. Zwanenburg, P. Kumaravadivel, *et al.*, *Fast microwave-driven two-qubit gates between fluxonium qubits with a transmon coupler*, Physical Review Applied **25**(2), 024020 (2026).

3. F. Yilmaz, **S. Singh**, M.F.S. Zwanenburg, J. Hu, T.V. Stefanski, C.K. Andersen, *Energy participation ratio analysis for very anharmonic superconducting circuits*, Physical Review Applied **25**, 044021 (2026).

2. M.F.S. Zwanenburg, **S. Singh**, E.Y. Huang, F. Yilmaz, T.V. Stefanski, J. Hu, *et al.*, *Single-qubit gates beyond the rotating-wave approximation for strongly anharmonic low-frequency qubits*, Physical Review Research **7**(4), 043290 (2025).

1. T.V. Stefanski, F. Yilmaz, E.Y. Huang, M.F.S. Zwanenburg, **S. Singh**, S. Wang, *et al.*, *Improved fluxonium readout through dynamic flux pulsing*, arXiv:2411.13437 (2024).